Parallel Metric Tree Embedding Based on an Algebraic View on Moore-Bellman-Ford

STEPHAN FRIEDRICHS, Max Planck Institute for Informatics, Germany, and Saarbrücken Graduate School of Computer Science, Germany

CHRISTOPH LENZEN, Max Planck Institute for Informatics, Germany

A *metric tree embedding* of expected stretch $\alpha \geq 1$ maps a weighted n -node graph $G = (V, E, \omega)$ to a weighted tree $T = (V_T, E_T, \omega_T)$ with $V \subseteq V_T$ such that, for all $v, w \in V$, $\text{dist}(v, w, G) \leq \text{dist}(v, w, T)$, and $\mathbb{E}[\text{dist}(v, w, T)] \leq \alpha \text{dist}(v, w, G)$. Such embeddings are highly useful for designing fast approximation algorithms as many hard problems are easy to solve on tree instances. However, to date, the best parallel (polylog n)-depth algorithm that achieves an asymptotically optimal expected stretch of $\alpha \in O(\log n)$ requires $\Omega(n^2)$ work and a metric as input.

In this article, we show how to achieve the same guarantees using polylog n depth and $\tilde{O}(m^{1+\varepsilon})$ work, where $m = |E|$ and $\varepsilon > 0$ is an arbitrarily small constant. Moreover, one may further reduce the work to $\tilde{O}(m + n^{1+\varepsilon})$ at the expense of increasing the expected stretch to $O(\varepsilon^{-1} \log n)$.

Our main tool in deriving these parallel algorithms is an algebraic characterization of a generalization of the classic Moore-Bellman-Ford algorithm. We consider this framework, which subsumes a variety of previous “Moore-Bellman-Ford-like” algorithms, to be of independent interest and discuss it in depth. In our tree embedding algorithm, we leverage it to provide efficient query access to an approximate metric that allows sampling the tree using polylog n depth and $\tilde{O}(m)$ work.

We illustrate the generality and versatility of our techniques by various examples and a number of additional results. Specifically, we (1) improve the state of the art for determining metric tree embeddings in the Congest model, (2) determine a $(1 + \varepsilon)$ -approximate metric regarding the distances in a graph G in polylogarithmic depth and $\tilde{O}(n(m + n^{1+\varepsilon}))$ work, and (3) improve upon the state of the art regarding the k -median and the buy-at-bulk network design problems.

CCS Concepts: • **Theory of computation** → **Shortest paths; Parallel algorithms; Routing and network design problems**; Distributed algorithms;

Additional Key Words and Phrases: MBF-like algorithms, hop sets, PRAM, approximate metric, buy-at-bulk network design, k -median

This work extends and subsumes the extended abstract that appeared in the *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2016)*, pages 455–466, 2016 [25].

Authors’ addresses: S. Friedrichs, Max Planck Institute for Informatics, Saarland Informatics Campus, 66123, Saarbrücken, Germany, Saarbrücken Graduate School of Computer Science, Saarland Informatics Campus, 66123, Saarbrücken, Germany; email: sfriedri@mpi-inf.mpg.de; C. Lenzen, Max Planck Institute for Informatics, Saarland Informatics Campus, 66123, Saarbrücken, Germany; email: clenzen@mpi-inf.mpg.de.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM

0004-5411/2018/11-ART43 \$15.00

<https://doi.org/10.1145/3231591>

ACM Reference format:

Stephan Friedrichs and Christoph Lenzen. 2018. Parallel Metric Tree Embedding Based on an Algebraic View on Moore-Bellman-Ford. *J. ACM* 65, 6, Article 43 (November 2018), 55 pages.
<https://doi.org/10.1145/3231591>

1 INTRODUCTION

In many graph problems the objective is closely related to distances in the graph. Prominent examples are shortest path problems, minimum weight spanning trees, a plethora of Steiner-type problems [27], the traveling salesman, finding a longest simple path, and many more.

If approximation is viable or mandatory, a successful strategy is to approximate the distance structure of the weighted graph G by a simpler graph G' , where “simpler” can mean fewer edges, smaller degrees, being from a specific family of graphs, or any other constraint making the considered problem easier to solve. One then proceeds to solve a related instance of the problem on G' and map the solution back to G , yielding an approximate solution to the original instance. Naturally, this requires a mapping with bounded impact on the objective value.

A standard tool is *metric embedding*, mapping $G = (V, E, \omega)$ to $G' = (V', E', \omega')$, such that $V \subseteq V'$ and $\text{dist}(v, w, G) \leq \text{dist}(v, w, G') \leq \alpha \text{dist}(v, w, G)$ for some $\alpha \geq 1$, referred to as *stretch*.¹ An especially convenient class of metric embeddings are *metric tree embeddings*, plainly because very few problems are hard to solve on tree instances. The utility of tree embeddings originates in the fact that, despite their extremely simple topology, it is possible to randomly construct an embedding of any graph G into a tree T so that the *expected stretch* $\alpha = \max\{\mathbb{E}_T[\text{dist}(v, w, T)] / \text{dist}(v, w, G) \mid v, w \in V\}$ satisfies $\alpha \in O(\log n)$ [23]. By linearity of expectation, this ensures an expected approximation ratio of $O(\log n)$ for most problems; repeating the process $\log(\varepsilon^{-1})$ times and taking the best result, one obtains an $O(\log n)$ -approximation with probability at least $1 - \varepsilon$.

A substantial advantage of tree embeddings lies in the simplicity of applying the machinery once they are computed: Translating the instance on G to one on T , solving the instance on T , and translating the solution back tends to be extremely efficient and highly parallelizable; we demonstrate this in Sections 9 and 10. Note also that the embedding can be computed as a preprocessing step, which is highly useful for online approximation algorithms [23]. Hence, a low-depth small-work parallel algorithm in the vein of Fakcharoenphol, Rao, and Talwar [23] (FRT) gives rise to fast and efficient parallel approximations for a large class of graph problems. Unfortunately, the tradeoff between depth and work achieved by state-of-the-art parallel algorithms for this purpose is suboptimal. Concretely, all algorithms of polylog n depth use $\Omega(n^2)$ work, whereas we are not aware of any stronger lower bound than the trivial $\Omega(m)$ work bound.²

Our Contribution. Our main contribution is to reduce the amount of work for sampling from the FRT distribution—a random distribution of tree embeddings—to $\tilde{O}(m^{1+\varepsilon})$ while maintaining polylog n depth. This article is organized in two parts. The first establishes the required techniques:

- Our key tool is an algebraic interpretation of Moore-Bellman-Ford-like (MBF-like) algorithms described in Section 2. As our framework subsumes a large class of known algorithms and explains them from a different perspective—we demonstrate this using numerous examples in Section 3—we consider it to be of independent interest.

¹ $\text{dist}(\cdot, \cdot, G)$ denotes the distance in G and defines a metric space. See definitions in Section 1.2.

²Partition $V = A \dot{\cup} B$ evenly and add spanning trees of A and B consisting of edges of weight 1. Connect A and B with $m - n + 2$ edges, all of weight $W \gg n \log n$, but with probability of $1/2$, pick one of the connecting edges uniformly at random and set its weight to 1. To approximate the distance between $a \in A$ and $b \in B$ better than factor $W/n \gg \log n$ with probability substantially larger than $1/2$, any algorithm must examine $\Omega(m)$ edges in expectation.

- Section 4 proposes a sampling technique for embedding a graph G in which d -hop distances $(1 + \hat{\varepsilon})$ approximate exact distances into a complete graph H , where H has polylogarithmic Shortest Path Diameter (SPD) and preserves G -distances $(1 + \hat{\varepsilon})^{O(\log n)}$ approximately.
- We devise an oracle that answers MBF-like queries by efficiently simulating an iteration of an MBF-like algorithm on H in Section 5. It uses only the edges of G and polylogarithmic overhead, resulting in $\tilde{O}(dm)$ work with respect to G (i.e., subquadratic work) per iteration; We use $d \in \text{polylog } n$.

The second part applies our techniques and establishes our results:

- A first consequence of our techniques is that we can query the oracle with All-Pairs Shortest Paths (APSP) to determine w.h.p.³ a $(1 + o(1))$ -approximate metric on G using $\tilde{O}(nm^{1+\varepsilon})$ work and $\text{polylog } n$ depth. We discuss this in Section 6.

THEOREM 6.2 ((1 + o(1))-APPROXIMATE METRIC). *Given a weighted graph $G = (V, E, \omega)$ and a constant $\varepsilon > 0$, we can w.h.p. compute, using $\tilde{O}(n(m + n^{1+\varepsilon}))$ work and $\text{polylog } n$ depth, a $(1 + 1/\text{polylog } n)$ -approximate metric of $\text{dist}(\cdot, \cdot, G)$ on V .*

Preprocessing the graph by computing a sparse spanner, one can get close to quadratic work at the expense of a constant approximation ratio.

THEOREM 6.3 (O(1)-APPROXIMATE METRIC). *For a weighted graph $G = (V, E, \omega)$ and a constant $\varepsilon > 0$, we can w.h.p. compute an $O(1)$ -approximate metric of $\text{dist}(\cdot, \cdot, G)$ on V using $\tilde{O}(n^{2+\varepsilon})$ work and $\text{polylog } n$ depth.*

- In Section 7, we show that, for any constant $\varepsilon > 0$, there is a randomized parallel algorithm of depth $\text{polylog } n$ and work $\tilde{O}(m^{1+\varepsilon})$ (w.h.p.) that samples from a metric tree embedding of expected stretch $O(\log n)$. This follows from the preceding techniques and the fact that sampling from the FRT distribution is MBF-like.

COROLLARY 7.10. *Given the weighted incidence list of a graph G and an arbitrary constant $\varepsilon > 0$, we can w.h.p. sample from a tree embedding of expected stretch $O(\log n)$ using depth $\text{polylog } n$ and work $\tilde{O}(m^{1+\varepsilon})$.*

Applying the spanner construction of Baswana and Sen [9] as a preprocessing step, the work can be reduced to $\tilde{O}(m + n^{1+\varepsilon})$ at the expense of stretch $O(\varepsilon^{-1} \log n)$.

COROLLARY 7.11. *Suppose we are given the weighted incidence list of a graph G . Then, for any constant $\varepsilon > 0$ and any $k \in \mathbb{N}$, we can w.h.p. sample from a tree embedding of G of expected stretch $O(k \log n)$ using depth $\text{polylog } n$ and work $\tilde{O}(m + n^{1+1/k+\varepsilon})$.*

- Our techniques allow us to improve over previous distributed algorithms computing tree embeddings in the Congest [41] model. We reduce the best-known round complexity for sampling from a tree embedding of expected stretch $O(\log n)$ from $\tilde{O}(n^{1/2+\varepsilon} + D(G))$, where $\varepsilon > 0$ is an arbitrary constant and $D(G)$ is the unweighted hop diameter of G , to $(n^{1/2} + D(G))n^{o(1)}$. This is detailed in Section 8.

THEOREM 8.1. *There is a randomized distributed algorithm that w.h.p. samples from a metric tree embedding of expected stretch $O(\log n)$ in $\min\{(\sqrt{n} + D(G))n^{o(1)}, \tilde{O}(\text{SPD}(G))\}$ rounds of the Congest model.*

³With high probability, see Section 1.2.

- We illustrate the utility of our main results by providing efficient approximation algorithms for the k -median and buy-at-bulk network design problems. Blleloch et al. [14] devise polylogarithmic depth parallel algorithms based on FRT embeddings for these problems assuming a metric as input. We provide polylogarithmic depth parallel algorithms for the more general case where the metric is given implicitly by G , obtaining more work-efficient solutions for a wide range of parameters. The details are given in Sections 9 and 10, respectively.

THEOREM 9.2. *For any fixed constant $\varepsilon > 0$, w.h.p., an expected $O(\log k)$ -approximation to k -median on a weighted graph can be computed using $\text{polylog } n$ depth and $\tilde{O}(m^{1+\varepsilon} + k^3)$ work.*

THEOREM 10.2. *For any constant $\varepsilon > 0$, w.h.p., an expected $O(\log n)$ -approximation to the buy-at-bulk network design problem can be computed using $\text{polylog } n$ depth and $\tilde{O}(\min\{m + n(k + n^\varepsilon), n^2\}) \subseteq \tilde{O}(n^2)$ work.*

Section 11 concludes this article.

Our Approach. The algorithm of Khan et al. [30], formulated for the Congest model [41], gives rise to an $\tilde{O}(\text{SPD}(G))$ -depth parallel algorithm sampling from the FRT distribution. The SPD is the maximum, over all $v, w \in V$, of the minimum hop-length of a shortest v - w -path. Intuitively, $\text{SPD}(G)$ captures the number of iterations of MBF-like algorithms in G : Each iteration updates distances until the $(\text{SPD}(G) + 1)$ -th iteration does not yield new information. Unfortunately, $\text{SPD}(G) = n - 1$ is possible, so a naive application of this algorithm results in poor performance.

A natural idea is to reduce the number of iterations by adding “shortcuts” to the graph. Cohen [17] provides an algorithm of depth $\text{polylog } n$ and work $\tilde{O}(m^{1+\varepsilon})$ that computes a $(d, \hat{\varepsilon})$ -hop set with $d \in \text{polylog } n$: This is a set E' of additional edges such that $\text{dist}(v, w, G) \leq \text{dist}^d(v, w, G') \leq (1 + \hat{\varepsilon}) \text{dist}(v, w, G)$ for all $v, w \in V$, where $\hat{\varepsilon} \in 1/\text{polylog } n$ and $\text{dist}^d(v, w, G')$ is the minimum weight of a v - w -path with at most d edges in G augmented with E' . Note carefully that ε is different from $\hat{\varepsilon}$. In other words, Cohen computes a metric embedding with the additional property that polylogarithmically many MBF-like iterations suffice to determine $(1 + 1/\text{polylog } n)$ -approximate distances.

The course of action might now seem obvious: Run Cohen’s algorithm, then run the algorithm by Khan et al. on the resulting graph for $d \in \text{polylog } n$ rounds and conclude that the resulting output corresponds to a tree embedding of the original graph G of stretch $O((1 + 1/\text{polylog } n) \log n) = O(\log n)$. Alas, this reasoning is flawed: Constructing FRT trees crucially relies on the fact that the distances form a metric (i.e., satisfy the triangle inequality). An approximate triangle inequality for approximate distances is insufficient since the FRT construction relies on the *subtractive* form of the triangle inequality; that is, $\text{dist}(v, w, G') - \text{dist}(v, u, G') \leq \text{dist}(w, u, G')$ for arbitrary $u, v, w \in V$.

Choosing a different hop set does not solve the problem. Hop sets guarantee that d -hop distances *approximate* distances, but any hop set that fulfills the triangle inequality on d -hop distances has to reduce the SPD to at most d (i.e., yield *exact* distances):

OBSERVATION 1.1. *Let G be a graph augmented with a $(d, \hat{\varepsilon})$ -hop set.⁴ If $\text{dist}^d(\cdot, \cdot, G)$ is a metric, then $\text{dist}^d(\cdot, \cdot, G) = \text{dist}(\cdot, \cdot, G)$, i.e., $\text{SPD}(G) \leq d$.*

PROOF. Let π be a shortest u - v -path in G . Since $\text{dist}^d(\cdot, \cdot, G)$ fulfills the triangle inequality,

$$\text{dist}(u, v, G) \leq \text{dist}^d(u, v, G) \leq \sum_{\{u_1, u_2\} \in \pi} \text{dist}^d(u_1, u_2, G) \leq \sum_{\{u_1, u_2\} \in \pi} \omega(u_1, u_2) = \text{dist}(u, v, G). \quad (1)$$

□

⁴By the definitions in Section 1.2.

We overcome this obstacle by embedding G' into a complete graph H on the same node set that $(1 + o(1))$ -approximates distances in G and fulfills $\text{SPD}(H) \in \text{polylog } n$. In other words, where Cohen preserves distances *exactly* and ensures existence of *approximately* shortest paths with few hops, we preserve distances *approximately* but guarantee that we obtain *exact* shortest paths with few hops. This yields a sequence of embeddings:

- (1) Start with the original graph G ,
- (2) augment G with a $(d, 1/\text{polylog } n)$ -hop set [17], yielding G' , and
- (3) modify G' to ensure a small SPD, resulting in H (Section 4).

Unfortunately, this introduces a new obstacle: As H is complete, we cannot explicitly compute H without incurring $\Omega(n^2)$ work.

MBF-like Algorithms. This is where our novel perspective on MBF-like algorithms comes into play. We can simulate an iteration of any MBF-like algorithm on H using only the edges of G' and polylogarithmic overhead, resulting in an oracle for MBF-like queries on H . Since $\text{SPD}(H) \in \text{polylog } n$, the entire algorithm runs in polylogarithmic time and with a polylogarithmic work overhead with respect to G' .

In an iteration of an MBF-like algorithm,

- the information stored at each node is *propagated* to its neighbors,
- each node *aggregates* the received information, and
- optionally *filters* out irrelevant parts.

For example, in order for each node to determine the k nodes closest to it, each node stores node–distance pairs (initially only themselves at distance 0) and then iterates the following steps:

- communicate the node–distance pairs to the neighbors (distances uniformly increased by the corresponding edge weight),
- aggregate the received values by picking the node-wise minimum, and
- discard all but the pairs corresponding to the k closest sources.

It is well-known [3, 39, 43] that distance computations can be performed by multiplication with the (weighted) adjacency matrix A over the min-plus semiring $S_{\min,+} = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +)$ (see Definition A.2 in Appendix A). For instance, if $B = A^h$ with $h \geq \text{SPD}(G)$, then $b_{vw} = \text{dist}(v, w, G)$. In terms of $S_{\min,+}$, propagation is the “multiplication” with an edge weight and aggregation is “summation.” The $(i + 1)$ -th iteration results in $x^{(i+1)} = r^V A x^{(i)}$, where r^V is the (node-wise) filter and $x \in M^V$ the node values. Both M and M^V form *semimodules*: A semimodule supports scalar multiplication (propagation) and provides a *semigroup* (representing aggregation). Compare Definition A.3 in Appendix A—over $S_{\min,+}$.

In other words, in an h -iteration MBF-like algorithm, each node determines its part of the output based on its h -hop distances to all other nodes. However, for efficiency reasons, various algorithms [4, 7, 8, 29, 33–35] compute only a subset of these distances. The role of the filter is to remove the remaining values to allow for better efficiency. The core feature of an MBF-like algorithm is that filtering is *compatible* with propagation and aggregation: If a node discards information and then propagates it, the discarded parts must be “uninteresting” at the receiving node as well. We model this using a *congruence relation* on the node states; filters pick a suitable (efficiently encodable) representative of the node state’s equivalence class.

Constructing FRT Trees. This helps us to sample from the FRT distribution as follows. First, we observe that an MBF-like algorithm can acquire the information needed to represent an FRT tree. Second, we can *simulate* any MBF-like algorithm on H —without explicitly storing H —using

polylogarithmic overhead and MBF-like iterations on G' . The previously mentioned sampling technique decomposes the vertices and edges of H into $\Lambda \in O(\log n)$ levels. We may rewrite the adjacency matrix of H as $A_H = \bigoplus_{\lambda=0}^{\Lambda} P_{\lambda} A_{\lambda}^d P_{\lambda}$, where $\oplus$ is the “addition” of functions induced by the semimodule, P_{λ} is a projection on nodes of at least level λ , and A_{λ} is a (slightly stretched) adjacency matrix of G' . We are interested in $r^V A_H^h x^{(0)} - h$ iterations on the graph H followed by applying the node-wise filter r^V . The key insight is that the congruence relation allows us to apply intermediate filtering steps without changing the outcome as filtering does not change the equivalence class of a state. Hence, we may compute $(r^V \bigoplus_{\lambda=0}^{\Lambda} P_{\lambda} (r^V A_{\lambda})^d P_{\lambda})^h x^{(0)}$ instead. This repeated application of r^V keeps the intermediate results small, ensuring that we can perform multiplication with A_{λ} with $\tilde{O}(|E| + |E'|) \subseteq \tilde{O}(m^{1+\epsilon})$ work. Since $d \in \text{polylog } n$, $\Lambda \in O(\log n)$, and each A_{λ} accounts for $|E| + |E'|$ edges, this induces only polylogarithmic overhead with respect to iterations in G' , yielding a highly efficient parallel algorithm of depth $\text{polylog } n$ and work $\tilde{O}(m^{1+\epsilon})$.

1.1 Related Work

We confine the discussion to undirected graphs.

Classical Distance Computations. The earliest—and possibly also most basic—algorithms for Single-Source Shortest Paths (SSSP) computations are Dijkstra’s algorithm [21] and the MBF algorithm [11, 24, 40]. From the perspective of parallel algorithms, Dijkstra’s algorithm performs excellently in terms of work, requiring $\tilde{O}(m)$ computational steps, but suffers from being inherently sequential, processing one vertex at a time.

Algebraic Distance Computations. The MBF algorithm can be interpreted as a fixed-point iteration $Ax^{(i+1)} = Ax^{(i)}$, where A is the adjacency matrix of the graph G and “addition” and “multiplication” are replaced by $\min$ and $+$, respectively. This structure is known as the min-plus semiring (a.k.a. tropical semiring) $\mathcal{S}_{\min,+} = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +)$ (compare Section 1.2), which is a well-established tool for distance computations [3, 39, 43]. From this point of view, $\text{SPD}(G)$ is the number of iterations until a fixed point is reached. MBF thus has depth $\tilde{O}(\text{SPD}(G))$ and work $\tilde{O}(m \text{SPD}(G))$, where small $\text{SPD}(G)$ are possible.

One may overcome the issue of large depth entirely by performing the fixed-point iteration on the matrix by setting $A^{(0)} := A$ and iterating $A^{(i+1)} := A^{(i)}A^{(i)}$; after $\lceil \log \text{SPD}(G) \rceil \leq \lceil \log n \rceil$ iterations, a fixed point is reached [19]. The final matrix then has as entries exactly the pairwise node distances and the computation has polylogarithmic depth. This comes at the cost of $\Omega(n^3)$ work (even if $m \ll n^2$) but is as work-efficient as n instances of Dijkstra’s algorithm for solving APSP in dense graphs without incurring depth $\Omega(n)$.

Mohri [39] solved various shortest-distance problems using the $\mathcal{S}_{\min,+}$ semiring and variants thereof. While Mohri’s framework is quite general, our approach is different in crucial aspects:

- (1) Mohri uses an individual semiring for each problem and then solves it by a general algorithm. Our approach, on the other hand, is more generic as well as easier to use: We use off-the-shelf semirings—usually just $\mathcal{S}_{\min,+}$ —and combine them with appropriate semimodules carrying problem-specific information. Further problem-specific customization happens in the definition of a congruence relation on the semi-ring; it specifies which parts of a node’s state can be discarded because they are irrelevant for the problem. We demonstrate the modularity and flexibility of the approach by various examples in Section 3, which cover a large variety of distance problems.
- (2) In our framework, node states are semimodule elements and edge weights are semiring elements; hence, there is no multiplication of node states. Mohri’s approach, however, does

not make this distinction and hence requires the introduction of an artificial “multiplication” of node states.

- (3) Mohri’s algorithm can be interpreted as a generalization of Dijkstra’s algorithm [21] because it maintains a queue and, in each iteration, applies a relaxation technique to the dequeued element and its neighbors. This strategy is inherently sequential; to the best of our knowledge, we are the first to present a general algebraic framework for distance computations that exploits the implicit parallelism of the MBF algorithm.
- (4) In Mohri’s approach, choosing the global queueing strategy is not only an integral part of an algorithm, but it also simplifies the construction of the underlying semirings as one may rule that elements are processed in a “convenient” order. Our framework is flexible enough to achieve counterparts even of Mohri’s more involved results without such assumptions; concretely, we propose a suitable semiring for solving the k -Shortest Distance Problem (k -SDP) and the k -Distinct-Shortest Distance Problem (k -DSDP) in Section 3.3.

Approximate Distance Computations. As metric embeddings reproduce distances only approximately, we may base them on approximate distance computation in the original graph. Using rounding techniques and embedding $\mathcal{S}_{\min,+}$ into a polynomial ring, this enables us to use fast matrix multiplication to speed up the aforementioned fixed-point iteration $A^{(i+1)} := A^{(i)}A^{(i)}$ [43]. This reduces the work to $\tilde{O}(n^\omega)$ at the expense of only $(1 + o(1))$ -approximating distances, where $\omega < 2.3729$ [32] denotes the fast matrix-multiplication exponent. However, even if the conjecture that $\omega = 2$ holds true, this technique must result in $\Omega(n^2)$ work simply because $\Omega(n^2)$ pairwise distances are computed.

Regarding SSSP, there was no work-efficient low-depth parallel algorithm for a long time, even when allowing approximation. This was referred to as the “sequential bottleneck”: Matrix-matrix multiplication was inefficient in terms of work, while sequentially exploring (shortest) paths resulted in depth $\Omega(\text{SPD}(G))$. Klein and Subramanian [31] showed that depth $\tilde{O}(\sqrt{n})$ can be achieved with $\tilde{O}(m\sqrt{n})$ work, beating the n^2 work barrier with sublinear depth in sparse graphs. As an aside, similar bounds were later achieved for exact SSSP computations by Shi and Spencer [42].

In a seminal paper, Cohen [17] proved that SSSP can be $(1 + o(1))$ -approximated with depth $\text{polylog } n$ and near-optimal $\tilde{O}(m^{1+\epsilon})$ work for any constant choice of $\epsilon > 0$; her approach is based on the aforementioned hop set construction. Similar guarantees can be achieved deterministically. Henzinger et al. [29] focus on Congest algorithms, which can be interpreted in our framework to yield hop sets $(1 + 1/\text{polylog } n)$ -approximating distances for $d \in 2^{O(\sqrt{\log n})} \subset n^{o(1)}$ and can be computed using depth $2^{O(\sqrt{\log n})} \subset n^{o(1)}$ and work $m2^{O(\sqrt{\log n})} \subset m^{1+o(1)}$. In a recent breakthrough, Elkin and Neiman obtained hop sets with substantially improved tradeoffs [22], both for the parallel setting and the Congest model. On the negative side, Abboud et al. [1] proved principal limitations of hop sets by providing lower bounds on the tradeoffs between the parameters.

Our embedding technique is formulated independently from the underlying hop set construction, whose performance is reflected in the depth and work bounds of our algorithms. While the improvements by Elkin and Neiman do not enable us to achieve a work bound of $m^{1+o(1)}$ when sticking to our goals of depth $\text{polylog } n$ and expected stretch $O(\log n)$, they can be used to obtain better tradeoffs between the parameters.

Metric Tree Embeddings. When metrically embedding into a tree, it is, in general, impossible to guarantee a small stretch. For instance, when the graph is a cycle with unit edge weights, it is impossible to embed it into a tree without having at least one edge with stretch $\Omega(n)$. However, *on average*, the edges in this example are stretched by a constant factor only, justifying the hope that one may be able to randomly embed into a tree such that, for each pair of nodes, the *expected*

stretch is small. A number of elegant algorithms [4, 7, 8, 23] compute tree embeddings, culminating in the one by Fakcharoenphol, Rao, and Talwar [23] (FRT) that achieves stretch $O(\log n)$ in expectation. This stretch bound is optimal in the worst case, as illustrated by expander graphs [8]. Mendel and Schwoab show how to sample from the FRT distribution in $O(m \log^3 n)$ steps [37]. This upper bound has recently been improved: Blelloch et al. present an algorithm that requires time $O(m \log n)$ w.h.p. Both algorithms match the trivial $\Omega(m)$ lower bound up to polylogarithmic factors. However, their approach relies on a pruned version of Dijkstra's algorithm for distance computations and hence does not lead to a low-depth parallel algorithm.

Several parallel and distributed algorithms compute FRT trees [14, 26, 30]. These algorithms and ours have in common that they represent the embedding by Least Element (LE) lists, which were first introduced in Cohen [16, 18]. In the parallel case, the state-of-the-art solution due to Blelloch et al. [14] achieves $O(\log^2 n)$ depth and $O(n^2 \log n)$ work. However, Blelloch et al. assume the input to be given as an n -point *metric*, where the distance between two points can be queried at constant cost. Note that our approach is more general, as a metric can be interpreted as a complete weighted graph of SPD 1; a single MBF-like iteration reproduces the result by Blelloch et al. Moreover, this point of view shows that the input required to achieve subquadratic work must be a sparse graph. Furthermore, Blelloch et al. determine LE lists in $O(D \log n)$ depth and $O(W \log n)$ work, where D and W are the depth and work of an SSSP computation that is used as a black box [12]. However, to date only approximate SSSP algorithms simultaneously achieve $W \in \tilde{O}(m)$ and $D \in \text{polylog } n$; because calling an approximate SSSP algorithm multiple times does not result in approximate distances that respect the triangle inequality, this approach cannot be used to efficiently determine an FRT-style embedding. For graph inputs, we are not aware of any metric tree embedding algorithm achieving $\text{polylog } n$ depth and a nontrivial work bound (i.e., not incurring the $\Omega(n^3)$ work caused by relying on matrix-matrix multiplication).

In the distributed setting, Khan et al. [30] show how to compute LE lists in $O(\text{SPD}(G) \log n)$ rounds in the Congest model [41]. On the lower bound side, trivially $\Omega(D(G))$ rounds are required, where $D(G)$ is the maximum hop distance (i.e., ignoring weights) between nodes. However, even if $D(G) \in O(\log n)$, $\tilde{O}(\sqrt{n})$ rounds are necessary [20, 26]. Extending the algorithm by Khan et al. [26], it is shown how to obtain a round complexity of $\tilde{O}(\min\{n^{1/2+\varepsilon}, \text{SPD}(G)\} + D(G))$ for any $\varepsilon > 0$ at the expense of increasing the stretch to $O(\varepsilon^{-1} \log n)$. We partly build on these ideas; specifically, the construction in Section 4 can be seen as a generalization of the key technique from Khan et al. [26]. As detailed in Section 8, our framework subsumes these algorithms and can be used to improve on the result from Khan et al. [26]: Leveraging further results [29, 35], we obtain a metric tree embedding with expected stretch $O(\log n)$ that is computed in $\min\{n^{1/2+o(1)} + D(G)^{1+o(1)}, \tilde{O}(\text{SPD}(G))\}$ rounds.

1.2 Notation and Preliminaries

We consider weighted, undirected graphs $G = (V, E, \omega)$ without loops or parallel edges with nodes V , edges E , and edge weights $\omega : E \rightarrow \mathbb{R}_{>0}$. Unless specified otherwise, we set $n := |V|$ and $m := |E|$. For an edge $e = \{v, w\} \in E$, we write $\omega(v, w) := \omega(e)$, $\omega(v, v) := 0$ for $v \in V$, and $\omega(v, w) := \infty$ for $\{v, w\} \notin E$. We assume that the ratio between maximum and minimum edge weight is polynomially bounded in n and that each edge weight and constant can be stored with sufficient precision in a single register.⁵ We assume that G is connected and given in the form of an adjacency list.

⁵As we are interested in approximation algorithms, $O(\log n)$ bits suffice to encode values with sufficient precision.

Let $p \subseteq E$ be a path. p has $|p|$ hops and weight $\omega(p) := \sum_{e \in p} \omega(e)$. For the nodes $v, w \in V$ let $P(v, w, G)$ denote the set of paths from v to w and $P^h(v, w, G)$ the set of such paths using at most h hops. We denote by $\text{dist}^h(v, w, G) := \min\{\omega(p) \mid p \in P^h(v, w, G)\}$ the minimum weight of an h -hop path from v to w , where $\min \emptyset := \infty$; the distance between v and w is $\text{dist}(v, w, G) := \text{dist}^n(v, w, G)$. The shortest path hop distance between v and w is $\text{hop}(v, w, G) := \min\{|p| \mid p \in P(v, w, G) \wedge \omega(p) = \text{dist}(v, w, G)\}$; $\text{MHSP}(v, w, G) := \{p \in P^{\text{hop}(v, w, G)}(v, w, G) \mid \omega(p) = \text{dist}(v, w, G)\}$ denotes all *min-hop shortest paths* from v to w . Finally, the *Shortest Path Diameter (SPD)* of G is $\text{SPD}(G) := \max\{\text{hop}(v, w, G) \mid v, w \in V\}$, and $D(G) := \min\{h \in \mathbb{N} \mid \forall v, w \in V : \text{dist}^h(v, w, G) < \infty\}$ is the unweighted hop diameter of G .

We sometimes use $\min$ and $\max$ as binary operators, assume $0 \in \mathbb{N}$, and define, for a set N and $k \in \mathbb{N}$, $\binom{N}{k} := \{M \subseteq N \mid |M| = k\}$ and denote by $\text{id} : N \rightarrow N$ the identity function. Furthermore, we use weak asymptotic notation hiding polylogarithmic factors in n : $O(f(n) \text{polylog}(n)) = \tilde{O}(f(n))$, etc.

Model of Computation. We use an abstract model of parallel computation similar to those used in circuit complexity; the goal here is to avoid distraction by details such as read or write collisions or load balancing issues typical to PRAM models, noting that these can be resolved with (at most) logarithmic overheads. The computation is represented by a Directed Acyclic Graph (DAG) with constantly bounded maximum indegree, where nodes represent words of memory that are given as input (indegree 0) or computed out of previously determined memory contents (non-zero indegree). Words are computed with a constant number of basic instructions (e.g., addition, multiplication, comparison, etc.); here, we also allow for the use of independent randomness. For simplicity, a memory word may hold any number computed throughout the algorithm. As pointed out earlier, $O(\log n)$ -bit words suffice for our purpose.

An algorithm defines, given the input, the DAG and how the nodes' content is computed as well as which nodes represent the output. Given an instance of the problem, the *work* is the number of nodes of the corresponding DAG and the *depth* is its longest path. Assuming that there are no read or write conflicts, the work is thus (proportional to) the time required by a single processor (of uniform speed) to complete the computation, whereas the depth lower-bounds the time required by an infinite number of processors. Note that the DAG may be a random graph because the algorithm may use randomness, implying that work and depth may be random variables. When making probabilistic statements, we require that they hold for all instances (i.e., the respective probability bounds are satisfied after fixing an arbitrary instance).

Probability. A claim holds w.h.p. if it occurs with a probability of at least $1 - n^{-c}$ for any fixed choice of $c \in \mathbb{R}_{\geq 1}$; c is a constant for the purposes of O -notation. We use the following basic statement frequently and implicitly throughout this article:

LEMMA 1.2. Let $\mathcal{E}_1, \dots, \mathcal{E}_k$ be events occurring w.h.p., and $k \in \text{poly } n$. $\mathcal{E}_1 \cap \dots \cap \mathcal{E}_k$ occurs w.h.p.

PROOF. We have $k \leq an^b$ for fixed $a, b \in \mathbb{R}_{>0}$ and choose that all $\mathcal{E}_i$ occur with a probability of at least $1 - n^{-c'}$ with $c' = c + b + \log_n a$ for some fixed $c \geq 1$. The union bound yields

$$\mathbb{P}[\overline{\mathcal{E}_1 \cap \dots \cap \mathcal{E}_k}] \leq \sum_{i=1}^k \mathbb{P}[\bar{\mathcal{E}}_i] \leq kn^{-c'} = an^b n^{-c-b-\log_n a} = n^{-c}, \quad (2)$$

hence $\mathcal{E}_1 \cap \dots \cap \mathcal{E}_k$ occurs w.h.p. as claimed. $\square$

Hop Sets. A graph $G = (V, E, \omega)$, contains a $(d, \hat{\epsilon})$ -hop set if

$$\forall v, w \in V : \quad \text{dist}^d(v, w, G) \leq (1 + \hat{\epsilon}) \text{dist}(v, w, G); \quad (3)$$

(i.e., if its d -hop distances are a $(1 + \varepsilon)$ -approximation of its distances). This definition is based on Cohen [17], who describes how to efficiently add edges to G to establish this property.

Distance Metrics. The min-plus semiring $\mathcal{S}_{\min,+} = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +)$, also referred to as the tropical semiring, forms a semiring (i.e., a ring without additive inverses; see Definition A.2 in Appendix A). Unless explicitly stated otherwise, we associate $\oplus$ and $\odot$ with the addition and multiplication of the underlying ring throughout the article; in this case, we use $a \oplus b := \min\{a, b\}$ and $a \odot b := a + b$. Observe that ∞ and 0 are the neutral elements with respect to $\oplus$ and $\odot$, respectively. We sometimes write $x \in \mathcal{S}_{\min,+}$ instead of $x \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ to refer to the elements of a semiring. Furthermore, we follow the standard convention to occasionally leave out $\odot$ and give it priority over $\oplus$ (e.g., interpret $ab \oplus c$ as $(a \odot b) \oplus c$ for all $a, b, c \in \mathcal{S}_{\min,+}$).

The min-plus semiring is a well-established tool to determine pairwise distances in a graph via the distance product (see, e.g., [3, 39, 43]). Let $G = (V, E, \omega)$ be a weighted graph and let $A \in \mathcal{S}_{\min,+}^{V \times V}$ be its *adjacency matrix* A , given by

$$(a_{vw}) := \begin{cases} 0 & \text{if } v = w \\ \omega(v, w) & \text{if } \{v, w\} \in E \\ \infty & \text{otherwise.} \end{cases} \quad (4)$$

Throughout this article, the operations involved in matrix addition and multiplication are the operations of the underlying semiring; that is, for square matrices A, B with row and column index set V we have

$$(A \oplus B)_{vw} = \min\{a_{vw}, b_{vw}\} \text{ and} \quad (5)$$

$$(AB)_{vw} = \min_{u \in V} \{a_{vu} + b_{uw}\}. \quad (6)$$

The *distance product* A^h corresponds to h -hop distances (i.e., $(A^h)_{vw} = \text{dist}^h(v, w, G)$) [3]. In particular, this corresponds to the exact distances between all pairs of nodes for $h \geq \text{SPD}(G)$.

2 MBF-LIKE ALGORITHMS

The MBF algorithm [11, 24, 40] is both fundamental and elegant. In its classical form, it solves the SSSP problem: In each iteration, each node communicates its current upper bound on its distance to the source node s (initially ∞ at all nodes but s) plus the corresponding edge weight to its neighbors, which then keep the minimum of the received values and their previously stored one. Iterating h times determines all nodes' h -hop distances to s .

Over the years, numerous algorithms emerged that use similar iterative schemes for distributing information [4, 7, 8, 23, 29, 33–35]. It is natural to ask for a characterization that captures all of these algorithms. In this section, we propose such a characterization: the class of *MBF-like algorithms*. The common denominator of these algorithms is the following:

- An initial state vector $x^{(0)} \in M^V$ contains information initially known to each node.
- In each *iteration*, each node first *propagates* information along all incident edges.
- All nodes then *aggregate* the received information. This and the previous step are precisely the same as updating the state vector $x^{(i)}$ by the matrix-vector product $x^{(i+1)} = Ax^{(i)}$ over the min-plus semiring.
- Finally, irrelevant information is *filtered* out before moving on to the next iteration.

As a concrete example consider k -Source Shortest Paths (k -SSP), the task of determining for each node the list of its k closest nodes. To this end, one needs to consider all nodes as sources (i.e., run

the multi-source variant of the classic MBF algorithm with all nodes as sources). Nodes store values in $(\mathbb{R}_{\geq 0} \cup \{\infty\})^V$, so that in iteration i each node $v \in V$ can store $\text{dist}^i(v, w, G) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ for all $w \in V$. Initially, $x_{v,w}^{(0)}$ is 0 if $v = w$ and ∞ everywhere else (the 0-hop distances).⁶ *Propagating* these distances over an edge of weight $\omega(e)$ means uniformly increasing them by $\omega(e)$. During *aggregation*, each node picks, for each target node, the smallest distance reported so far. This is costly since each node might learn non- ∞ distances values for *all* other nodes. To increase efficiency, we *filter* out, in each iteration and at each node, all source–distance pairs but the k pairs with smallest distance. This reduces the amount of work per iteration from $\tilde{O}(mn)$ to $\tilde{O}(mk)$.

The filtering step generalizes from classic MBF to an MBF-like algorithm, with the goal of reducing work. The crucial characteristics exploited by this idea are the following:

- Propagation and aggregation are interchangeable. It makes no difference whether two pieces of information are propagated separately or as a single aggregated piece of information.
- Filtering or not filtering after aggregation has no impact on the correctness (i.e., the output) of an algorithm, only on its efficiency.

In this section, we formalize this approach for later use in more advanced algorithms. To this end, we develop a characterization of MBF-like algorithms in Sections 2.1–2.3 and establish basic properties in Section 2.4. We demonstrate that our approach applies to a wide variety of known algorithms in Section 3. In order to maintain self-containment without obstructing presentation, basic algebraic definitions are given in Appendix A.

2.1 Propagation and Aggregation

Let M be the set of node states, i.e., the possible values that an MBF-like algorithm can store at a vertex. We represent propagation of $x \in M$ over an edge of weight $s \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ by $s \odot x$, where $\odot : \mathbb{R}_{\geq 0} \cup \{\infty\} \times M \rightarrow M$, and aggregation of $x, y \in M$ at some node by $x \oplus y$, where $\oplus : M \times M \rightarrow M$; the discussion of filtering is deferred to Section 2.2. Concerning the aggregation of information, we demand that $\oplus$ is associative and has a neutral element $\perp \in M$ encoding “no available information,” hence $(M, \oplus)$ is a semigroup with neutral element $\perp$. Furthermore, we require for all $s, t \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ and $x, y \in M$ (note that we “overload” $\oplus$ and $\odot$):

$$0 \odot x = x \tag{7}$$

$$\infty \odot x = \perp \tag{8}$$

$$s \odot (x \oplus y) = (s \odot x) \oplus (s \odot y) \tag{9}$$

$$(s \oplus t) \odot x = (s \odot x) \oplus (t \odot x) \tag{10}$$

$$(s \odot t) \odot x = s \odot (t \odot x). \tag{11}$$

Our requirements are quite natural: Equations (7) and (8) state that propagating information over zero distance (e.g., keeping it at a vertex) does not alter it and that propagating it infinitely far away (i.e., “propagating” it over a nonexistent edge) means losing it, respectively. Note that 0 and ∞ are the neutral elements with respect to $\odot$ and $\oplus$ in $\mathcal{S}_{\min,+}$. Equation (9) says that propagating aggregated information is equivalent to aggregating propagated information (along identical distances), Equation (10) means that propagating information over a shorter of two edges is equivalent to moving it along both edges and then aggregating it (information “becomes obsolete” with

⁶Whenever nodes store vectors of distances, we use double-index notation: $x_v^{(0)} \in (\mathbb{R}_{\geq 0} \cup \{\infty\})^V$ is the “distance vector” stored at node v , and $x_{v,w}^{(0)} \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ is the distance regarding w stored at node v .

increasing distance), and Equation (11) states that propagating propagated information can be done in a single step.

Altogether, this is equivalent to demanding that $\mathcal{M} = (M, \oplus, \odot)$ is a zero-preserving semimodule (see Definition A.3 in Appendix A) over $\mathcal{S}_{\min,+}$. A straightforward choice of $\mathcal{M}$ is the direct product of $|V|$ copies of $\mathbb{R}_{\geq 0} \cup \{\infty\}$, which is suitable for most of the applications we consider.

Definition 2.1 (Distance Map). The *distance map semimodule* $\mathcal{D} := ((\mathbb{R}_{\geq 0} \cup \{\infty\})^V, \oplus, \odot)$ is given by, for all $s \in \mathcal{S}_{\min,+}$ and $x, y \in \mathcal{D}$,

$$(x \oplus y)_v := x_v \oplus y_v = \min\{x_v, y_v\} \quad (12)$$

$$(s \odot x)_v := s \odot x_v = s + x_v \quad (13)$$

where $\perp := (\infty, \dots, \infty)^\top \in \mathcal{D}$ is the neutral element with respect to $\oplus$.

COROLLARY 2.2. $\mathcal{D}$ is a zero-preserving semimodule over $\mathcal{S}_{\min,+}$ with zero $\perp = (\infty, \dots, \infty)^\top$ by Lemma A.4.

Distance maps can be represented by only storing the non- ∞ distances (and their indices from V). This is of interest when there are few non- ∞ entries, which can be ensured by filtering (see below). In the following, we denote by $|x|$ the number of non- ∞ entries of $x \in \mathcal{D}$. The following lemma shows that this representation allows for efficient aggregation.

LEMMA 2.3. Suppose $x_1, \dots, x_n \in \mathcal{D}$ are stored in lists of index-distance pairs as above. Then $\bigoplus_{i=1}^n x_i$ can be computed with $O(\log n)$ depth and $O(\sum_{i=1}^n |x_i| \log n)$ work.

PROOF. We sort $\bigcup_{i=1}^n x_i$ in ascending lexicographical order. This can be done in parallel with $O(\log(\sum_{i=1}^n |x_i|)) \leq O(\log n)$ depth and $O(\sum_{i=1}^n |x_i| \log n)$ work [2]. Then we delete each pair for which the next smaller pair has the same index; the resulting list thus contains, for each $v \in V$ for which there is a non- ∞ value in some list x_i , the minimum such value. As this operation is easy to implement with $O(\log n)$ depth and $O(\sum_{i=1}^n |x_i| \log n)$ work, the claim follows. $\square$

While $\mathcal{S}_{\min,+}$ and $\mathcal{D}$ suffice for most applications and are suitable to convey our ideas, it is sometimes necessary to use a different semiring. We elaborate on this in Section 3. Hence, rather than confining the discussion to semimodules over $\mathcal{S}_{\min,+}$, in the following we make general statements about an arbitrary semimodule $\mathcal{M} = (M, \oplus, \odot)$ over an arbitrary semiring $\mathcal{S} = (S, \oplus, \odot)$ wherever it does not obstruct the presentation. It is, however, helpful to keep $\mathcal{S} = \mathcal{S}_{\min,+}$ and $\mathcal{M} = \mathcal{D}$ in mind.

2.2 Filtering

MBF-like algorithms achieve efficiency by maintaining and propagating—instead of the full amount of information nodes are exposed to—only a filtered (small) representative of the information they obtained. Our goal in this section is to capture the properties that a filter must satisfy to not affect output correctness. We start with a congruence relation, an equivalence relation compatible with propagation and aggregation, on $\mathcal{M}$. A filter $r : \mathcal{M} \rightarrow \mathcal{M}$ is a projection mapping all members of an equivalence class to the same representative within that class; compare Definition 2.6.

Definition 2.4 (Congruence Relation). Let $\mathcal{M} = (M, \oplus, \odot)$ be a semimodule over the semiring $\mathcal{S}$ and $\sim$ an equivalence relation on M . We call $\sim$ a *congruence relation on $\mathcal{M}$* if and only if

$$\forall s \in \mathcal{S}, \forall x, x' \in M : \quad x \sim x' \Rightarrow sx \sim sx' \quad (14)$$

$$\forall x, x', y, y' \in M : \quad x \sim x' \wedge y \sim y' \Rightarrow x \oplus y \sim x' \oplus y'. \quad (15)$$

A congruence relation induces a quotient semimodule.

OBSERVATION 2.5. Denote by $[x]$ the equivalence class of $x \in \mathcal{M}$ under the congruence relation $\sim$ on the semimodule $\mathcal{M}$. Set $\mathcal{M}/\sim := \{[x] \mid x \in \mathcal{M}\}$. Then $\mathcal{M}/\sim := (\mathcal{M}/\sim, \oplus, \odot)$ is a semimodule with the operations $[x] \oplus [y] := [x \oplus y]$ and $s \odot [x] := [sx]$.

An MBF-like algorithm performs efficient computations by implicitly operating on this quotient semimodule (i.e., on suitable, typically small, representatives of the equivalence classes). Such representatives are obtained in the filtering step using a *representative projection*, also referred to as a *filter*. We refer to this step as filtering since, in all our applications and examples, it discards a subset of the available information that is irrelevant to the problem at hand.

Definition 2.6 (Representative Projection). Let $\mathcal{M} = (M, \oplus, \odot)$ be a semimodule over the semiring $\mathcal{S}$ and $\sim$ a congruence relation on $\mathcal{M}$. Then $r : \mathcal{M} \rightarrow \mathcal{M}$ is a *representative projection with respect to $\sim$* if and only if

$$\forall x \in \mathcal{M} : \quad x \sim r(x) \quad (16)$$

$$\forall x, y \in \mathcal{M} : \quad x \sim y \Rightarrow r(x) = r(y). \quad (17)$$

OBSERVATION 2.7. A *representative projection* is a *projection* (i.e., $r^2 = r$).

In the following, we typically first define a suitable projection r ; this projection in turn defines equivalence classes $[x] := \{y \in \mathcal{M} \mid r(x) = r(y)\}$. The following lemma is useful when we need to show that equivalence classes defined this way yield a congruence relation (i.e., are suitable for MBF-like algorithms).

LEMMA 2.8. Let $\mathcal{M}$ be a semimodule over the semiring $\mathcal{S}$, let $r : \mathcal{M} \rightarrow \mathcal{M}$ be a projection, and for $x, y \in \mathcal{M}$, let $x \sim y :\Leftrightarrow r(x) = r(y)$. Then $\sim$ is a congruence relation with representative projection r if:

$$\forall s \in \mathcal{S}, \forall x, x' \in \mathcal{M} : \quad r(x) = r(x') \Rightarrow r(sx) = r(sx'), \text{ and} \quad (18)$$

$$\forall x, x', y, y' \in \mathcal{M} : \quad r(x) = r(x') \wedge r(y) = r(y') \Rightarrow r(x \oplus y) = r(x' \oplus y'). \quad (19)$$

PROOF. Obviously, $\sim$ is an equivalence relation, and r fulfills Equations (16) and (17). Conditions (2.8) and (2.9) directly follow from the preconditions of the lemma. $\square$

An MBF-like algorithm has to behave in a compatible way for all vertices in that each vertex follows the same propagation, aggregation, and filtering rules. This induces a semimodule structure on the (possible) state vectors of the algorithm in a natural way.

Definition 2.9 (Power Semimodule). Given a node set V and a zero-preserving semimodule $\mathcal{M} = (M, \oplus, \odot)$ over the semiring $\mathcal{S}$, we define $\mathcal{M}^V = (M^V, \oplus, \odot)$ by applying the operations of $\mathcal{M}$ coordinatewise (i.e., $\forall v, w \in M^V, \forall s \in \mathcal{S}$):

$$(x \oplus y)_v := x_v \oplus y_v \text{ and} \quad (20)$$

$$(s \odot x)_v := s \odot x_v. \quad (21)$$

Furthermore, by r^V we denote the componentwise application of a representative projection r of $\mathcal{M}$,

$$(r^V x)_v := r(x_v). \quad (22)$$

This induces the equivalence relation $\sim$ on $\mathcal{M}$ via $x \sim y$ if and only if $x_v \sim y_v$ for all $v \in V$.

OBSERVATION 2.10. $\mathcal{M}^V$ is a zero-preserving semimodule over $\mathcal{S}$ and $\perp^V := (\perp, \dots, \perp)^\top \in \mathcal{M}^V$ is its neutral element with respect to $\oplus$, where $\perp$ is the neutral element of $\mathcal{M}$. The equivalence relation $\sim$ induced by r^V is a congruence relation on $\mathcal{M}^V$ with representative projection r^V .

2.3 The Class of MBF-like Algorithms

The following definition connects the properties introduced and motivated earlier:

Definition 2.11 (MBF-like Algorithm). A MBF-like algorithm $\mathcal{A}$ is determined by

- (1) a zero-preserving semimodule $\mathcal{M}$ over a semiring $\mathcal{S}$,
- (2) a congruence relation on $\mathcal{M}$ with representative projection $r : \mathcal{M} \rightarrow \mathcal{M}$, and
- (3) initial values $x^{(0)} \in \mathcal{M}^V$ for the nodes (which may depend on the input graph).

On a graph G with adjacency matrix A , h iterations of $\mathcal{A}$ determine

$$\mathcal{A}^h(G) := x^{(h)} := r^V A^h x^{(0)}. \quad (23)$$

Since $\mathcal{A}$ reaches a fixed point after at most $i = \text{SPD}(G) < n$ iterations (i.e., a state where $x^{(i+1)} = x^{(i)}$), we abbreviate $\mathcal{A}(G) := A^n(G)$.

Note that the definition of the adjacency matrix $A \in \mathcal{S}^{V \times V}$ depends on the choice of the semiring $\mathcal{S}$. For the standard choice of $\mathcal{S} = \mathcal{S}_{\min,+}$, which suffices for all our core results, we define A in Equation (4); examples using different semirings and the associated adjacency matrices are discussed in Sections 3.2–3.4.

The $(i+1)$ -th iteration of an MBF-like algorithm $\mathcal{A}$ determines $x^{(i+1)} := r^V A x^{(i)}$ (propagate, aggregate, and filter). Thus, h iterations yield $(r^V A)^h x^{(0)}$, which we show to be identical to $r^V A^h x^{(0)}$ in Corollary 2.17 of Section 2.4.

2.4 Preserving State-Equivalence Across Iterations

As motivated earlier, MBF-like algorithms filter intermediate results; a representative projection r^V determines a small representative of each node state. This maintains efficiency: Nodes propagate and aggregate only small representatives of the relevant information instead of the full amount of information they are exposed to. However, as motivated in, for example, Section 2.2, filtering is only relevant regarding efficiency, but not the correctness of MBF-like algorithms.

In this section, we formalize this concept in the following steps.

- (1) We introduce the functions needed to iterate MBF-like algorithms without filtering (i.e., multiplications with [adjacency] matrices). These *Simple Linear Functions (SLFs)* are a proper subset of the linear⁷ functions on $\mathcal{M}^V$.
- (2) The next step is to observe that SLFs are well-behaved with respect to the equivalence classes $\mathcal{M}^V / \sim$ of node states.
- (3) Equivalence classes of SLFs mapping equivalent inputs to equivalent outputs yield the functions required for the study of MBF-like algorithms. These form a semiring of (a subset of) the functions on $\mathcal{M}^V / \sim$.
- (4) Finally, we observe that $r^V \sim \text{id}$, formalizing the concepts of “operating on equivalence classes of node states” and “filtering being optional with respect to correctness.”

An SLF f is “simple” in the sense that it corresponds to matrix-vector multiplications; that is, it maps $x \in \mathcal{M}^V$ such that $(f(x))_v$ is a linear combination of the coordinates x_w , $w \in V$, of x .

Definition 2.12 (Simple Linear Function). Let $\mathcal{M}$ be a semimodule over the semiring $\mathcal{S}$. Each matrix $A \in \mathcal{S}^{V \times V}$ defines an SLF $A : \mathcal{M}^V \rightarrow \mathcal{M}^V$ (and vice versa) by

$$A(x)_v := (Ax)_v = \bigoplus_{w \in V} a_{vw} x_w. \quad (24)$$

⁷A linear function $f : \mathcal{M} \rightarrow \mathcal{M}$ on the semimodule $\mathcal{M}$ over the semiring $\mathcal{S}$ satisfies, for all $x, y \in \mathcal{M}$ and $s \in \mathcal{S}$, that $f(x \oplus y) = f(x) \oplus f(y)$ and $f(s \odot x) = s \odot f(x)$.

Thus, each iteration of an MBF-like algorithm is an application of an SLF given by an adjacency matrix followed by an application of the filter r^V . In the following, fix a semiring $\mathcal{S}$, a semimodule $\mathcal{M}$ over $\mathcal{S}$, and a congruence relation $\sim$ on $\mathcal{M}$. Furthermore, let F denote the set of SLFs (i.e., matrices $A \in \mathcal{S}^{V \times V}$), each defining a function $A : \mathcal{M}^V \rightarrow \mathcal{M}^V$.

Example 2.13 (Non-Simple Linear Function). We remark that not all linear functions on $\mathcal{M}^V$ are SLFs. Choose $V = \{1, 2\}$, $\mathcal{S} = \mathcal{S}_{\min, +}$, and $\mathcal{M} = \mathcal{D}$. Consider $f : \mathcal{M}^V \rightarrow \mathcal{M}^V$ given by

$$f \left(\begin{pmatrix} x_{11}, x_{12} \\ x_{21}, x_{22} \end{pmatrix} \right) := \begin{pmatrix} (x_{11} \oplus x_{12}, \infty) \\ \perp \end{pmatrix}. \quad (25)$$

While f is linear, $f(x)_1$ is not a linear combination of x_1 and x_2 . Hence, f is not an SLF.

Let $A, B \in F$ be SLFs. Denote by $A(x) \mapsto Ax$ the application of the SLF A to the argument $x \in \mathcal{M}^V$. Furthermore, we write $(A \oplus B)(x) \mapsto A(x) \oplus B(x)$ and $(A \circ B)(x) \mapsto A(B(x))$ for the addition and concatenation of SLFs, respectively. We proceed to Lemma 2.14, in which we show that matrix addition and multiplication are equivalent to the addition and concatenation of SLF functions, respectively. It follows that the SLFs form a semiring that is isomorphic to the matrix semiring of SLF matrices. Hence, we may use $A(x)$ and Ax interchangeably in the following.

LEMMA 2.14. $\mathcal{F} := (F, \oplus, \circ)$, where $\oplus$ denotes the addition of functions and $\circ$ their concatenation, is a semiring. Furthermore, $\mathcal{F}$ is isomorphic to the matrix-semiring over $\mathcal{S}$; that is, for all $A, B \in F$ and $x \in \mathcal{M}^V$,

$$(A \oplus B)(x) = (A \oplus B)x \text{ and} \quad (26)$$

$$(A \circ B)(x) = ABx. \quad (27)$$

PROOF. Let $A, B \in F$ and $x \in \mathcal{M}^V$ be arbitrary. Regarding Equations (26) and (27), observe that we have

$$(A \oplus B)x = Ax \oplus Bx = A(x) \oplus B(x) = (A \oplus B)(x) \text{ and} \quad (28)$$

$$ABx = A(Bx) = A(B(x)) = (A \circ B)(x), \quad (29)$$

respectively; addition and concatenation of SLFs are equivalent to addition and multiplication of their respective matrices. It follows that $\mathcal{F}$ is isomorphic to the matrix semiring $(\mathcal{S}^{V \times V}, \oplus, \odot)$ and hence $\mathcal{F}$ is a semiring as claimed. $\square$

Recall that MBF-like algorithms project node states to appropriate equivalent node states. SLFs correspond to matrices and (adjacency) matrices correspond to MBF-like iterations. Hence, it is important that SLFs are well-behaved with respect to the equivalence classes $\mathcal{M}^V / \sim$ of node states. Lemma 2.15 states that this is the case (i.e., that $Ax \sim Ax'$ for all $x' \in [x]$).

LEMMA 2.15. Let $A \in F$ be an SLF. Then we have, for all $x, x' \in \mathcal{M}^V$,

$$x \sim x' \quad \Rightarrow \quad Ax \sim Ax'. \quad (30)$$

PROOF. First, for $k \in \mathbb{N}$, let $x_1, \dots, x_k, x'_1, \dots, x'_k \in \mathcal{M}$ be such that $x_i \sim x'_i$ for all $1 \leq i \leq k$. We show that for all $s_1, \dots, s_k \in \mathcal{S}$ it holds that

$$\bigoplus_{i=1}^k s_i x_i \sim \bigoplus_{i=1}^k s_i x'_i. \quad (31)$$

We argue that Equation (31) holds by induction over k . For $k = 1$, the claim trivially follows from Equation (14). Regarding $k \geq 2$, suppose the claim holds for $k - 1$. Since $x_k \sim x'_k$, we have that

$s_k x_k \sim s_k x'_k$ by (2.8). The induction hypothesis yields $\bigoplus_{i=1}^{k-1} s_i x_i \sim \bigoplus_{i=1}^{k-1} s_i x'_i$. Hence,

$$\bigoplus_{i=1}^k s_i x_k = \left(\bigoplus_{i=1}^{k-1} s_i x_i \right) \oplus s_k x_k \stackrel{(2.9)}{\sim} \left(\bigoplus_{i=1}^{k-1} s_i x'_i \right) \oplus s_k x'_k = \bigoplus_{i=1}^k s_i x'_k. \quad (32)$$

As for the original claim, let $v \in V$ be arbitrary and note that we have

$$(Ax)_v = \bigoplus_{w \in V} a_{vw} x_w \stackrel{(2.25)}{\sim} \bigoplus_{w \in V} a_{vw} x'_w = (Ax')_v. \quad (33)$$

□

Due to Lemma 2.15, each SLF $A \in F$ not only defines a function $A : \mathcal{M}^V \rightarrow \mathcal{M}^V$, but also a function $A : \mathcal{M}^V / \sim \rightarrow \mathcal{M}^V / \sim$ with $A[x] := [Ax]$ ($A[x]$ does not depend on the choice of the representant $x' \in [x]$). This is important since MBF-like algorithms implicitly operate on $\mathcal{M}^V / \sim$ and because they do so using adjacency matrices, which are SLFs. As a natural next step, we rule for SLFs $A, B \in F$ that

$$A \sim B \quad :\Leftrightarrow \quad \forall x \in \mathcal{M}^V : Ax \sim Bx; \quad (34)$$

that is, that they are equivalent if and only if they yield equivalent results when presented with the same input. This yields equivalence classes $F / \sim = \{[A] \mid A \in F\}$. This implies, by Equation (34), that $[A][x] := [Ax]$ is well-defined. In Theorem 2.16, we show that the equivalence classes of SLFs with respect to summation and concatenation form a semiring $\mathcal{F} / \sim$. As MBF-like algorithms implicitly work on $\mathcal{M}^V / \sim$, we obtain with $\mathcal{F} / \sim$ precisely the structure that may be used to manipulate the state of MBF-like algorithms, which we leverage throughout this article.

THEOREM 2.16. *Each $[A] \in F / \sim$ defines an SLF on $\mathcal{M}^V / \sim$. Furthermore, $\mathcal{F} / \sim := (F / \sim, \oplus, \circ)$, where $\oplus$ denotes the addition and $\circ$ the concatenation of functions, is a semiring of SLFs on $\mathcal{M}^V / \sim$ with*

$$[A] \oplus [B] = [A \oplus B] \text{ and} \quad (35)$$

$$[A] \circ [B] = [AB]. \quad (36)$$

PROOF. As argued earlier, for any $A \in F$, $[A] \in F / \sim$ is well-defined on $\mathcal{M}^V / \sim$ by Lemma 2.15. Equations (35) and (36) follow from Equations (26) and (27), respectively:

$$[A \oplus B][x] = [(A \oplus B)x] \stackrel{(2.20)}{=} [(A \oplus B)(x)] = ([A] \oplus [B])([x]) \quad (37)$$

$$[AB][x] = [ABx] \stackrel{(2.21)}{=} [(A \circ B)(x)] = [A \circ B]([x]). \quad (38)$$

To see that $[A]$ is linear, let $s \in \mathcal{S}$ and $x, y \in \mathcal{M}^V$ be arbitrary and compute

$$\begin{aligned} [A][x] \oplus [A][y] &= [Ax] \oplus [Ay] = [Ax \oplus Ay] = [A(x \oplus y)] = [A][x \oplus y] \\ &= [A]([x] \oplus [y]) \text{ and} \end{aligned} \quad (39)$$

$$[A](s[x]) = [A(sx)] = [s(Ax)] = s[Ax] = s[A][x]. \quad (40)$$

This implies that $\mathcal{F} / \sim$ is a semiring of linear functions. As each function $[A]$ is represented by multiplication with (any) SLF $A' \in [A]$, $[A]$ is an SLF. □

The following corollary is a key property used throughout this article. It allows us to apply filter steps whenever convenient. We later use this to simulate MBF-like iterations on an implicitly represented graph whose edges correspond to entire paths in the original graph. This is efficient only because we have the luxury of applying intermediate filtering repeatedly without affecting the output.

COROLLARY 2.17 ($r^V \sim \text{id}$). For any representative projection r on $\mathcal{M}$, we have $r^V \sim \text{id}$; that is, for any SLF $A \in F$ it holds that

$$r^V A \sim A r^V \sim A. \quad (41)$$

In particular (as promised in Section 2.3) for any MBF-like algorithm $\mathcal{A}$, we have

$$\mathcal{A}^h(G) \stackrel{(2.17)}{=} r^V A^h x^{(0)} \stackrel{(2.35)}{=} (r^V A)^h x^{(0)}. \quad (42)$$

Finally, we stress that both the restriction to SLFs and the componentwise application of r in r^V are crucial for Corollary 2.17.

Example 2.18 (Non-Simple Linear Functions break Corollary 2.17). Consider V , $\mathcal{M}$, and f from Example 2.13. If $r(x) = (x_1, \infty)$ for all $x \in \mathcal{M}$, we have that

$$r^V f \left(\begin{pmatrix} (2, 1) \\ \perp \end{pmatrix} \right) = \left(\begin{pmatrix} (1, \infty) \\ \perp \end{pmatrix} \right) \approx \left(\begin{pmatrix} (2, \infty) \\ \perp \end{pmatrix} \right) = f r^V \left(\begin{pmatrix} (2, 1) \\ \perp \end{pmatrix} \right), \quad (43)$$

implying that $r^V f \approx f r^V$.

Example 2.19 (Non-Component-Wise Filtering Breaks Corollary 2.17). Consider $V = \{1, 2\}$, $\mathcal{S} = \mathcal{S}_{\min, +}$, and $\mathcal{M} = \mathcal{D}$. Suppose f is the SLF given by $f x := \begin{pmatrix} x_1 \oplus x_2 \\ \perp \end{pmatrix}$ and $r^V(x) := \begin{pmatrix} x_1 \\ \perp \end{pmatrix}$; that is, r^V is not a component-wise application of some representative projection r on $\mathcal{M}$, but still a representative projection on $\mathcal{M}^V$. Then we have that

$$r^V f \left(\begin{pmatrix} (2, \infty) \\ (1, \infty) \end{pmatrix} \right) = r^V \left(\begin{pmatrix} (1, \infty) \\ \perp \end{pmatrix} \right) = \left(\begin{pmatrix} (1, \infty) \\ \perp \end{pmatrix} \right) \approx \left(\begin{pmatrix} (2, \infty) \\ \perp \end{pmatrix} \right) = f \left(\begin{pmatrix} (2, \infty) \\ \perp \end{pmatrix} \right) = f r^V \left(\begin{pmatrix} (2, \infty) \\ (1, \infty) \end{pmatrix} \right), \quad (44)$$

again implying that $r^V f \approx f r^V$.

3 A COLLECTION OF MBF-LIKE ALGORITHMS

For the purpose of illustration and to demonstrate the generality of our framework, we show that a variety of standard algorithms are MBF-like algorithms; due to the machinery established earlier, this is a trivial task in many cases. In order to provide an unobstructed view on the machinery—and since this section is not central to our contributions—we defer proofs to Appendix B.

We demonstrate that some more involved distributed algorithms in the Congest model have a straightforward and compact interpretation in our framework in Section 8. They compute metric tree embeddings based on the FRT distribution; we present them alongside an improved distributed algorithm based on the other results of this work.

MBF-like algorithms are specified by a zero-preserving semimodule $\mathcal{M}$ over a semiring $\mathcal{S}$, a representative projection of a congruence relation on $\mathcal{M}$, initial states $x^{(0)}$, and the number of iterations h (compare Definition 2.11). While this might look like a lot, typically, a standard semiring and semimodule can be chosen; the general-purpose choices of $\mathcal{S} = \mathcal{S}_{\min, +}$ and $\mathcal{M} = \mathcal{D}$ (see Definition 2.1 and Corollary 2.2) or $\mathcal{M} = \mathcal{S}_{\min, +}$ (every semiring is a zero-preserving semimodule over itself) usually are up to the task. Refer to Sections 3.2 and 3.3 for examples that require different semirings. However, even in these cases, the semirings and semimodules specified in Sections 3.2 and 3.3 can be reused. Hence, all that is left to do in most cases is to pick an existing semiring and semimodule, choose $h \in \mathbb{N}$, and specify a representative projection r .

3.1 MBF-like Algorithms over the Min-Plus Semiring

We demonstrate that the min-plus semiring $\mathcal{S}_{\min, +}$ (a.k.a. the tropical semiring) is the semiring of choice to capture many standard distance problems. Note that we also use $\mathcal{S}_{\min, +}$ in our core result (i.e., for sampling FRT trees). For the sake of completeness, first recall the adjacency matrix A of

the weighted graph G in the semiring $\mathcal{S}_{\min,+}$ from Equation (4) and the distance-map semimodule $\mathcal{D}$ from Definition 2.1; consider the initialization $x^{(0)} \in \mathcal{D}^V$ with

$$x_{vw}^{(0)} := \begin{cases} 0 & \text{if } v = w \text{ and} \\ \infty & \text{otherwise,} \end{cases} \quad (45)$$

and observe that the entries of

$$x^{(h)} := A^h x^{(0)} \quad (46)$$

correspond to the h -hop distances in G :

LEMMA 3.1. *For $h \in \mathbb{N}$ and $x^{(h)}$ from Equation (46), we have*

$$x_{vw}^{(h)} = \text{dist}^h(v, w, G). \quad (47)$$

It is well-known that the min-plus semiring can be used for distance computations [3, 39, 43]. Nevertheless, for the sake of completeness, we prove Lemma 3.1 in terms of our notation in Appendix B.

As a first example, we turn our attention to source detection. It generalizes all examples covered in this section, saving us from proving each one of them correct; well-established examples like SSSP and APSP follow. Source detection was introduced by Lenzen and Peleg [36]. Note, however, that we include a maximum considered distance d in the definition.

Example 3.2 (Source Detection [36]). Given a weighted graph $G = (V, E, \omega)$, sources $S \subseteq V$, hop and result limits $h, k \in \mathbb{N}$, and a maximum distance $d \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, (S, h, d, k) -source detection is the following problem: For each $v \in V$, determine the k smallest elements of $\{(\text{dist}^h(v, s, G), s) \mid s \in S, \text{dist}(v, s, G) \leq d\}$ with respect to lexicographical ordering, or all of them if there are fewer than k .

Source detection is solved by h iterations of an MBF-like algorithm with $\mathcal{S} = \mathcal{S}_{\min,+}$, $\mathcal{M} = \mathcal{D}$,

$$r(x)_v \mapsto \begin{cases} x_v & \text{if } v \in S, x_v \leq d, \text{ and } x_v \text{ is among } k \text{ smallest entries of } x \text{ (ties broken by index),} \\ \infty & \text{otherwise,} \end{cases} \quad (48)$$

and $x_{vv}^{(0)} = 0$ if $v \in S$ and $x_{vw}^{(0)} = \infty$ in all other cases.

Since it may not be obvious that r is a representative projection, we prove it in Appendix B.

Example 3.3 (SSSP). Single-Source Shortest Paths (SSSP) requires us to determine the h -hop distance to $s \in V$ for all $v \in V$. It is solved by an MBF-like algorithm with $\mathcal{S} = \mathcal{M} = \mathcal{S}_{\min,+}$, $r = \text{id}$, and $x_s^{(0)} = 0$, $x_v^{(0)} = \infty$ for all $v \neq s$.

Equivalently, one may use $(\{s\}, h, \infty, 1)$ -source detection, effectively resulting in $\mathcal{M} = \mathcal{S}_{\min,+}$. When only storing the non- ∞ entries, only the s -entry is relevant; however, the vertex ID of s is stored as well—and $r = \text{id}$, too.

Example 3.4 (k -SSP). The k -SSP requires us to determine, for each node, the k closest nodes in terms of the h -hop distance $\text{dist}^h(\cdot, \cdot, G)$. It is solved by an MBF-like algorithm, as it corresponds to (V, h, ∞, k) -source detection.

Example 3.5 (APSP). All-Pairs Shortest Paths (APSP) is the task of determining the h -hop distance between all pairs of nodes. It is solved by an MBF-like algorithm because we can use (V, h, ∞, n) -source detection, resulting in $\mathcal{M} = \mathcal{D}$, $r = \text{id}$, and $x^{(0)}$ from Equation (45).

Example 3.6 (MSSP). In the *Multi-Source Shortest Paths (MSSP)* problem, each node is looking for the h -hop distances to all nodes in a designated set $S \subseteq V$ of source nodes. This is solved by the MBF-like algorithm for $(S, h, \infty, |S|)$ -source detection.

Example 3.7 (Forest Fires). The nodes in a graph G form a distributed sensor network, the edges represent communication channels, and edge weights correspond to distances. Our goal is to detect, for each node v , if there is a node w on fire within distance $\text{dist}(v, w, G) \leq d$ for some $d \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, where every node initially knows if it is on fire. As a suitable MBF-like algorithm, pick $h = n$, $\mathcal{S} = \mathcal{M} = \mathcal{S}_{\min, +}$,

$$r(x) \mapsto \begin{cases} x & \text{if } x \leq d \text{ and} \\ \infty & \text{otherwise,} \end{cases} \quad (49)$$

and $x_v^{(0)} = 0$ if v is on fire and $x_v^{(0)} = \infty$ otherwise.

Example 3.7 can be handled differently by using $(\mathcal{S}, n, d, 1)$ -source detection, where $\mathcal{S}$ are the nodes on fire. This also reveals the closest node on fire, whereas the solution from Example 3.7 works in anonymous networks. One can interpret both solutions as instances of SSSP with a virtual source $s \notin V$ that is connected to all nodes on fire by an edge of weight 0. This, however, requires a simulation argument and additional reasoning if the closest node on fire is to be determined.

3.2 MBF-like Algorithms over the Max-Min Semiring

Some problems require using a semiring other than $\mathcal{S}_{\min, +}$. As an example, consider the Widest Path Problem (WPP), also referred to as the bottleneck shortest path problem: Given two nodes v and w in a weighted graph, find a v - w -path maximizing the lightest edge in the path. More formally, we are interested in the widest-path distance between v and w :

Definition 3.8 (Widest-Path Distance). Given a weighted graph $G = (V, E, \omega)$, a path p has width $\text{width}(p) := \min\{\omega(e) \mid e \in p\}$. The h -hop widest-path distance between $v, w \in V$ is

$$\text{width}^h(v, w, G) := \max_{p \in \mathcal{P}^h(v, w, G)} \{\text{width}(p)\}. \quad (50)$$

We abbreviate $\text{width}(v, w, G) := \text{width}^n(v, w, G)$.

An application of the WPP are trust networks: The nodes of a graph are entities, and an edge $\{v, w\}$ of weight $0 < \omega(v, w) \leq 1$ encodes that v and w trust each other with $\omega(v, w)$. Assuming trust to be transitive, v trusts w with $\max_{p \in \mathcal{P}(v, w, G)} \min_{e \in p} \omega(e) = \text{width}(v, w, G)$. The WPP requires a semiring supporting the max and min operations:

Definition 3.9 (Max-Min Semiring). We refer to $\mathcal{S}_{\max, \min} := (\mathbb{R}_{\geq 0} \cup \{\infty\}, \max, \min)$ as the *max-min semiring*.

LEMMA 3.10. $\mathcal{S}_{\max, \min}$ is a semiring with neutral elements 0 and ∞ .

Proof in Appendix B.

COROLLARY 3.11. $\mathcal{S}_{\max, \min}$ is a zero-preserving semimodule over itself. Furthermore, we have that $\mathcal{W} := ((\mathbb{R}_{\geq 0} \cup \{\infty\})^V, \oplus, \odot)$ with, for all $x, y \in (\mathbb{R}_{\geq 0} \cup \{\infty\})^V$ and $s \in \mathbb{R}_{\geq 0} \cup \{\infty\}$,

$$(x \oplus y)_v := \max\{x_v, y_v\} \quad (51)$$

$$(s \odot x)_v := \min\{s, x_v\} \quad (52)$$

is a zero-preserving semimodule over $\mathcal{S}_{\max, \min}$ with zero $\perp = (0, \dots, 0)^\top$ by Lemma A.4.

As adjacency matrix of $G = (V, E, \omega)$ with respect to $\mathcal{S}_{\max, \min}$, we propose $A \in \mathcal{S}_{\max, \min}^{V \times V}$ with

$$(a_{vw}) := \begin{cases} \infty & \text{if } v = w, \\ \omega(v, w) & \text{if } \{v, w\} \in E, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (53)$$

This is a straightforward adaptation of the adjacency matrix with respect to $\mathcal{S}_{\min,+}$ in Equation (4). As an initialization, $x^{(0)} \in \mathcal{W}^V$ in which each node knows the trivial path of unbounded width to itself, but nothing else, is given by

$$x_{vw}^{(0)} := \begin{cases} \infty & \text{if } v = w \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (54)$$

Then, $1 \leq h \in \mathbb{N}$ multiplications with A (i.e., h iterations) yield

$$x^{(h)} := A^h x^{(0)} \quad (55)$$

which corresponds to the h -hop widest-path distance:

LEMMA 3.12. *Given $x^{(h)}$ from Equation (55), we have*

$$x_{vw}^{(h)} = \text{width}^h(v, w, G). \quad (56)$$

Proof in Appendix B.

Example 3.13 (Single-Source Widest Paths). *Single-Source Widest Paths (SSWP)* asks for, given a weighted graph $G = (V, E, \omega)$, a designated source node $s \in V$, and $h \in \mathbb{N}$, the h -hop widest-path distance $\text{width}^h(s, v, G)$ for every $v \in V$. It is solved by an MBF-like algorithm with $\mathcal{S} = \mathcal{M} = \mathcal{S}_{\max, \min}$, $r = \text{id}$, and $x_s^{(0)} = \infty$ and $x_v^{(0)} = 0$ for all $v \neq s$.

Example 3.14 (All-Pairs Widest Paths). *All-Pairs Widest Paths (APWP)* asks for, given $G = (V, E, \omega)$ and $h \in \mathbb{N}$, $\text{width}^h(v, w, G)$ for all $v, w \in V$. APWP is MBF-like; it is solved by choosing $\mathcal{S} = \mathcal{S}_{\max, \min}$, $\mathcal{M} = \mathcal{W}$, $r = \text{id}$, and $x^{(0)}$ from Equation (54) by Lemma 3.12.

Example 3.15 (Multi-Source Widest Paths). In the *Multi-Source Widest Paths (MSWP)* problem, each node is looking for the h -hop widest path distance to all nodes in a designated set $S \subseteq V$ of source nodes. This is solved by the same MBF-like algorithm as for APWP in Example 3.14 when changing $x^{(0)}$ to $x_{vw}^{(0)} = \infty$ if $v = w \in S$ and $x_{vw}^{(0)} = 0$ otherwise.

3.3 MBF-like Algorithms over the All-Paths Semiring

Mohri discusses k -SDP, where each $v \in V$ is required to find the k shortest paths to a designated source node $s \in V$, in the light of his algebraic framework for distance computations [39]. Our framework captures this application as well but requires a different semiring than $\mathcal{S}_{\min,+}$: While $\mathcal{S}_{\min,+}$ suffices for many applications (see Section 3.1), it cannot distinguish between different paths of the same length. This is a problem in the k -SDP because there may be multiple paths of the same length among the k shortest.

OBSERVATION 3.16. *No semimodule $\mathcal{M}$ over $\mathcal{S}_{\min,+}$ can overcome this issue: The left-distributive law (A.9) requires, for all $x \in \mathcal{M}$ and $s, s' \in \mathcal{S}_{\min,+}$, that $sx \oplus s'x = (s \oplus s')x$. Consider different paths $\pi \neq \pi'$ ending in the same node with $\omega(\pi) = s = s' = \omega(\pi')$. With respect to $\mathcal{S}_{\min,+}$ and $\mathcal{M}$, the left-distributive law yields $sx \oplus s'x = \min\{s, s'\} \odot x$; that is, propagating x over π , over π' , or over both and then aggregating must be indistinguishable in the case of $s = s'$.*

This does not mean that the framework of MBF-like algorithms cannot be applied, but rather it indicates that the toolbox needs a more powerful semiring than $\mathcal{S}_{\min,+}$. The motivation of this section is to add such a semiring, the *all-paths semiring* $\mathcal{P}_{\min,+}$, to the toolbox. Having established $\mathcal{P}_{\min,+}$, the advantages of the previously established machinery are available: pick a semimodule (or use $\mathcal{P}_{\min,+}$ itself) and define a representative projection. We demonstrate this for k -SDP and a variant.

The basic concept of $\mathcal{P}_{\min,+}$ is simple: remember *paths* instead of adding up “anonymous” distances. Instead of storing the sum of the traversed edges’ weight, store the string of edges. We also add the ability to remember multiple paths into the semiring. This includes enough features in $\mathcal{P}_{\min,+}$; we do not require dedicated semimodules for k -SDP, and we use the fact that $\mathcal{P}_{\min,+}$ is a zero-preserving semimodule over itself.

We begin the technical part with a convenient representation of paths: Let $P \subset V^+$ denote the set of non-empty, loop-free, directed paths on V , denoted as tuples of nodes. Furthermore, let $\circ \subseteq P^2$ be the relation of *concatenable* paths defined by

$$(v_1, \dots, v_k) \circ (w_1, \dots, w_\ell) \quad :\Leftrightarrow \quad v_k = w_1. \quad (57)$$

By abuse of notation, when and if its operands are concatenable, we occasionally use $\circ$ as the concatenation operator. Furthermore, we use $\{(\pi^1, \pi^2) \mid \pi = \pi^1 \circ \pi^2\}$ as a shorthand for the rather cumbersome $\{(\pi^1, \pi^2) \mid \pi^1, \pi^2 \in P \wedge \pi^1 \circ \pi^2 \wedge \pi \text{ is the concatenation of } \pi^1 \text{ and } \pi^2\}$ to iterate over all two-splits of π . We call a path π *valid with respect to G* if $\pi \in P(G)$ and *invalid with respect to G* otherwise.

As motivated earlier, the all-paths semiring can store multiple paths. We represent this using vectors in $(\mathbb{R}_{\geq 0} \cup \{\infty\})^P$ storing a non- ∞ weight for every encountered path and ∞ for all paths not encountered so far. This can be efficiently represented by implicitly leaving out all ∞ entries.

Definition 3.17 (All-Paths Semiring). We call $\mathcal{P}_{\min,+} = ((\mathbb{R}_{\geq 0} \cup \{\infty\})^P, \oplus, \odot)$ the *all-paths semiring*, where $\oplus$ and $\odot$ are defined, for all $\pi \in P$ and $x, y \in \mathcal{P}_{\min,+}$, by

$$(x \oplus y)_\pi := \min\{x_\pi, y_\pi\} \text{ and} \quad (58)$$

$$(x \odot y)_\pi := \min\{x_{\pi^1} + y_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\}. \quad (59)$$

We say that x *contains π* (with weight x_π) if and only if $x_\pi < \infty$.

Summation picks the smallest weight associated with each path in either operand; multiplication $(x \odot y)_\pi$ finds the lightest estimate for π composed of two-splits $\pi = \pi^1 \circ \pi^2$, where π^1 is picked from x and π^2 from y . Observe that $\mathcal{P}_{\min,+}$ supports upper bounds on path lengths; we do not, however, use this feature. Intuitively, $\mathcal{P}_{\min,+}$ stores all encountered paths with their exact weights; in this mindset, summation corresponds to the union and multiplication to the concatenability-obeying Cartesian product of the paths contained in x and y .

LEMMA 3.18. $\mathcal{P}_{\min,+}$ is a semiring with neutral elements

$$0 := (\infty, \dots, \infty)^\top \text{ and} \quad (60)$$

$$1_\pi := \begin{cases} 0 & \text{if } \pi = (v) \text{ for some } v \in V \text{ and} \\ \infty & \text{otherwise} \end{cases} \quad (61)$$

with respect to $\oplus$ and $\odot$, respectively.

Proof in Appendix B.

COROLLARY 3.19. $\mathcal{P}_{\min,+}$ is a zero-preserving semimodule over itself.

Computations on a graph $G = (V, E, \omega)$ with respect to $\mathcal{P}_{\min,+}$ require (this is a generalization of Equation (4)) an adjacency matrix $A \in \mathcal{P}_{\min,+}^{V \times V}$ defined by

$$(a_{vw})_\pi := \begin{cases} 1_\pi & \text{if } v = w, \\ \omega(v, w) & \text{if } \pi = (v, w), \text{ and} \\ \infty & \text{otherwise.} \end{cases} \quad (62)$$

On the diagonal, $a_{vv} = 1_\pi$ contains exactly the zero-hop paths of weight 0; all nontrivial paths are “unknown” in a_{vv} (i.e., accounted for with an infinite weight). An entry a_{vw} with $v \neq w$ contains,

if present, only the edge $\{v, w\}$, represented by the path (v, w) of weight $\omega(v, w)$; all other paths are not contained in a_{vw} . An initialization where each node v knows only about the zero-hop path (v) is represented by the vector $x^{(0)} \in \mathcal{P}_{\min,+}^V$ with

$$(x_v^{(0)})_\pi := \begin{cases} 0 & \text{if } \pi = (v) \text{ and} \\ \infty & \text{otherwise.} \end{cases} \quad (63)$$

Then, $1 \leq h \in \mathbb{N}$ multiplications of $x^{(0)}$ with A (i.e., h iterations) yield $x^{(h)}$ with

$$x^{(h)} := A^h x^{(0)}. \quad (64)$$

As expected, $x_v^{(h)}$ contains exactly the h -hop paths beginning in v with their according weights:

LEMMA 3.20. *Let $x^{(h)}$ be defined as in Equation (63), with respect to the graph $G = (V, E, \omega)$. Then, for all $v \in V$ and $\pi \in P$*

$$(x_v^{(h)})_\pi = \begin{cases} \omega(\pi) & \text{if } \pi \in P^h(v, \cdot, G) \text{ and} \\ \infty & \text{otherwise.} \end{cases} \quad (65)$$

Proof in Appendix B.

With the all-paths semiring $\mathcal{P}_{\min,+}$ established, we turn to the k -SDP, our initial motivation for adding $\mathcal{P}_{\min,+}$ to the toolbox of MBF-like algorithms in the first place.

Definition 3.21 (k -Shortest Distance Problem [39]). Given a graph $G = (V, W, \omega)$ and a designated source vertex $s \in V$, the k -SDP asks: For each node $v \in V$ and considering all v - s -paths, what are the weights of the k lightest such paths? In the k -DSDP, the path weights have to be distinct.

Observe that with the preceding definitions of A and $x^{(h)}$, we always associate a path π with either its weight $\omega(\pi)$ or with ∞ ; in particular, invalid paths always are associated with ∞ . Formally, G induces a subsemiring of $\mathcal{P}_{\min,+}$. In addition to being an interesting observation, these properties are required for the representative projections defined later (r breaks for the k -DSDP when facing inconsistent non- ∞ values for the same path) so we formalize them. Let $G = (V, E, \omega)$ be a graph and let $D(G) \subset (\mathbb{R}_{\geq 0} \cup \{\infty\})^P$ be the restriction of $(\mathbb{R}_{\geq 0} \cup \{\infty\})^P$ to exact path weights and ∞ :

$$x_\pi \in \{\omega(\pi), \infty\}; \quad (66)$$

recall that $\omega(\pi) = \infty$ for all paths π invalid with respect to G .

Definition 3.22 (Graph-Induced All-Paths Semiring). Let G be a graph and $D(G) \subset (\mathbb{R}_{\geq 0} \cup \{\infty\})^P$ as above. Then we refer to $\mathcal{P}_{\min,+}(G) := (D(G), \oplus, \odot)$ as the *all-paths semiring induced by G* , where $\oplus$ and $\odot$ are the same as in Definition 3.17.

The next step is to show that $\mathcal{P}_{\min,+}(G)$ is a semiring.⁸

LEMMA 3.23. *$\mathcal{P}_{\min,+}(G)$ is a semiring.*

Proof in Appendix B.

COROLLARY 3.24. *$\mathcal{P}_{\min,+}(G)$ is a zero-preserving semimodule over itself.*

Observe that we have $A \in \mathcal{P}_{\min,+}^{V \times V}(G)$ as well as $x^{(0)} \in \mathcal{P}_{\min,+}^V(G)$. It follows that Lemma 3.20 holds for $\mathcal{P}_{\min,+}(G)$ as much as it does for $\mathcal{P}_{\min,+}$. Furthermore, observe that the restriction to $\mathcal{P}_{\min,+}(G)$ happens implicitly, simply by starting with the preceding initialization. There is no information

⁸It is a subsemiring of $\mathcal{P}_{\min,+}$. We do not, however, use that property and hence refrain from formally introducing the according definitions.

about $\mathcal{P}_{\min,+}(G)$ that needs to be distributed in the graph in order to run an MBF-like algorithm over $\mathcal{P}_{\min,+}(G)$.

In order to solve the k -SDP, we require a representative projection that reduces the abundance of paths stored in an unfiltered $x^{(h)}$ to the relevant ones. Relevant in this case simply means to keep the k shortest v - s -paths in $x_v^{(h)}$. In order to formalize this, let $P(v, w, x)$ denote, for $x \in \mathcal{P}_{\min,+}(G)$ and $v, w \in V$, the set of all v - w -paths contained in x :

$$P(v, w, x) := \{\pi \in P \mid \pi \text{ is a } v\text{-}w\text{-path with } x_\pi \neq \infty\}. \quad (67)$$

Order $P(v, w, x)$ ascendingly with respect to the weights x_π , breaking ties using lexicographical order on P . Then let $P_k(v, w, x)$ denote the set of the first (at most) k entries of that sequence:

$$P_k(v, w, x) := \{\pi \mid (x_\pi, \pi) \text{ is among the } k \text{ smallest of } \{(x_{\pi'}, \pi') \mid \pi' \in P(v, w, x)\}\}. \quad (68)$$

We define the (representative, see below) projection $r : \mathcal{P}_{\min,+}(G) \rightarrow \mathcal{P}_{\min,+}(G)$ by

$$r(x)_\pi \mapsto \begin{cases} x_\pi & \text{if } \pi \in P_k(v, s, x) \text{ for some } v \in V \text{ and} \\ \infty & \text{otherwise.} \end{cases} \quad (69)$$

If $x_\pi = r(x)_\pi$, we say that r keeps π and otherwise that r discards π . The projection r keeps, for each $v \in V$, exactly the k shortest v - s -paths contained in x . Following the standard approach (see Lemma 2.8) we define vectors $x, y \in \mathcal{P}_{\min,+}$ to be equivalent if and only if their entries for $P_k(\cdot, s, x)$ do not differ:

$$\forall x, y \in \mathcal{P}_{\min,+}(G) : \quad x \sim y \quad :\Leftrightarrow \quad r(x) = r(y). \quad (70)$$

LEMMA 3.25. $\sim$ is a congruence relation on $\mathcal{P}_{\min,+}(G)$ with representative projection r .

Proof in Appendix B.

Observe that r is defined to maintain the k shortest v - s -paths for all $v \in V$, potentially storing $k|V|$ paths instead of just k . Intuitively, one could argue that $r^V x_v^{(h)}$ only needs to contain k paths since they all start in v , which is what the algorithm should actually be doing. This objection is correct in that this is what actually happens when running the algorithm with initialization $x^{(0)}$: By Lemma 3.20, $x_v^{(h)}$ contains the h -hop shortest paths starting in v and r removes all that do not end in s or are too long. On the other hand, the objection is flawed. In order for r to behave correctly with respect to all $x \in \mathcal{P}_{\min,+}$, especially those less nicely structured than $x_v^{(h)}$ where all paths start at v , we must define r as it is; otherwise, the proof of Lemma 3.25 fails for mixed starting-node inputs.

Example 3.26 (k -Shortest Distance Problem). k -SDP (compare Definition 3.21) is solved by an MBF-like algorithm $\mathcal{A}$ with $\mathcal{S} = \mathcal{M} = \mathcal{P}_{\min,+}(G)$, the representative projection and congruence relation defined in Equations (69) and (70), the choices of A and $x^{(0)}$ from Equations (62) and (63), and $h = \text{SPD}(G)$ iterations.

By Lemma 3.20 and due to $h = \text{SPD}(G)$, $x_v^{(h)}$ contains all paths that start in v , associated with their weights. Since $\mathcal{A}^h(G) = r^V x^{(h)}$, by definition of r in Equation (69), $(r^V x^{(h)})_v = r(x_v^{(h)})$ contains the subset of those paths that have the k smallest weights and start in v (i.e., precisely what k -SDP asks for).

We remark that solving a generalization of k -SDP looking for the k shortest h -hop distances is straightforward using h iterations. Furthermore, note that our approach reveals the actual paths along with their weights.

Example 3.27 (k -Distinct-Shortest Distance Problem). k -DSDP from Definition 3.21 can be solved analogously to k -SDP in Example 3.26.

In order for this to work, the definition of $P_k(v, w, x)$ in Equation (68) needs to be adjusted. For each of the k smallest weights in x , the modified $\bar{P}_k(v, w, x)$ contains only one representative: the path contained in x of that weight that is first with respect to lexicographical order on P . This results in

$$\bar{P}'_k(v, w, x) := \{\pi \mid x_\pi \text{ is among the } k \text{ smallest of } \{x_{\pi'} \mid \pi' \in P(v, w, x)\}\} \text{ and} \quad (71)$$

$$\bar{P}_k(v, w, x) := \{\pi \mid \pi \text{ is lexicographically first of } \{\pi' \in \bar{P}'_k(v, w, x) \mid x_{\pi'} = x_\pi\}\}. \quad (72)$$

The proof of Lemma 3.25 works without modification when replacing Equation (68) with Equations (71)–(72).

3.4 MBF-like Algorithms over the Boolean Semiring

A well-known semiring is the Boolean semiring $\mathcal{B} = (\{0, 1\}, \vee, \wedge)$. By Lemma A.4, $\mathcal{B}^V$ is a zero-preserving semimodule over $\mathcal{B}$. It can be used to check for connectivity in a graph⁹ using the adjacency matrix

$$(a_{vw}) := \begin{cases} 1 & \text{if } v = w \text{ or } \{v, w\} \in E \text{ and} \\ 0 & \text{otherwise} \end{cases} \quad (73)$$

together with initial values

$$x_{vw}^{(0)} := \begin{cases} 1 & \text{if } v = w \text{ and} \\ 0 & \text{otherwise} \end{cases} \quad (74)$$

indicating that each node $v \in V$ is connected to itself. An inductive argument reveals that

$$(A^h x^{(0)})_{vw} = 1 \iff P^h(v, w, G) \neq \emptyset. \quad (75)$$

Example 3.28 (Connectivity). Given a graph, we want to check which pairs of nodes are connected by paths of at most h hops. This is solved by an MBF-like algorithm using $\mathcal{S} = \mathcal{B}$, $\mathcal{M} = \mathcal{B}^V$, $r = \text{id}$, and $x^{(0)}$ from Equation (74). This example easily generalizes to single-source and multi-source connectivity variants.

4 THE SIMULATED GRAPH

In order to sample from a tree embedding of the graph G , we need to determine LE lists (compare Section 7) for a random permutation of the nodes. These are the result of an MBF-like algorithm using $\mathcal{S}_{\min,+}$ and $\mathcal{D}$; its filter r ensures that $|r(x^{(i)})_v| \in O(\log n)$ w.h.p. for all i (i.e., that intermediate results are small). This allows for performing an iteration with $\tilde{O}(m)$ work. However, doing so requires SPD(G) iterations, which in general can be as large as $n - 1$, conflicting with our goal of polylogarithmic depth.

To resolve this problem, we reduce the SPD, accepting a slight increase in stretch. The first step is to use Cohen's $(d, 1/\text{polylog } n)$ -hop set [17]: a small number of additional (weighted) edges for G , such that for all $v, w \in V$, $\text{dist}^d(v, w, G') \leq (1 + \hat{\epsilon}) \text{dist}(v, w, G)$, where G' is G augmented with the additional edges and $\hat{\epsilon} \in 1/\text{polylog } n$. Her algorithm is sufficiently efficient in terms of depth, work, and number of additional edges. Yet our problem is not solved: The d -hop distances in G' only *approximate* distances (compare Observation 1.1), but constructing FRT trees critically depends on the triangle inequality and thus on the use of *exact* distances.

In this section, we resolve this issue. After augmenting G with the hop set, we embed it into a complete graph H on the same node set so that $\text{SPD}(H) \in O(\log^2 n)$, keeping the stretch limited. Where hop sets preserve distances *exactly* and ensure the existence of *approximately* shortest paths with few hops, H preserves distances *approximately* but guarantees that we obtain *exact* shortest

⁹For this problem, we drop the assumption that graphs are connected.

paths with few hops. Note that explicitly constructing H causes $\Omega(n^2)$ work; we circumnavigate this obstacle in Section 5 with the help of the machinery developed in Section 2.

Since our construction requires to first add the hop set to G , assume for the sake of presentation that G already contains a $(d, \hat{\varepsilon})$ -hop set for fixed $\hat{\varepsilon} \in \mathbb{R}_{>0}$ and $d \in \mathbb{N}$ throughout this section. We begin our construction of H by sampling levels for the vertices V : Every vertex starts at level 0. In step $\lambda \geq 1$, each vertex in level $\lambda - 1$ is raised to level λ with probability $\frac{1}{2}$. We continue until the first step $\Lambda + 1$ where no node is sampled. $\lambda(v)$ refers to the level of $v \in V$. We define the level of an edge $e \in E$ as $\lambda(e) := \min\{\lambda(v) \mid v \in e\}$, the minimal level of its incident vertices; as H will be a complete graph, we thus have $\lambda(\{v, w\}) = \min\{\lambda(v), \lambda(w)\}$ for each $v, w \in V, v \neq w$.

LEMMA 4.1. *W.h.p., $\Lambda \in O(\log n)$.*

PROOF. For $c \in \mathbb{R}_{\geq 1}$, $v \in V$ has $\lambda(v) < c \log n$ with probability $1 - (\frac{1}{2})^{c \log n} = 1 - n^{-c}$ (i.e., w.h.p.). Lemma 1.2 yields that all nodes have a level of $O(\log n)$ w.h.p. and the claim follows. $\square$

The idea is to use the levels in the following way. We devise a complete graph H on V . An edge of H of level λ is weighted with the d -hop distance between its endpoints in G —a $(1 + \hat{\varepsilon})$ -approximation of their exact distance because G contains a $(d, \hat{\varepsilon})$ -hop set by assumption and multiplied with a penalty of $(1 + \hat{\varepsilon})^{\Lambda - \lambda}$. This way, high-level edges are “more attractive” for shortest paths because they receive smaller penalties.

Definition 4.2 (Simulated graph H). Let $G = (V, E, \omega)$ be a graph that contains a $(d, \hat{\varepsilon})$ -hop set with levels sampled as above. We define the complete graph H as

$$H := \left(V, \binom{V}{2}, \omega_\Lambda \right) \quad (76)$$

$$\omega_\Lambda(\{v, w\}) \mapsto (1 + \hat{\varepsilon})^{\Lambda - \lambda(v, w)} \text{dist}^d(v, w, G). \quad (77)$$

We formalize the notion of high-level edges being “more attractive” than low-level paths: In H , any min-hop shortest path between two nodes of level λ is exclusively comprised of edges of level λ or higher; no min-hop shortest path’s level locally decreases. Therefore, all min-hop shortest paths can be split into two subpaths, the first of monotonically increasing and the second of monotonically decreasing levels.

LEMMA 4.3. *Consider $v, w \in V, \lambda = \lambda(v, w)$, and $p \in \text{MHSP}(v, w, H)$. Then all edges of p have level at least λ .*

PROOF. The case $\lambda = 0$ is trivial. Consider $1 \leq \lambda \leq \Lambda$ and, for the sake of contradiction, let q be a nontrivial maximal subpath of p containing only edges of level strictly less than λ . Observe that $q \in \text{MHSP}(v', w', H)$ for some $v', w' \in V$ with $\lambda(v'), \lambda(w') \geq \lambda$. We have

$$\omega_\Lambda(q) \geq (1 + \hat{\varepsilon})^{\Lambda - (\lambda - 1)} \text{dist}(v', w', G). \quad (78)$$

However, the edge $e = \{v', w'\}$ has level $\lambda(v', w') \geq \lambda$ and weight

$$\omega_\Lambda(e) \leq (1 + \hat{\varepsilon})^{\Lambda - \lambda} \text{dist}^d(v', w', G) \leq (1 + \hat{\varepsilon})^{\Lambda - (\lambda - 1)} \text{dist}(v', w', G) \leq \omega_\Lambda(q) \quad (79)$$

by construction. Since $|q|$ is maximal and $\lambda(v'), \lambda(w') \geq \lambda$, q can only be a single edge of level λ or higher, contradicting the assumption. $\square$

Knowing that edge levels in min-hop shortest paths are first monotonically increasing and then monotonically decreasing, the next step is to limit the number of hops spent on each level.

LEMMA 4.4. *Consider vertices v and w of H with $\lambda(v), \lambda(w) \geq \lambda$. Then w.h.p., one of the following statements holds:*

$$\text{hop}(v, w, H) \in O(\log n) \text{ or} \quad (80)$$

$$\forall p \in \text{MHSP}(v, w, H) \exists e \in p : \lambda(e) \geq \lambda + 1. \quad (81)$$

PROOF. Condition on the event $\mathcal{E}_{V_\lambda}$ that $V_\lambda \subseteq V$ is the set of nodes with level λ or higher (with level $\lambda + 1$ not yet sampled). Let $H_\lambda := (V_\lambda, (\frac{V_\lambda}{2}, \omega_\lambda))$ with $\omega_\lambda(\{v, w\}) \mapsto (1 + \varepsilon)^{\Lambda - \lambda} \text{dist}^d(v, w, G)$ denote the subgraph of H spanned by V_λ and capped at level λ .

Consider $p \in \text{MHSP}(v, w, H_\lambda)$. Observe that $\mathbb{P}[\lambda(u) \geq \lambda + 1 \mid \mathcal{E}_{V_\lambda}] = \frac{1}{2}$ independently for all $u \in V_\lambda$, and hence $\mathbb{P}[\lambda(e) \geq \lambda + 1 \mid \mathcal{E}_{V_\lambda}] = \frac{1}{4}$ for all $e \in p$. This probability holds independently for every other edge of p . If $|p| \geq 2c \log_{4/3} n$ for some choice of $c \in \mathbb{R}_{\geq 1}$, the probability that p contains no edge of level $\lambda + 1$ or higher is bounded from above by $(\frac{3}{4})^{|p|/2} \leq (\frac{3}{4})^{c \log_{4/3} n} = n^{-c}$, so p contains such an edge w.h.p.

Fix a single arbitrary $p \in \text{MHSP}(v, w, H_\lambda)$. Let $\mathcal{E}_p$ denote the event that p fulfills $|p| \in O(\log n)$ or contains an edge of level $\lambda + 1$ or higher; as argued earlier, $\mathcal{E}_p$ occurs w.h.p. Note that we cannot directly apply the union bound to deduce a similar statement for all $q \in \text{MHSP}(v, w, H_\lambda)$: There are more than polynomially many v - w -paths. Instead, we argue that if $\mathcal{E}_p$ holds, it follows that all $q \in \text{MHSP}(v, w, H)$ must behave as claimed.

To show that all $q \in \text{MHSP}(v, w, H)$ fulfill Equations (80) or (81) under the assumption that $\mathcal{E}_p$ holds, first recall that q only uses edges of level λ or higher by Lemma 4.3. Furthermore, observe that $\omega_\Lambda(q) \leq \omega_\Lambda(p)$, as q is a shortest path with respect to ω_Λ . If q contains an edge of level $\lambda + 1$ or higher, Equation (81) holds for q . Otherwise, we have $\omega_\lambda(q) = \omega_\Lambda(q)$, and distinguish two cases:

Case 1 ($|p| \in O(\log n)$): We have

$$\omega_\Lambda(p) \leq \omega_\lambda(p) \leq \omega_\lambda(q) = \omega_\Lambda(q), \quad (82)$$

so $\omega_\Lambda(q) = \omega_\Lambda(p)$ and $|q| \leq |p| \in O(\log n)$ follows from $q \in \text{MHSP}(v, w, H)$.

Case 2 (p contains an edge of level $\lambda + 1$ or higher): This yields $\omega_\Lambda(p) < \omega_\lambda(p)$, implying

$$\omega_\Lambda(p) < \omega_\lambda(p) \leq \omega_\lambda(q) = \omega_\Lambda(q), \quad (83)$$

which contradicts $q \in \text{MHSP}(v, w, H)$.

So far, we condition on $\mathcal{E}_{V_\lambda}$. In order to remove this restriction, let $\mathcal{E}_{vw}$ denote the event that Equation (80) or Equation (81) holds for $v, w \in V$. The above case distinction shows that $\mathbb{P}[\mathcal{E}_{vw} \mid \mathcal{E}_{V_\lambda}] \geq 1 - n^{-c}$ for an arbitrary $c \in \mathbb{R}_{\geq 1}$. We conclude that

$$\mathbb{P}[\mathcal{E}_{vw} \mid \lambda(v, w) \geq \lambda] = \sum_{V_\lambda \subseteq V} \mathbb{P}[\mathcal{E}_{V_\lambda} \mid \lambda(v, w) \geq \lambda] \mathbb{P}[\mathcal{E}_{vw} \mid \mathcal{E}_{V_\lambda}] \quad (84)$$

$$= \sum_{\{v, w\} \subseteq V_\lambda \subseteq V} \mathbb{P}[\mathcal{E}_{V_\lambda} \mid \lambda(v, w) \geq \lambda] \mathbb{P}[\mathcal{E}_{vw} \mid \mathcal{E}_{V_\lambda}] \quad (85)$$

$$\geq \sum_{\{v, w\} \subseteq V_\lambda \subseteq V} \mathbb{P}[\mathcal{E}_{V_\lambda} \mid \lambda(v, w) \geq \lambda] (1 - n^{-c}) \quad (86)$$

$$= (1 - n^{-c}) \sum_{\{v, w\} \subseteq V_\lambda \subseteq V} \mathbb{P}[\mathcal{E}_{V_\lambda} \mid \lambda(v, w) \geq \lambda] \quad (87)$$

$$= 1 - n^{-c}, \quad (88)$$

which is the statement of the lemma. $\square$

We argue above that any min-hop shortest path in H traverses every level at most twice, Lemma 4.4 states that each such traversal, w.h.p., only has a logarithmic number of hops, and Lemma 4.1 asserts that, w.h.p., there are only logarithmically many levels. Together, this means that min-hop shortest paths in H have $O(\log^2 n)$ hops w.h.p. Additionally, our construction limits the stretch of shortest paths in H as compared to G by $(1 + \hat{\varepsilon})^{\Lambda+1}$ (i.e., by $(1 + \hat{\varepsilon})^{O(\log n)}$) w.h.p.

THEOREM 4.5. *W.h.p., $\text{SPD}(H) \in O(\log^2 n)$ and, for all $v, w \in V$,*

$$\text{dist}(v, w, G) \leq \text{dist}(v, w, H) \leq (1 + \hat{\varepsilon})^{O(\log n)} \text{dist}(v, w, G). \quad (89)$$

PROOF. Fix a level λ . Any fixed pair of vertices of level λ or higher fulfills, w.h.p., Equation (80) or (81) by Lemma 4.4. Since there are at most $\binom{n}{2}$ such pairs, w.h.p., all of them fulfill Equation (80) or (81) by Lemma 1.2.

Let $\mathcal{E}_{\log}$ denote the event that there is no higher level than $\Lambda \in O(\log n)$, which holds w.h.p. by Lemma 4.1. Furthermore, let $\mathcal{E}_\lambda$ denote the event that all pairs of vertices of level λ or higher fulfill Equation (80) or (81), which holds w.h.p. as argued above. Then $\mathcal{E} := \mathcal{E}_{\log} \cap \mathcal{E}_0 \cap \dots \cap \mathcal{E}_\Lambda$ holds w.h.p. by Lemma 1.2.

Condition on $\mathcal{E}$; in particular, no min-hop shortest path whose edges all have the same level has more than $O(\log n)$ hops. Consider some min-hop shortest path p in H . By Lemma 4.3, p has two parts: The edge level monotonically increases in the first and monotonically decreases in the second part. Hence, p can be split up into at most $2\Lambda - 1$ segments, in each of which all edges have the same level. As this holds for all min-hop shortest paths, we conclude that $\text{SPD}(H) \in O(\Lambda \log n) \subseteq O(\log^2 n)$ w.h.p., as claimed.

As for Inequality (4.14), recall that H is constructed from $G = (V, E, \omega)$ and that G contains a $(d, \hat{\varepsilon})$ -hop set. For all $v, w \in V$, we have

$$\text{dist}(v, w, H) \leq \omega_\Lambda(v, w) \leq (1 + \hat{\varepsilon})^\Lambda \text{dist}^d(v, w, G) \leq (1 + \hat{\varepsilon})^{\Lambda+1} \text{dist}(v, w, G) \quad (90)$$

by construction of H . Recalling that $\Lambda \in O(\log n)$ due to $\mathcal{E}$ completes the proof. $\square$

We use Cohen's construction to obtain a $(d, \hat{\varepsilon})$ -hop set with $\hat{\varepsilon} \in 1/\text{polylog } n$, where the exponent of $\text{polylog } n$ is under our control [17]. A sufficiently large exponent yields $(1 + \hat{\varepsilon})^{O(\log n)} \subseteq e^{\hat{\varepsilon} O(\log n)} \subseteq e^{1/\text{polylog } n} = 1 + 1/\text{polylog } n$, upper-bounding Inequality (4.14) by

$$\text{dist}(v, w, G) \leq \text{dist}(v, w, H) \in (1 + 1/\text{polylog } n) \text{dist}(v, w, G) \subseteq (1 + o(1)) \text{dist}(v, w, G). \quad (91)$$

To wrap things up: Given a weighted graph G , we augment G with a $(d, 1/\text{polylog } n)$ -hop set. After that, the d -hop distances in G approximate the actual distances in G , but these approximations may violate the triangle inequality. We fix this by embedding into H , using geometrically sampled node levels and an exponential penalty on the edge weights with decreasing levels. Since H is a complete graph, explicitly constructing it is prohibitively costly in terms of work. The next section shows how to avoid this issue by efficiently simulating MBF-like algorithms on H .

5 AN ORACLE FOR MBF-LIKE QUERIES

Given a weighted graph G and $\hat{\varepsilon} \in 1/\text{polylog } n$, Section 4 introduces a complete graph H that $(1 + o(1))$ -approximates the distances of G and w.h.p. has a polylogarithmic SPD using a $(d, \hat{\varepsilon})$ -hop set. H would solve our problem, but we cannot explicitly write H into memory as this requires an unacceptable $\Omega(n^2)$ work.

Instead, we dedicate this section to an *oracle that answers MBF-like queries*; that is, to an oracle that, given a weighted graph G , an MBF-like algorithm $\mathcal{A}$, and a number of iterations h , returns $\mathcal{A}^h(H)$. Note that while the oracle can answer distance queries in polylogarithmic depth (when, e.g., queried by SSSP, k -SSP, or APSP), MBF-like queries are more general (compare Section 3)

and allow for more work-efficient algorithms (as in Sections 6 and 7). The properties of MBF-like algorithms discussed in Section 2 allow the oracle to internally work on G and simulate iterations of $\mathcal{A}$ on H using d (i.e., polylogarithmically many) iterations on G .

Throughout this section, we denote by A_G and A_H the adjacency matrices of G and H , respectively. Furthermore, we fix the semiring to be $\mathcal{S}_{\min,+}$, since we explicitly calculate distances; generalizations to other semirings are possible but require appropriate generalizations of adjacency matrices and hence obstruct presentation.

We establish this section's results in two steps: Section 5.1 derives a representation of A_H in terms of A_G , which is then used to efficiently implement the oracle in Section 5.2. The oracle is used to approximate the metric of G in Section 6 and to construct an FRT tree in Section 7, both with polylogarithmic depth.

5.1 Decomposing H

The idea is to simulate each iteration of an MBF-like algorithm $\mathcal{A}$ on H using d iterations on G . This is done for each level $\lambda \in \{0, \dots, \Lambda\}$ in parallel. For level λ , we run $\mathcal{A}$ for d iterations on G with edge weights scaled up by $(1 + \varepsilon)^{\Lambda - \lambda}$, where the initial vector is obtained by discarding all information at nodes of level smaller than λ . Afterward, we again discard everything stored at vertices with a level smaller than λ . Since $(A_G^d)_{vw} = \text{dist}^d(v, w, G)$, this ensures that we propagate information between nodes $v, w \in V$ with $\lambda(v, w) = \lambda$ with the corresponding edge weight while discarding any exchange between nodes with $\lambda(v, w) < \lambda$ (which is handled by the respective parallel run). While we also propagate information between v and w if $\lambda(v, w) > \lambda$ (over too long a distance because edge weights are scaled by $(1 + \varepsilon)^{\Lambda - \lambda} > (1 + \varepsilon)^{\Lambda - \lambda(v, w)}$) the parallel run for $\lambda(v, w)$ correctly propagates values. Therefore, aggregating the results of all levels (i.e., applying $\oplus$, the source-wise minimum) and applying r^V completes the simulation of an iteration of $\mathcal{A}$ on H .

This approach resolves two complexity issues. First, we multiply (polylogarithmically often) with A_G , which, as opposed to the dense A_H , has $O(m)$ non- ∞ entries only. Second, Corollary 2.17 shows that we are free to filter using r^V at any time, keeping the entries of intermediate state vectors small.

We formalize the preceding intuition. Recall that

$$(A_H)_{vw} = \omega_\Lambda(v, w) = (1 + \varepsilon)^{\Lambda - \lambda(v, w)} \text{dist}^d(v, w, G) = (1 + \varepsilon)^{\Lambda - \lambda(v, w)} (A_G^d)_{vw}. \quad (92)$$

For $\lambda \in \{0, \dots, \Lambda\}$, denote by P_λ the $\mathcal{M}^V$ -projection to coordinates $V_\lambda := \{v \in V \mid \lambda(v) \geq \lambda\}$:

$$(P_\lambda x)_v := \begin{cases} x_v & \text{if } \lambda(v) \geq \lambda \text{ and} \\ \perp & \text{otherwise.} \end{cases} \quad (93)$$

Observe that P_λ is an SLF on $\mathcal{M}^V$, where $(P_\lambda)_{vw} = 0$ if $v = w \in V_\lambda$ and $(P_\lambda)_{vw} = \infty$ otherwise. This gives us the tools to decompose A_H as motivated earlier.

LEMMA 5.1. *With $(A_\lambda)_{vw} := (1 + \varepsilon)^{\Lambda - \lambda} (A_G)_{vw}$ (with respect to multiplication in $\mathbb{R}$, not $\odot$), we have*

$$A_H = \bigoplus_{\lambda=0}^{\Lambda} P_\lambda A_\lambda^d P_\lambda. \quad (94)$$

PROOF. Since $(A_G^d)_{vw} = \text{dist}^d(v, w, G)$, it holds that $(A_\lambda^d)_{vw} = (1 + \varepsilon)^{\Lambda - \lambda} \text{dist}^d(v, w, G)$. Therefore,

$$(A_\lambda^d P_\lambda)_{vw} = \min_{u \in V} \{ (A_\lambda^d)_{vu} + (P_\lambda)_{uw} \} = \begin{cases} (1 + \varepsilon)^{\Lambda - \lambda} \text{dist}^d(v, w, G) & \text{if } w \in V_\lambda \text{ and} \\ \infty & \text{otherwise,} \end{cases} \quad (95)$$

and hence

$$(P_\lambda A_\lambda^d P_\lambda)_{vw} = \min_{u \in V} \{(P_\lambda)_{vu} + (A_\lambda^d P_\lambda)_{uw}\} = \begin{cases} (1 + \hat{\varepsilon})^{\Lambda-\lambda} \text{dist}^d(v, w, G) & \text{if } v, w \in V_\lambda \text{ and} \\ \infty & \text{otherwise.} \end{cases} \quad (96)$$

We conclude that

$$\left(\bigoplus_{\lambda=0}^{\Lambda} P_\lambda A_\lambda^d P_\lambda \right)_{vw} = \min \left\{ (1 + \hat{\varepsilon})^{\Lambda-\lambda} \text{dist}^d(v, w, G) \mid \lambda \in \{0, \dots, \lambda(v, w)\} \right\} \quad (97)$$

$$= (1 + \hat{\varepsilon})^{\Lambda-\lambda(v, w)} \text{dist}^d(v, w, G) \quad (98)$$

$$= (A_H)_{vw}. \quad \square$$

Having decomposed A_H , we continue with $\mathcal{A}^h(H)$, taking the freedom to apply filters intermediately. For all $h \in \mathbb{N}$, we have

$$A_H^h \stackrel{(5.3)}{=} \left(\bigoplus_{\lambda=0}^{\Lambda} P_\lambda A_\lambda^d P_\lambda \right)^h \stackrel{(2.35)}{\sim} \left(r^V \left(\bigoplus_{\lambda=0}^{\Lambda} P_\lambda (r^V A_\lambda)^d P_\lambda \right) \right)^h r^V, \quad (99)$$

and hence

$$\mathcal{A}^h(H) = r^V A_H^h x^{(0)} \stackrel{(2.11), (5.8)}{=} \left(r^V \left(\bigoplus_{\lambda=0}^{\Lambda} P_\lambda (r^V A_\lambda)^d P_\lambda \right) \right)^h r^V x^{(0)}. \quad (100)$$

Observe that we can choose $h = \text{SPD}(H) \in O(\log^2 n)$ w.h.p. by Theorem 4.5 and recall that $d \in \text{polylog } n$. Overall, this allows us to determine $\mathcal{A}(H)$ with polylogarithmic depth and $\tilde{O}(m)$ work provided we can implement the individual steps (see below) at this complexity.

5.2 Implementing the Oracle

The oracle determines iterations of $\mathcal{A}$ on H using iterations on G while only introducing a polylogarithmic overhead with respect to iterations in G . With the decomposition from Lemma 5.1 at hand, it can be implemented as follows.

Given a state vector $x^{(i)} \in \mathcal{M}^V$, simulate one iteration of $\mathcal{A}$ on H for edges of level λ ; that is, determine $y_\lambda := P_\lambda (r^V A_\lambda)^d P_\lambda x^{(i)}$ by

- discarding entries at nodes of a level smaller than λ ,
- running d iterations of $\mathcal{A}$ with distances stretched by $(1 + \hat{\varepsilon})^{\Lambda-\lambda}$ on G , applying the filter after each iteration, and
- again discarding entries at nodes with levels smaller than λ .

After running this procedure in parallel for all $0 \leq \lambda \leq \Lambda$, perform the $\oplus$ -operation and apply the filter; that is, for each node $v \in V$ determine $x_v^{(i+1)} = r(\bigoplus_{\lambda=0}^{\Lambda} y_\lambda)_v$.

THEOREM 5.2 (ORACLE). *Consider an MBF-like algorithm $\mathcal{A}$ using a semimodule $\mathcal{M}$ over the semiring $\mathcal{S}_{\min,+}$, the representative projection $r : \mathcal{M} \rightarrow \mathcal{M}$, and the initialization $x^{(0)} \in \mathcal{M}^V$. If we can, for all $1 \leq f \leq d$, $1 \leq i \leq h$, and $0 \leq \lambda \leq \Lambda$, where we may assume $\Lambda \in O(\log n)$,*

- (1) *compute $r^V x^{(0)}$ from $x^{(0)}$ with depth D and work W ,*
- (2) *determine $r^V A_\lambda y$ from any intermediate state vector $y = (r^V A_\lambda)^{f-1} P_\lambda x^{(i-1)}$ —corresponding to the f th iteration with respect to A_λ starting at state $x^{(i-1)}$ —with depth D and work W , and*

- (3) compute $r^V(\bigoplus_{\lambda=0}^{\Lambda} y_{\lambda})$ from the individual $y_{\lambda} = (P_{\lambda} r^V A_{\lambda}^d P_{\lambda}) x^{(i-1)}$, $0 \leq \lambda \leq \Lambda$ (reflecting aggregation over all levels to complete a simulated iteration in H) using depth $D_{\oplus}$ and work $W_{\oplus}$,

then we can w.h.p.

- (1) determine $\mathcal{A}^h(H)$ using $O((dW \log n + W_{\oplus})h) \subseteq \tilde{O}((dW + W_{\oplus})h)$ work and a depth of $O((dD + D_{\oplus})h)$ and thus
 (2) calculate $\mathcal{A}(H)$ using $O((dW \log n + W_{\oplus}) \log^2 n) \subseteq \tilde{O}(dW + W_{\oplus})$ work and a depth of $O((dD + D_{\oplus}) \log^2 n) \subseteq \tilde{O}(dD + D_{\oplus})$.

PROOF. Condition on $\Lambda \in O(\log n)$ and $\text{SPD}(H) \in O(\log^2 n)$; both events occur w.h.p. by Lemma 4.1 and Theorem 4.5. By Equation (100), we have to compute

$$\mathcal{A}^h(H) = \left(r^V \left(\bigoplus_{\lambda=0}^{\Lambda} P_{\lambda} (r^V A_{\lambda})^d P_{\lambda} \right) \right)^h r^V x^{(0)}. \quad (101)$$

Concerning P_{λ} , note that we can evaluate $(P_{\lambda} y)_{v \in V}$ lazily; that is, we can determine whether $(P_{\lambda} y)_v$ evaluates to $\perp$ or to y_v only if it is accessed. Thus, the total work and depth required increase by at most a constant factor due to all applications of P_{λ} . Together with the prerequisites, this means that $(r^V A_{\lambda} P_{\lambda}) y$ can be determined in $O(W)$ work and $O(D)$ depth and that evaluating $P_{\lambda} (r^V A_{\lambda})^d P_{\lambda} y$ sequentially in d requires $O(dW)$ work and $O(dD)$ depth.

The set of summands of $\bigoplus_{\lambda=0}^{\Lambda} P_{\lambda} (r^V A_{\lambda})^d P_{\lambda} y$ can be determined using $O(\Lambda dW)$ work and $O(dD)$ depth since this is independent for each λ . Performing the aggregation and applying the filter is possible in $D_{\oplus}$ depth and $W_{\oplus}$ work by assumption. We arrive at $O(\Lambda dW + W_{\oplus})$ work and $O(dD + D_{\oplus})$ depth for determining $x^{(i)} = r^V \bigoplus_{\lambda=0}^{\Lambda} P_{\lambda} (r^V A_{\lambda})^d P_{\lambda} x^{(i-1)}$ from $x^{(i-1)}$.

Sequentially iterating this h times to determine $\mathcal{A}^h(H)$ increases work and depth by a factor of h , yielding $O((\Lambda dW + W_{\oplus})h)$ work and $O((dD + D_{\oplus})h)$ depth. Computing $r^V x^{(0)}$ requires work W and depth D by the prerequisites and does not change the asymptotic complexity accumulated so far. We arrive at $O((dW \log n + W_{\oplus})h)$ work and $O((dD + D_{\oplus})h)$ depth, which is the first claim; $\text{SPD}(H) \in O(\log^2 n)$ yields the second claim. As we only condition on two events that occur w.h.p., this concludes the proof by Lemma 1.2. $\square$

6 APPROXIMATE METRIC CONSTRUCTION

As a consequence of the machinery in Section 5, we can efficiently determine approximate metrics. In fact, our metric approximations are fast enough to improve the state of the art regarding the FRT embedding from $\tilde{O}(n^3)$ to $\tilde{O}(n^{2+\varepsilon})$ work when combined with a result of Blelloch et al. [14]; see below.

An approximate metric is stronger than approximate distances (i.e., stronger than hop sets) as it requires consistency with the triangle inequality:

Definition 6.1 (Approximate Metric). Consider $\alpha \in \mathbb{R}_{\geq 1}$ and a metric $d : V \times V \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ on V . We refer to d' as α -approximate metric of d if

- (1) d' is a metric on V and
 (2) for all $v, w \in V$,

$$d(v, w) \leq d'(v, w) \leq \alpha d(v, w). \quad (102)$$

We can determine a $(1 + o(1))$ -approximate metric of $\text{dist}(\cdot, \cdot, G)$ for an arbitrary graph G by querying the oracle with APSP on H using polylogarithmic depth and $\tilde{O}(nm^{1+\varepsilon})$ work. This is much

more work-efficient on sparse graphs than the naive approach using $O(n^3 \log n)$ work (squaring the adjacency matrix $\lceil \log_2 n \rceil$ times) for obtaining $\text{dist}(\cdot, \cdot, G)$ exactly.

THEOREM 6.2 ($((1 + o(1))$ -APPROXIMATE METRIC). *Given a weighted graph $G = (V, E, \omega)$ and a constant $\varepsilon > 0$, we can w.h.p. compute, using $\tilde{O}(n(m + n^{1+\varepsilon}))$ work and $\text{polylog } n$ depth, a $(1 + 1/\text{polylog } n)$ -approximate metric of $\text{dist}(\cdot, \cdot, G)$ on V .*

PROOF. First augment G with a $(d, 1/\text{polylog } n)$ -hop set using $\tilde{O}(m^{1+\varepsilon})$ work and $\text{polylog } n$ depth with $d \in \text{polylog } n$ using Cohen's hop set construction [17]. The resulting graph has $\tilde{O}(m + n^{1+\varepsilon})$ edges. An iteration of APSP (compare Example 3.5) incurs $O(\log n)$ depth and $O(\delta_v n \log n)$ work at a node v of degree δ_v by Lemma 2.3. Hence, $D \in O(\log n)$ depth and $W \in O(\sum_{v \in V} \delta_v n \log n) \subseteq \tilde{O}(n(m + n^{1+\varepsilon}))$ work suffice for an entire iteration. Aggregation and filtering over the individual y_λ , $0 \leq \lambda \leq \Lambda \in O(\log n)$, takes $D_\oplus \in O(\log n)$ depth and $W_\oplus \in O(n^2 \log^2 n)$ work. The trivial filter $r^V = \text{id}$ does not induce any overhead. By Theorem 5.2, we can w.h.p. simulate SPD(H) iterations of APSP on H using $\tilde{O}(n(m + n^{1+\varepsilon}))$ work and $\tilde{O}(1)$ depth. By Theorem 4.5 and Equation (91), this yields a metric which $(1 + 1/\text{polylog } n)$ -approximates $\text{dist}(\cdot, \cdot, G)$. $\square$

Using the sparse spanner algorithm of Baswana and Sen [9], we can obtain a metric with a different work-approximation tradeoff. Note that this is near-optimal in terms of work due to the trivial lower bound of $\Omega(n^2)$ for writing down the solution.

THEOREM 6.3 ($O(1)$ -APPROXIMATE METRIC). *For a weighted graph $G = (V, E, \omega)$ and a constant $\varepsilon > 0$, we can w.h.p. compute an $O(1)$ -approximate metric of $\text{dist}(\cdot, \cdot, G)$ on V using $\tilde{O}(n^{2+\varepsilon})$ work and $\text{polylog } n$ depth.*

PROOF. Baswana and Sen show how to compute a $(2k - 1)$ -spanner of $G = (V, E, \omega)$; that is, $E' \subseteq E$ such that $G' := (V, E', \omega)$ fulfills, for all $v, w \in V$,

$$\text{dist}(v, w, G) \leq \text{dist}(v, w, G') \leq (2k - 1) \text{dist}(v, w, G), \quad (103)$$

using $\tilde{O}(1)$ depth and $\tilde{O}(m)$ work. Baswana and Sen argue that $|E'| \in O(kn^{1+1/k})$ in expectation [9], and we obtain $|E'| \in O(kn^{1+1/k} \log n) \subseteq \tilde{O}(kn^{1+1/k})$ w.h.p., as, for example, argued in Appendix A of Becker et al. [10]. Furthermore, without loss of generality $k \in O(\log n)$, because $kn^{1/k} = k2^{\log n/k}$ starts growing beyond that point. This results in $|E'| \in \tilde{O}(n^{1+1/k})$ w.h.p.

We compute an $O(1)$ -approximate metric as follows. (1) Compute a $(2k - 1)$ -spanner for $k = \lceil 1/\varepsilon \rceil$. This is possible within the given bounds on work and depth, yielding $|E'| \in \tilde{O}(n^{1+1/k}) = \tilde{O}(n^{1+\varepsilon})$ edges w.h.p. and a stretch that is constant with respect to n and m . (2) Apply Theorem 6.2 to $G' := (V, E', \omega)$ and ε . This induces $\tilde{O}(1)$ depth and $\tilde{O}(n^{2+\varepsilon})$ work.

By construction, the resulting metric has stretch $(2k - 1)(1 + o(1)) \subseteq O(1)$. $\square$

Blelloch et al. [14] show how to construct an FRT tree from a metric using $O(n^2)$ work and $O(\log^2 n)$ depth. Combining this with Theorem 6.3 enables us to w.h.p. construct an FRT tree from a graph G using polylogarithmic depth and $\tilde{O}(n^{2+\varepsilon})$ work. This already improves upon the state of the art of using $\tilde{O}(n^3)$ work to compute $\text{dist}(\cdot, \cdot, G)$ exactly and then applying the algorithm of Blelloch et al. [14]. We can, however, achieve this even more efficiently on sparse graphs: Constructing FRT trees is an MBF-like algorithm and solving the problem directly—using the oracle—reduces the work to $\tilde{O}(m^{1+\varepsilon})$; this is the goal of Section 7.

7 FRT CONSTRUCTION

Given a weighted graph G , determining a metric that $O(1)$ -approximates $\text{dist}(\cdot, \cdot, G)$ (using polylogarithmic depth and $\tilde{O}(n^{2+\varepsilon})$ work) is straightforward; see Theorem 6.3. The oracle is queried with the MBF-like APSP algorithm, implicitly enjoying the benefits of the SPD-reducing sampling

technique of Section 4. In this section, we show that collecting the information required to construct FRT trees—LE lists—is an MBF-like algorithm; that is, a query that can be directly answered by the oracle. Since collecting LE lists is more work-efficient than APSP, this leads to our main result: w.h.p. sampling from the FRT distribution using polylogarithmic depth and $\tilde{O}(m^{1+\varepsilon})$ work.

We begin with a formal definition of metric (tree) embeddings in general and the FRT embedding in particular in Section 7.1, proceed to show that the underlying algorithm is MBF-like (Section 7.2) and that all intermediate steps are sufficiently efficient in terms of depth and work (Section 7.3), and present our main results in Section 7.4. Section 7.5 describes how to retrieve the original paths in G that correspond to the edges of the sampled FRT tree.

7.1 Metric Tree Embeddings

We use this section to introduce the (distribution over) metric tree embeddings of Fakcharoenphol, Rao, and Talwar, referred to as FRT embedding, which has expected stretch $O(\log n)$ [23].

Definition 7.1 (Metric Embedding). Let $G = (V, E, \omega)$ be a graph. A *metric embedding of stretch α* of G is a graph $G' = (V', E', \omega')$, such that $V \subseteq V'$ and

$$\forall v, w \in V : \quad \text{dist}(v, w, G) \leq \text{dist}(v, w, G') \leq \alpha \text{dist}(v, w, G), \quad (104)$$

for some $\alpha \in \mathbb{R}_{\geq 1}$. If G' is a tree, we refer to it as *metric tree embedding*. For a random distribution of metric embeddings G' , we require $\text{dist}(v, w, G) \leq \text{dist}(v, w, G')$ and define the *expected stretch* as

$$\alpha := \max_{v \neq w \in V} \left\{ \mathbb{E} \left[\frac{\text{dist}(v, w, G')}{\text{dist}(v, w, G)} \right] \right\}. \quad (105)$$

We show how to efficiently sample from the FRT distribution for the graph H introduced in Section 4. As H is an embedding of G with a stretch in $1 + o(1)$, this results in a tree embedding of G of stretch $O(\log n)$. Khan et al. [30] show that a suitable representation of (a tree sampled from the distribution of) the FRT embedding [23] can be constructed as follows.

- (1) Choose $\beta \in [1, 2)$ uniformly at random.
- (2) Choose uniformly at random a total order of the nodes (i.e., a uniformly random permutation). In the following, $v < w$ means that v is smaller than w with respect to this order.
- (3) Determine for each node $v \in V$ its *LE list*: This is the list obtained by deleting from $\{(\text{dist}(v, w, H), w) \mid w \in V\}$ all pairs $(\text{dist}(v, w, H), w)$ for which there is some $u \in V$ with $\text{dist}(v, u, H) \leq \text{dist}(v, w, H)$ and $u < w$. Essentially, v learns, for every distance d , the smallest node within distance at most d , i.e., $\min\{w \in V \mid \text{dist}(v, w, G) \leq d\}$.
- (4) Denote by $\omega_{\min} := \min_{e \in E} \{\omega(e)\}$ and $\omega_{\max} := \max_{e \in E} \{\omega(e)\}$ the minimum and maximum edge weight, respectively; recall that $\omega_{\max} / \omega_{\min} \in \text{poly } n$ by assumption. From the LE lists, determine for each $v \in V$ and distance $\beta 2^i \in [\omega_{\min} / 2, 2\omega_{\max}]$, $i \in \mathbb{Z}$, the node $v_i := \min\{w \in V \mid \text{dist}(v, w, H) \leq \beta 2^i\}$. Without loss of generality, we assume that $i \in \{0, \dots, k\}$ for $k \in O(\log n)$ (otherwise, we shift the indices of the nodes v_i accordingly). Hence, for each $v \in V$, we obtain a sequence of nodes $(v_0, v_1, \dots, v_k)$. $(v_0, v_1, \dots, v_k)$ is the leaf corresponding to $v = v_0$ of the tree embedding, $(v_1, \dots, v_k)$ is its parent, and so on; the root is (v_k) . The edge from $(v_i, \dots, v_k)$ to $(v_{i+1}, \dots, v_k)$ has weight $\beta 2^i$.

We refer to Ghaffari and Lenzen [26] for a more detailed summary.

The preceding procedure implicitly specifies a random distribution over tree embeddings with expected stretch $O(\log n)$ [23], which we call the *FRT distribution*. We refer to following the procedure 1–4 as *sampling* from the FRT distribution. Once the randomness is fixed (i.e., steps 1–2 are

completed), the tree resulting from steps 3–4 is unique; we refer to them as *constructing an FRT tree*.

The next lemma shows that step 4 (i.e., constructing the FRT tree from the LE lists) is easy.

LEMMA 7.2. *Given LE lists of length $O(\log n)$ for all vertices, the corresponding FRT tree can be determined using $O(n \log^3 n)$ work and $O(\log^2 n)$ depth.*

PROOF. Determining $\omega_{\max}$, $\omega_{\min}$, and the range of indices i is straightforward at this complexity, as is sorting of each node's list in ascending order with respect to distance. Note that in each resulting list of distance–node pairs, the nodes are strictly decreasing in terms of the random order on the nodes, and each list ends with an entry for the minimal node. For each node v and entry (d, u) in its list in parallel, we determine the values of $i \in \{0, \dots, k\}$ such that u is the smallest node within distance $\beta 2^i$ of v . This is done by reading the distance value d' of the next entry of the list (using $d' = \beta 2^k + 1$ if (d, u) is the last entry) and writing to memory $v_i = u$ for each i satisfying that $d \leq \beta 2^i < d'$. Since $\omega_{\max} / \omega_{\min} \in \text{poly } n$, this has depth $O(\log n)$ and a total work of $O(n \log^2 n)$.

Observe that we computed the list $(v_0, \dots, v_k)$ for each $v \in V$. Recall that the ancestors of the leaf $(v_0, \dots, v_k)$ are determined by its k suffixes. It remains to remove duplicates wherever nodes share a parent. To this end, we sort the list (possibly with duplicates) of $(k + 1)n \in O(n \log n)$ suffixes (each with $O(\log n)$ entries) lexicographically, requiring $O(n \log^3 n)$ work and depth $O(\log^2 n)$, as comparing two suffixes requires depth and work $O(\log n)$. Then duplicates can be removed by comparing each key to its successor in the sorted sequence, taking another $O(n \log^2 n)$ work and $O(\log n)$ depth.

Note that tree edges and their weights are encoded implicitly as the parent of each node is given by removing the first node from the list, and the level of a node (and thus the edge to its parent) is given by the length of the list representing it. If required, it is thus trivial to determine, for example, an adjacency list with $O(n \log^2 n)$ work and depth $O(\log^2 n)$. Overall, we spent $O(n \log^3 n)$ work at $O(\log^2 n)$ depth. $\square$

7.2 Computing LE Lists is MBF-like

Picking β is trivial and choosing a random order of the nodes can be done w.h.p. by assigning to each node a string of $O(\log n)$ uniformly and independently chosen random bits. Hence, in the following, we assume this step to be completed, without loss of generality, resulting in a random assignment of the vertex IDs $\{1, \dots, n\}$. It remains to establish how to efficiently compute LE lists.

We establish that LE lists can be computed by an MBF-like algorithm (compare Definition 2.11) using the parameters in Definition 7.3; the claim that Equations (106) and (107) define a representative projection and a congruence relation is shown in Lemma 7.5.

Definition 7.3. For constructing LE lists, use the semiring $\mathcal{S} = \mathcal{S}_{\min,+}$ and the distance map $\mathcal{M} = \mathcal{D}$ from Definition 2.1 as zero-preserving semimodule. For all $x \in \mathcal{D}$, define

$$r(x)_v := \begin{cases} \infty & \exists w < v : x_w \leq x_v \text{ and} \\ x_v & \text{otherwise, and} \end{cases} \quad (106)$$

$$x \sim y \quad :\Leftrightarrow \quad r(x) = r(y) \quad (107)$$

as representative projection and congruence relation, respectively. As initialization $x^{(0)} \in \mathcal{D}^V$ use

$$x_{vw}^{(0)} := \begin{cases} 0 & \text{if } v = w \text{ and} \\ \infty & \text{otherwise.} \end{cases} \quad (108)$$

Hence, $r(x)$ is the LE list of $v \in V$ if $x_w = \text{dist}(v, w, H)$ for all $w \in V$, and we consider two lists equivalent if and only if they result in the same LE list. This allows us to prepare the proof that

retrieving LE lists can be done by an MBF-like algorithm in the following lemma. It states that filtering keeps the relevant information: If a node–distance pair is dominated by an entry in a distance map, the filtered distance map also contains a—possibly different—dominating entry.

LEMMA 7.4. *Consider arbitrary $x, y \in \mathcal{D}$, $v \in V$, and $s \in \mathbb{R}_{\geq 0} \cup \{\infty\}$. Then*

$$\exists w < v : x_w \leq s \quad \Leftrightarrow \quad \exists w < v : r(x)_w \leq s \quad (109)$$

PROOF. Observe that the necessity “ $\Leftarrow$ ” is trivial. As for sufficiency “ $\Rightarrow$,” suppose that there is $w < v$ such that $x_w \leq s$. If $r(x)_w = x_w$, we are done. Otherwise, there must be some $u < w < v$ satisfying $x_u \leq x_w \leq s$. Since $|V|$ is finite, an inductive repetition of the argument yields that there is some $w' < v$ with $r(x)_{w'} = x_{w'} \leq s$. $\square$

Equipped with this lemma, we can prove that $\sim$ is a congruence relation on $\mathcal{D}$ with representative projection r . We say that a node–distance pair (v, d) *dominates* (v', d') if and only if $v < v'$ and $d \leq d'$; in the context of $x \in \mathcal{D}$, we say that x_w dominates x_v if and only if (w, x_w) dominates (v, x_v) .

LEMMA 7.5. *The equivalence relation $\sim$ from Equation (107) of Definition 7.3 is a congruence relation. The function r from Equation (106) Definition 7.3 is a representative projection with respect to $\sim$.*

PROOF. Trivially, r is a projection (i.e., $r^2(x) = r(x)$ for all $x \in \mathcal{D}$). By Lemma 2.8, it hence suffices to show that Equations (18) and (19) hold. In order to do that, let $s \in \mathcal{S}_{\min,+}$ be arbitrary, and $x, x', y, y' \in \mathcal{D}$ such that $r(x) = r(x')$ and $r(y) = r(y')$. As we have $x_v \leq x_w \Leftrightarrow s + x_v \leq s + x_w$ for all $v, w \in V$, Equation (18) immediately follows from Equation (109).

Regarding Equation (19), we show that

$$r(x \oplus y) = r(r(x) \oplus r(y)) \quad (110)$$

which implies Equation (19) due to $r(x \oplus y) = r(r(x) \oplus r(y)) = r(r(x') \oplus r(y')) = r(x' \oplus y')$. Let $v \in V$ be an arbitrary vertex and observe that $(x \oplus y)_v$ is dominated if and only if

$$\exists w < v : (x \oplus y)_w \leq (x \oplus y)_v \quad (111)$$

$$\Leftrightarrow \exists w < v : \min\{x_w, y_w\} \leq (x \oplus y)_v \quad (112)$$

$$\Leftrightarrow \exists w < v : x_w \leq (x \oplus y)_v \vee y_w \leq (x \oplus y)_v \quad (113)$$

$$\stackrel{(7.6)}{\Leftrightarrow} \exists w < v : r(x)_w \leq (x \oplus y)_v \vee r(y)_w \leq (x \oplus y)_v. \quad (114)$$

In order to show Equation (110), we distinguish two cases.

Case 1 ($(x \oplus y)_v$ is dominated): By Definition 7.3, we have $r(x \oplus y)_v = \infty$. Additionally, we know that $(r(x) \oplus r(y))_v = \min\{r(x)_v, r(y)_v\} \geq \min\{x_v, y_v\} = (x \oplus y)_v$ must be dominated due to Equation (114), and hence $r(r(x) \oplus r(y))_v = \infty = r(x \oplus y)_v$.

Case 2 ($(x \oplus y)_v$ is not dominated): This means that, by Definition 7.3, $r(x \oplus y)_v = (x \oplus y)_v = \min\{x_v, y_v\}$. Furthermore, the negation of Equation (114) holds; that is, $\forall w < v : \min\{r(x)_w, r(y)_w\} > (x \oplus y)_v = \min\{x_v, y_v\}$. Assuming without loss of generality that $x_v \leq y_v$ (the other case is symmetric), we have that $x_v = (x \oplus y)_v = r(x \oplus y)_v$ and that $x_v \geq r(x)_v = (r(x) \oplus r(y))_v$, where $x_v = r(x)_v$ is implied by Equation (109) because $r(x)_w \geq \min\{r(x)_w, r(y)_w\} > \min\{x_v, y_v\} = x_v$ for any $w < v$. Thus, $\forall w < v : (r(x) \oplus r(y))_w > (r(x) \oplus r(y))_v$, yielding by applying Definition 7.3 once more that

$$r(r(x) \oplus r(y))_v = (r(x) \oplus r(y))_v = x_v = r(x \oplus y)_v. \quad (115)$$

Altogether, this shows Equation (110) and, as demonstrated above, implies Equation (19). $\square$

Having established that determining LE lists can be done by an MBF-like algorithm allows us to apply the machinery developed in Sections 2–5. Next, we establish that LE list computations can be performed efficiently, which we show by bounding the length of LE lists.

7.3 Computing LE Lists is Efficient

Our course of action is to show that LE list computations are efficient using Theorem 5.2 (i.e., the oracle theorem). The purpose of this section is to prepare the lemmas required to apply Theorem 5.2. We stress that the key challenge is to perform each iteration in polylogarithmic depth; this allows us to determine $\mathcal{A}(H)$ in polylogarithmic depth due to $\text{SPD}(H) \in O(\log^2 n)$. To this end, we first establish the length of intermediate LE lists to be logarithmic w.h.p. (Lemma 7.6). This permits us to apply r^V and determine the matrix-vector multiplication with A_λ (the scaled version of A_G , the adjacency matrix of G from Section 5) in a sufficiently efficient manner (Lemmas 7.7 and 7.8). Section 7.4 plugs these results into Theorem 5.2 to establish our main result.

We remark that LE lists are known to have length $O(\log n)$ w.h.p. throughout intermediate computations [26, 30], assuming that LE lists are assembled using h -hop distances. Lemma 7.6, while using the same key argument, is more general since it makes no assumption about x except for its independence of the random node order; we need the more general statement due to our decomposition of A_H .

Recall that by $|x|$ we denote the number of non- ∞ entries of $x \in \mathcal{D}$ and that we only need to keep the non- ∞ entries in memory. Lemma 7.6 shows that any LE list $r(x) \in \mathcal{D}$ has length $|r(x)| \in O(\log n)$ w.h.p., provided that x does not depend on the random node ordering. Observe that, in fact, the lemma is quite powerful as it suffices that there is *any* $y \in [x]$ that does not depend on the random node ordering: as $r(x) = r(y)$, then $|r(x)| = |r(y)| \in O(\log n)$ w.h.p.

LEMMA 7.6. *Let $x \in \mathcal{D}$ be arbitrary but independent of the random order of the nodes. Then $|r(x)| \in O(\log n)$ w.h.p.*

PROOF. Order the non- ∞ values of x by ascending distance, breaking ties independently of the random node order. Denote for $i \in \{1, \dots, |x|\}$ by $v_i \in V$ the i th node with respect to this order (i.e., x_{v_i} is the i th smallest entry in x). Furthermore, denote by X_i the indicator variable which is 1 if $v_i < v_j$ for all $j \in \{1, \dots, i-1\}$ and 0 otherwise. As the node order and x are independent, we obtain $\mathbb{E}[X_i] = 1/i$. For $X := \sum_{i=1}^{|x|} X_i$, this implies

$$\mathbb{E}[X] = \sum_{i=1}^{|x|} \frac{1}{i} \leq \sum_{i=1}^n \frac{1}{i} \in \Theta(\log n). \quad (116)$$

Observe that X_i is independent of $\{X_1, \dots, X_{i-1}\}$, as whether $v_i < v_j$ for all $j < i$ is independent of the internal order of the set $\{v_1, \dots, v_{i-1}\}$. We conclude that all $\{X_1, \dots, X_{|x|}\}$ are independent; this can be checked by inductively verifying that $\mathbb{P}[(X_1, \dots, X_k) = (b_1, \dots, b_k)] = \prod_{i=1}^k \mathbb{P}[X_i = b_i]$ for any possible assignment $(b_1, \dots, b_k) \in \{0, 1\}^k$. Applying Chernoff's bound yields $X \in \Theta(\log n)$ w.h.p. As $\mathbb{P}[X = k] = \mathbb{P}[|r(x)| = k]$, this concludes the proof. $\square$

Hence, filtered, possibly intermediate LE lists $r(x)$ w.h.p. comprise $O(\log n)$ entries. We proceed to show that, under these circumstances, $r(x)$ can be computed efficiently.

LEMMA 7.7. *Let $x \in \mathcal{D}$ be arbitrary. Then $r(x)$ can be computed using $O(|r(x)| \log n)$ depth and $O(|r(x)||x|)$ work.*

PROOF. We use one iteration per non- ∞ entry of $r(x)$. In each iteration, the smallest non-dominated entry of x_v is copied to $r(x)_v$, and all entries of x dominated by x_v are marked as dominated. This yields $|r(x)|$ iterations as follows:

- (1) Initialize $r(x) \leftarrow \perp$. Construct a balanced binary tree on the non- ∞ elements of x and identify its leaves with their indices $v \in V$ ($O(\log n)$ depth and $O(|x|)$ work).
- (2) Find the element with the smallest node index v with respect to the random node order whose corresponding leaf is not marked as *discarded* by propagating the minimum up the tree ($O(\log n)$ depth and $O(|x|)$ work). Set $r(x)_v \leftarrow x_v$.
- (3) Mark each leaf w for which $x_v \leq x_w$, including v , as *discarded* ($O(1)$ depth and $O(|x|)$ work).
- (4) If there are non-*discarded* leaves ($O(\log n)$ depth and $O(|x|)$ work), continue at step 2.

Note that for each $w \neq v$ for which the corresponding node is *discarded*, we have $r(x)_w = \infty$. On the other hand, by construction, we have for all v for which we stored $r(x)_v = x_v$ that there is no $w \in V$ satisfying both $x_w \leq x_v$ and $w < v$. Thus, the computed list is indeed $r(x)$.

The depth and work bounds follow from the above bounds on the complexities of the individual steps and by observing that, in each iteration, we add a distinct index-value pair (with non- ∞ value) to the list that after termination equals $r(x)$. $\square$

Any intermediate result used by the oracle is of the form $r^V A_\lambda y$ with

$$y = (r^V A_\lambda)^f P_\lambda x^{(h)}, \quad (117)$$

where $x^{(h)} = r^V A_H^h x^{(0)}$ is the intermediate result of h iterations on H , $\lambda \in \{0, \dots, \Lambda\}$ is a level, and $(r^V A_\lambda)^f P_\lambda$ represents another f iterations in G with edge weights stretched according to level λ . The oracle uses this to simulate the $(h+1)$ -th iteration in H (compare Section 5 and Theorem 5.2 in particular).

LEMMA 7.8. Consider $x^{(0)} \in \mathcal{D}^V$ from Equation (108). For arbitrary $h, f, \lambda \in \mathbb{N}_0$, we can w.h.p.

- (1) determine $r^V x^{(0)}$ from $x^{(0)}$ using $O(n)$ work and $O(1)$ depth,
- (2) compute $r^V A_\lambda y$ from y as defined in Equation (117) using $W \in O(m \log^2 n)$ work and $D \in O(\log^2 n)$ depth, and
- (3) compute $r^V (\bigoplus_{\lambda=0}^\Lambda y_\lambda)$ from the individual $y_\lambda = (P_\lambda r^V A_\lambda^d P_\lambda) x^{(i)}$, $0 \leq \lambda \leq \Lambda$, with $W_\oplus \in O(n \log^4 n)$ work and $D_\oplus \in O(\log^2 n)$ depth.

PROOF. We establish the claims in order.

- (1) Regarding the first claim, observe that $r^V x^{(0)} = x^{(0)}$. Hence, we can copy $x^{(0)}$ using constant depth and $O(\sum_{v \in V} |x_v^{(0)}|) = O(n)$ work.
- (2) As for the second claim, we expand $x^{(h)}$ in Equation (117) and remove all intermediate filtering steps, obtaining

$$y \stackrel{(5.8)}{=} (r^V A_\lambda)^f P_\lambda \left(r^V \left(\bigoplus_{\lambda=0}^\Lambda P_\lambda (r^V A_\lambda)^d P_\lambda \right) \right)^h r^V x^{(0)} \quad (118)$$

$$\stackrel{(2.35)}{=} \underbrace{r^V A_\lambda^f P_\lambda \left(\bigoplus_{\lambda=0}^\Lambda P_\lambda A_\lambda^d P_\lambda \right)^h}_{=: y'} x^{(0)}. \quad (119)$$

The key observation is that, since the random order of V only plays a role for r and we removed all intermediate applications of r^V , y' does not depend on that order. Hence, we

may apply Lemma 7.6, which yields for each $v \in V$ that $|y_v| = |r(y'_v)| \in O(\log n)$ w.h.p. Condition on $|y_v| \in O(\log n)$ for all $v \in V$ in the following, which happens w.h.p. by Lemma 1.2.

Regarding the computation of $r^V A_\lambda y$, we first compute each $(A_\lambda y)_v$ in parallel for all $v \in V$. By Lemma 2.3 and because $|y_v| \in O(\log n)$, this can be done using $O(\log n)$ depth and work

$$O\left(\sum_{v \in V} \sum_{\substack{w \in V \\ \{v, w\} \in E}} |y_w| \log n\right) \subseteq O\left(\sum_{\{v, w\} \in E} \log^2 n\right) = O(m \log^2 n). \quad (120)$$

Here, we use that propagation with respect to $\mathcal{D}$ —uniformly increasing weights—requires, due to $|y_v| \in O(\log n)$, no more than $O(1)$ depth and $O(m \log n)$ work and is thus dominated by aggregation. To bound the cost of computing $r^V A_\lambda y$ from $A_\lambda y$, observe that we have

$$|(A_\lambda y)_v| \in O\left(\sum_{\substack{w \in V \\ \{v, w\} \in E}} |y_w|\right). \quad (121)$$

Hence, by Lemma 7.7 and due to $|y_v| \in O(\log n)$, we can compute $(r^V A_\lambda y)_v$ in parallel for all $v \in V$ using $O(\log^2 n)$ depth and

$$O\left(\sum_{v \in V} |(A_\lambda y)_v| \log n\right) \stackrel{(7.18)}{\subseteq} O\left(\sum_{v \in V} \sum_{\substack{w \in V \\ \{v, w\} \in E}} |y_w| \log n\right) \quad (122)$$

$$\subseteq O\left(\sum_{\{v, w\} \in E} \log^2 n\right) \quad (123)$$

$$\subseteq O(m \log^2 n) \quad (124)$$

work. All operations are possible using $D \in O(\log^2 n)$ depth and $W \in O(m \log^2 n)$ work. As we condition only on an event that occurs w.h.p., this concludes the proof of the second claim.

- (3) Regarding the last claim, condition on logarithmic length of all LE lists (i.e., on $|(y_\lambda)_v| \in O(\log n)$ for all $0 \leq \lambda \leq \Lambda$). We can compute $\bigoplus_{\lambda=0}^{\Lambda} (y_\lambda)_v \in \mathcal{D}$, the aggregation for a single vertex v , using $O(\sum_{\lambda=0}^{\Lambda} |(y_\lambda)_v| \log n) = O(\log^3 n)$ work and $O(\log n)$ depth by Lemma 2.3. As the work bounds the length of the resulting list, we can determine $r(\bigoplus_{\lambda=0}^{\Lambda} (y_\lambda)_v)$ using $O(\log^4 n)$ work and $O(\log^2 n)$ depth by Lemma 7.7. Doing this in parallel for all $v \in V$ yields $W_\oplus \in O(n \log^4 n)$ work and $D_\oplus \in O(\log^2 n)$ depth. As we condition on two events that occur w.h.p., this concludes the last claim. $\square$

7.4 Metric Tree Embedding in Polylogarithmic Time and Near-Linear Work

Determining LE lists on H yields a probabilistic tree embedding of G with expected stretch $O(\log n)$ (Section 7.1), is the result of an MBF-like algorithm (Section 7.2), and each iteration of this algorithm is efficient (Theorem 5.2 and Section 7.3). We assemble these pieces in Theorem 7.9, which relies on G containing a suitable hop set. Corollaries 7.10 and 7.11 remove this assumption by

invoking known algorithms to establish this property first. Note that Theorem 7.9 serves as a blueprint yielding improved tree embedding algorithms when provided with improved hop set constructions.

THEOREM 7.9. *Suppose we are given the weighted incidence list of a graph $G = (V, E, \omega)$ satisfying for some $\alpha \in \mathbb{R}_{\geq 1}$ and $d \in \mathbb{N}$ that $\text{dist}^d(v, w, G) \leq \alpha \text{dist}(v, w, G)$ for all $v, w \in V$. Then, w.h.p., we can sample from a tree embedding of G of expected stretch $O(\alpha^{O(\log n)} \log n)$ with depth $O(d \log^4 n) \subset \tilde{O}(d)$ and work $O(m(d + \log n) \log^5 n) \subset \tilde{O}(md)$.*

PROOF. We first w.h.p. compute the LE lists of H . To this end, by Lemma 7.8, we may apply Theorem 5.2 with parameters $D \in O(\log^2 n)$, $W \in O(m \log^2 n)$, $D_{\oplus} \in O(\log^2 n)$, and $W_{\oplus} \in O(n \log^4 n)$; we arrive at depth $O(d \log^4 n)$ and work $O(m(d + \log n) \log^5 n)$. As shown in Fakcharoenphol et al. [23], the FRT tree T represented by these lists has expected stretch $O(\log n)$ with respect to the distance metric of H . By Theorem 4.5, w.h.p. $\text{dist}(v, w, G) \leq \text{dist}(v, w, H) \leq \alpha^{O(\log n)} \text{dist}(v, w, G)$ and hence

$$\text{dist}(v, w, G) \leq \text{dist}(v, w, T) \in O\left(\alpha^{O(\log n)} \log n \text{dist}(v, w, G)\right) \quad (125)$$

in expectation (compare Definition 7.1). Observe that, by Lemma 7.2, explicitly constructing the FRT tree is possible within the stated bounds. $\square$

As stated earlier, we require G to contain a $(d, 1/\text{polylog } n)$ -hop set with $d \in \text{polylog } n$ in order to achieve polylogarithmic depth. We also need to determine such a hop set using polylog n depth and near-linear work in m and that it does not significantly increase the problem size by adding too many edges. Cohen's hop sets [17] meet all these requirements, yielding the following corollary.

COROLLARY 7.10. *Given the weighted incidence list of a graph G and an arbitrary constant $\varepsilon > 0$, we can w.h.p. sample from a tree embedding of expected stretch $O(\log n)$ using depth $\text{polylog } n$ and work $\tilde{O}(m^{1+\varepsilon})$.*

PROOF. We apply the hop set construction by Cohen [17] to $G = (V, E, \omega)$ to w.h.p. determine an intermediate graph G' with vertices V and an additional $\tilde{O}(m^{1+\varepsilon})$ edges. The algorithm guarantees $\text{dist}^d(v, w, G) \leq \alpha \text{dist}(v, w, G')$ for $d \in \text{polylog } n$ and $\alpha \in 1 + 1/\text{polylog } n$ (where the polylog n term in α is under our control) and has depth $\text{polylog } n$ and work $\tilde{O}(m^{1+\varepsilon})$. Choosing $\alpha \in 1 + O(1/\log n)$ and applying Theorem 7.9, the claim follows. $\square$

Adding a hop set to G , embedding the resulting graph in H , and sampling an FRT tree on H is a three-step sequence of embeddings of G . Still, in terms of stretch, the embedding of Corollary 7.10 is—up to a factor in $1 + o(1)$ —as good as directly constructing an FRT tree of G : (1) Hop sets do not stretch distances. (2) By Theorem 4.5 and Equation (91), H introduces a stretch of $1 + 1/\text{polylog } n$. (3) Together, this ensures that the expected stretch of the FRT embedding with respect to G is $O(\log n)$.

It is possible to reduce the work at the expense of an increased stretch by first applying the spanner construction by Baswana and Sen [9]:

COROLLARY 7.11. *Suppose we are given the weighted incidence list of a graph G . Then, for any constant $\varepsilon > 0$ and any $k \in \mathbb{N}$, we can w.h.p. sample from a tree embedding of G of expected stretch $O(k \log n)$ using depth $\text{polylog } n$ and work $\tilde{O}(m + n^{1+1/k+\varepsilon})$.*

PROOF. The algorithm of Baswana and Sen [9] computes a $(2k - 1)$ -spanner of $G = (V, E, \omega)$; that is, a subgraph $G' = (V, E', \omega)$ satisfying for all $v, w \in V$ that $\text{dist}(v, w, G) \leq \text{dist}(v, w, G') \leq (2k - 1) \text{dist}(v, w, G)$ using $\text{polylog } n$ depth and $\tilde{O}(m)$ work. We argue in the proof of Theorem 6.3 that $|E'| \in \tilde{O}(n^{1+1/k})$ w.h.p. The claim follows from applying Corollary 7.10 to G' . $\square$

7.5 Reconstructing Paths from Virtual Edges

Given that we only deal with distances and not with paths in the FRT construction, there is one concern: Consider an arbitrary graph $G = (V, E, \omega)$, its augmentation with a hop set resulting in G' , which is then embedded into the complete graph H and finally into an FRT tree $T = (V_T, E_T, \omega_T)$. How can an edge $e \in E_T$ of weight $\omega_T(e)$ be mapped to a path p in G with $\omega(p) \leq \omega_T(H)$? Note that this question has to be answered in polylogarithmic depth and without incurring too much memory overhead. Our purpose is not to provide specifically tailored data structures, but we propose a three-step approach that maps edges in T to paths in H , edges in H to paths in G , and finally edges from G' to paths in G .

Concerning a tree edge $e \in E_T$, observe that e maps back to a path p of at most $\text{SPD}(H)$ hops in H with $\omega_H(p) \leq 3\omega_T(e)$ as follows. First, to keep the notation simple, identify each tree node—given as tuple $(v_i, \dots, v_j)$ —with its “leading” node $v_i \in V$; in particular, each leaf has $i = 0$ and is identified with the node in V that is mapped to it. A leaf v_0 has an LE entry $(\text{dist}(v_0, v_1, H), v_1)$, and we can trace the shortest v_0 - v_1 -path in H based on the LE lists (nodes locally store the predecessor of shortest paths just like in APSP). Moreover, $\text{dist}(v_i, v_{i+1}, H) \leq \omega_T(v_i, v_{i+1})$ (i.e., we may map the tree edge back to the path without incurring larger cost than in T). If $i > 0$, v_i and v_{i+1} are inner nodes. Choose an arbitrary leaf v_0 that is a common descendant (this choice can, e.g., be fixed when constructing the tree from the LE list without increasing the asymptotic bounds on depth or work). We then can trace shortest paths from v_0 to v_i and from v_0 to v_{i+1} in H , respectively. The cost of their concatenation is $\text{dist}(v_0, v_i, H) + \text{dist}(v_0, v_{i+1}, H) \leq \beta 2^i + \beta 2^{i+1} = 3(\beta 2^i) = 3\omega_T(v, w)$ by the properties of LE lists and the FRT embedding. Note that, due to the identification of each tree node with its “leading” graph node, paths in T map to concatenable paths in H .

Regarding the mapping from edges in H to paths in G , recall that we compute the LE lists of H by repeated application of the operations r^V , $\oplus$, P_λ , and A_λ with $0 \leq \lambda \leq \Lambda$. Observe that r^V , $\oplus$, and P_λ discard information; that is, distances to nodes that do not make it into the final LE lists and therefore are irrelevant to routing. A_λ , on the other hand, is an MBF step. Thus, we may store the necessary information for backtracing the induced paths at each node; specifically, we can store, for each iteration $h \in O(\log^2 n)$ with respect to H , each of the intermediate d iterations in G , and each $\lambda \in O(\log n)$, the state vector y of the form in Equation (117) in a lookup table. This requires $\tilde{O}(d)$ memory and efficiently maps edges of H to d -hop paths in G —or rather to d -hop paths in G' if we construct H after augmenting G to G' using a hop set.

Mapping edges of G' to edges in G depends on the hop set. Cohen [17] does not discuss this in her article, but her hop set edges can be efficiently mapped to paths in the original graph by a lookup table: Hop set edges either correspond to a shortest path in a small cluster or to a cluster that has been explored using polylogarithmic depth. Regarding other hop set algorithms, we note that many techniques constructing hop set edges using depth D allow for reconstruction of corresponding paths at depth $O(D)$ (i.e., that polylogarithmic-depth algorithms are compatible analogously to Cohen’s hop sets). For instance, this is the case for the hop set construction by Henzinger et al. [29], which we leverage in Section 8.3.

8 DISTRIBUTED FRT CONSTRUCTION

Distributed algorithms for constructing FRT-type tree embeddings in the Congest model are covered by our framework as well. In the following, we recap two existing algorithms [26, 30] (our framework allows doing this in a very compact way), and we improve upon the state of the art, reducing a factor of n^ϵ in the currently best-known round complexity for expected stretch $O(\log n)$ [26] to $n^{o(1)}$. We use the hop set of Henzinger et al. [29] instead of Cohen’s [17], because it is compatible with the Congest model. Note that replacing the hop set is straightforward since our theorems in the previous sections are formulated with respect to generic $(d, \hat{\epsilon})$ -hop sets.

The Congest Model. We refer to Peleg [41] for a formal definition of the Congest model but briefly outline its core aspects. The Congest model is a model of computation that captures distributed computations performed by the nodes of a graph, where communication is restricted to its edges. Each node is initialized with a unique ID of $O(\log n)$ bits, knows the IDs of its adjacent nodes along with the weights of the corresponding incident edges, and “its” part of the input (in our case the input is empty); each node has to compute “its” part of the output (in our case, as detailed in Section 7.1, its LE list). Computations happen in rounds, and we are interested in how many rounds it takes for an algorithm to complete. In each round, each node does the following:

- (1) Perform finite, but otherwise arbitrary local computations.
- (2) Send a message of $O(\log n)$ bits to each neighboring node.
- (3) Receive the messages sent by neighbors.

Recall that, by assumption, edge weights can be encoded using $O(\log n)$ bits; that is, an index-distance pair can be encoded in a single message.

Overview. Throughout this section, let $G = (V, E, \omega)$ be a weighted graph and denote, for any graph G , by $A_G \in (\mathbb{R}_{\geq 0} \cup \{\infty\})^{V \times V}$ its adjacency matrix according to Equation (4). Fix the semiring $\mathcal{S} = \mathcal{S}_{\min,+}$ and the zero-preserving semimodule $\mathcal{M} = \mathcal{D}$ from Definition 2.1, as well as $r, \sim$, and $x^{(0)}$ as given in Definition 7.3.

Sections 8.1 and 8.2 briefly summarize the distributed FRT algorithms by Kahn et al. [30] and Ghaffari and Lenzen [26], respectively. We use these preliminaries, our machinery, and a distributed hop set construction due to Henzinger et al. [29] in Section 8.3 to propose an algorithm that reduces a multiplicative overhead of n^ϵ in the round complexity of [26] to $n^{o(1)}$.

8.1 The Algorithm by Khan et al.

In our terminology, the algorithm of Khan et al. [30] performs $\text{SPD}(G)$ iterations of the MBF-like algorithm for collecting LE lists implied by Definition 7.3; that is,

$$r^V A_G^{\text{SPD}(G)} x^{(0)} \stackrel{(2.35)}{=} (r^V A_G)^{\text{SPD}(G)} x^{(0)}. \quad (126)$$

It does so in $\text{SPD}(G) + 1$ iterations by initializing $x^{(0)}$ as in Equation (108) and iteratively computing $x^{(i+1)} := r^V A_G x^{(i)}$ until a fixed point is reached (i.e., until $x^{(i+1)} = x^{(i)}$). As $(r^V A_G)^i x^{(0)} = r^V A_G^i x^{(0)}$, Lemma 7.6 shows, that w.h.p. $|x_v^{(i)}| \in O(\log n)$ for all $0 \leq i \leq \text{SPD}(G)$ and all $v \in V$. Therefore, $v \in V$ can w.h.p. transmit $x_v^{(i)}$ to all of its neighbors using $O(\log n)$ messages, and, upon reception of its neighbors' lists, locally compute $x_v^{(i+1)}$. Thus, each iteration takes $O(\log n)$ rounds w.h.p., implying the round complexity of $O(\text{SPD}(G) \log n)$ w.h.p. shown in Khan et al. [30].

8.2 The Algorithm by Ghaffari and Lenzen

The strongest lower bound regarding the round complexity for constructing a (low-stretch) metric tree embedding of G in the Congest model is $\tilde{\Omega}(\sqrt{n} + D(G))$ [20, 26]. If $\text{SPD}(G) \gg \max\{D(G), \sqrt{n}\}$, one may thus hope for a solution that runs in $\tilde{o}(\text{SPD}(G))$ rounds. For any $\epsilon \in \mathbb{R}_{>0}$, in Ghaffari and Lenzen [26], it is shown that expected stretch $O(\epsilon^{-1} \log n)$ can be achieved in $\tilde{O}(n^{1/2+\epsilon} + D(G))$ rounds; below, we summarize this algorithm.

The strategy is to first determine the LE lists of a constant-stretch metric embedding of (the induced submetric of) an appropriately sampled subset of V . The resulting graph is called the *skeleton spanner*, and its LE lists are then used to jump-start the computation on the remaining graph. When sampling the skeleton nodes in the right way, stretching non-skeleton edges analogously to Section 4, and fixing a shortest path for each pair of vertices, w.h.p. all of these paths contain

a skeleton node within a few hops. Ordering skeleton nodes before non-skeleton nodes with respect to the random ordering implies that each LE list has a short prefix accounting for the local neighborhood, followed by a short suffix containing skeleton nodes only. This is due to the fact that skeleton nodes dominate all non-skeleton nodes for which the respective shortest path passes through them. Hence, no node has to learn information that is further away than $d_S \in \tilde{O}(\sqrt{n})$, an upper bound on the number of hops when a skeleton node is encountered on a shortest path that holds w.h.p.

The Graph H . In Ghaffari and Lenzen [26], G is embedded into H and an FRT tree is sampled on H , where H is derived as follows. Abbreviate $\ell := \lceil \sqrt{n} \rceil$. For a sufficiently large constant c , sample $\lceil c\ell \log n \rceil$ nodes uniformly at random; call this set S . Define the *skeleton graph*

$$G_S := (S, E_S, \omega_S), \text{ where} \quad (127)$$

$$E_S := \left\{ \{s, t\} \in \binom{S}{2} \mid \text{dist}^\ell(s, t, G) < \infty \right\} \text{ and} \quad (128)$$

$$\omega_S(s, t) \mapsto \text{dist}^\ell(s, t, G). \quad (129)$$

Then w.h.p. $\text{dist}(s, t, G_S) = \text{dist}(s, t, G)$ for all $s, t \in S$ (Lemma 4.6 of [33]). For $k \in \Theta(\varepsilon^{-1})$, construct a $(2k - 1)$ -spanner

$$G'_S := (S, E'_S, \omega_S) \quad (130)$$

of the skeleton graph G_S that has $\tilde{O}(\ell^{1+1/k}) \subseteq \tilde{O}(n^{1/2+\varepsilon})$ edges w.h.p. (Lemma 4.9 of [33]). Define

$$H := (V, E_H, \omega_H), \text{ where} \quad (131)$$

$$E_H := E'_S \cup E, \text{ and} \quad (132)$$

$$\omega_H(e) \mapsto \begin{cases} \omega_S(e) & \text{if } e \in E'_S \text{ and} \\ (2k - 1) \omega(e) & \text{otherwise.} \end{cases} \quad (133)$$

By construction, G embeds into H with a stretch of $2k - 1$ w.h.p.; that is, $\text{dist}(v, w, G) \leq \text{dist}(v, w, H) \leq (2k - 1) \text{dist}(v, w, G)$. Computing an FRT tree T of H of expected stretch $O(\log n)$ thus implies that G embeds into T with expected stretch $O(k \log n) = O(\varepsilon^{-1} \log n)$.

FRT Trees of H . Observe that min-hop shortest paths in H contain only a single maximal subpath consisting of spanner edges, where the maximal subpaths of non-spanner edges have at most ℓ hops w.h.p. This follows analogously to Lemma 4.4 with two levels and a sampling probability of $\tilde{\Theta}(1/\ell)$. Assuming $s < v$ for all $s \in S$ and $v \in V \setminus S$ (we discuss this below) for each $v \in V$ and each entry $(w, \text{dist}(v, w, H))$ of its LE list, w.h.p. there is a min-hop shortest v - w -path with a prefix of ℓ non-spanner edges followed by a shortest path in G'_S . This entails that w.h.p.

$$r^V A_H^{\text{SPD}(H)} x^{(0)} = r^V A_{G, 2k-1}^\ell A_{G'_S}^{|S|} x^{(0)} = r^V A_{G, 2k-1}^\ell \underbrace{\left(r^V A_{G'_S}^{|S|} x^{(0)} \right)}_{=: \tilde{x}^{(0)}}, \quad (134)$$

where $A_{G, s}$ is A_G with entries stretched by a factor of $s \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, and we extend $A_{G'_S}$ to be a $V \times V$ matrix by setting $(A_{G'_S})_{vw} = \infty$ if $v \neq w \in V \setminus S$ and $(A_{G'_S})_{vv} = 0$ for $v \in V \setminus S$.

In order to construct an FRT tree, suppose we have sampled uniform permutations of S and $V \setminus S$ and a random choice of β . We extend the permutations to a permutation of V by ruling that, for all $s \in S$ and $v \in V \setminus S$, we have $s < v$, fulfilling the above assumption. Lemma 4.9 of Ghaffari and Lenzen [26] shows that the introduced dependence between the topology of H and the resulting permutation on V does not increase the expected stretch of the embedding beyond $O(\log n)$. The crucial advantage of this approach lies in the fact that now the LE lists of nodes in S may be used to jump-start the construction of LE lists for H , in accordance with Equation (134).

The Algorithm. In Ghaffari and Lenzen [26], it is shown that LE lists of H can be determined quickly in the Congest model as follows.

- (1) Some node v_0 starts by broadcasting k and a random choice of β , constructing a BFS tree on the fly. Upon receipt, each node generates a random ID of $O(\log n)$ bits which is unique w.h.p. Querying the amount of nodes with an ID of less than some threshold via the BFS tree, v_0 determines the bottom ℓ node IDs via binary search; these nodes form the set S and satisfy the assumption that went into Equation (134). All of these operations can be performed in $\tilde{O}(D(G))$ rounds.
- (2) The nodes in S determine G'_S , which is possible in $\tilde{O}(D(G) + \ell^{1+1/k}) \subseteq \tilde{O}(D(G) + n^{1/2+\varepsilon})$ rounds, such that all $v \in V$ learn E'_S and ω_S [26, 33]. After that, G'_S is global knowledge and each $v \in V$ can locally compute $\bar{x}_v^{(0)}$.
- (3) Subsequently, nodes w.h.p. determine their component of $r^V A_{G,2k-1}^\ell \bar{x}^{(0)} = (r^V A_{G,2k-1})^\ell \bar{x}^{(0)}$ via ℓ MBF-like iterations of

$$\bar{x}^{(i+1)} := r^V A_{G,2k-1} \bar{x}^{(i)}. \quad (135)$$

Here, one exploits that, for all i , $|\bar{x}_v^{(i)}| \in O(\log n)$ w.h.p. by Lemma 7.6,¹⁰ and thus each iteration can be performed by sending $O(\log n)$ messages over each edge (i.e., in $O(\log n)$ rounds); the entire step thus requires $\tilde{O}(\ell) \subseteq \tilde{O}(n^{1/2})$ rounds.

Together, this w.h.p. implies the round complexity of $\tilde{O}(n^{1/2+\varepsilon} + D(G))$ for an embedding of expected stretch $O(\varepsilon^{-1} \log n)$.

8.3 Achieving Stretch $O(\log n)$ in Near-Optimal Time

The multiplicative overhead of n^ε in the round complexity is due to constructing and broadcasting the skeleton spanner G'_S . We can improve upon this by relying on hop sets, just as we do in our parallel construction. Henzinger et al. [29] show how to compute an $(n^{o(1)}, o(1))$ -hop set of the skeleton graph in the Congest model using $n^{1/2+o(1)} + D(G)^{1+o(1)}$ rounds.

Our approach is similar to the one outlined in Section 8.2. The key difference is that we replace the use of a spanner by combining a hop set of the skeleton graph with the construction from Section 4; using the results from Section 5, we can then efficiently construct the LE lists on S to jump-start the construction of LE lists for all nodes.

The Graph H . Let ℓ , c , and the skeleton graph $G_S = (S, E_S, \omega_S)$ be defined as in Section 8.2 and Equations (127)–(128), w.h.p. yielding $\text{dist}(s, t, G_S) = \text{dist}(s, t, G)$ for all $s, t \in S$. Suppose for all $s, t \in S$, we know approximate weights $\omega'_S(s, t)$ with

$$\text{dist}(s, t, G) \leq \omega'_S(s, t) \in (1 + o(1)) \omega_S(s, t)$$

(our algorithm has to rely on an approximation to meet the stated round complexity) and add an $(n^{o(1)}, o(1/\log n))$ -hop set to G_S using the construction of Henzinger et al. [29]. Together, this results in a graph

$$G'_S := (S, E'_S, \omega'_S), \quad (136)$$

¹⁰We apply Lemma 7.6 twice, as it requires $x \in \mathcal{D}$ to be independent of the permutation. First consider a computation initialized with $y_{vw}^{(0)} := 0$ if $v = w \in S$ and $y_{vw}^{(0)} := \infty$ else. By Lemma 7.6, we have $|y_v^{(i)}| \in O(\log n)$ w.h.p. for all $y^{(i)} := r^V A_{H_S}^i y^{(0)}$ and iterations $i \in \{1, \dots, |S|\}$. Analogously, apply Lemma 7.6 to $z^{(i)} := r^V A_{G,2k-1}^i z^{(0)}$, $i \in \{1, \dots, \ell\}$ with $z_{vw}^{(0)} := 0$ if $v = w \in V \setminus S$ and $z_{vw}^{(0)} := \infty$ else; this yields that $|z_v^{(i)}| \in O(\log n)$ for all $v \in V$ w.h.p., too. As we have $x_v^{(i)} = r^V (y_v^{(j)} \oplus z_v^{(k)})$ for all $v \in V$ and appropriate $i, j, k \in \mathbb{N}$, we obtain $|x_v^{(i)}| \in O(\log n)$ w.h.p.

where E'_S contains the skeleton edges E_S and some additional edges, and w.h.p. it holds for all $s, t \in S$ that

$$\text{dist}(s, t, G_S) \leq \text{dist}^d(s, t, G'_S) \in (1 + o(1/\log n)) \text{dist}(v, w, G_S) \quad (137)$$

for some $d \in n^{o(1)}$ and $\text{dist}(v, w, G) \leq \text{dist}(v, w, G_S) \in (1 + o(1)) \text{dist}(v, w, G)$. Next, embed G'_S into H_S as in Section 4, yielding node and edge levels $\lambda(e) \in \{0, \dots, \Lambda\}$:

$$H_S := \left(S, \binom{S}{2}, \omega_{H_S} \right) \text{ with} \quad (138)$$

$$\omega_{H_S}(\{s, t\}) \mapsto (1 + \hat{\varepsilon})^{\Lambda - \lambda(s, t)} \text{dist}^d(s, t, G'_S) \quad (139)$$

with d as above, $\hat{\varepsilon} \in o(1/\log n)$. By Theorem 4.5, w.h.p. we have that $\text{SPD}(G) \in O(\log^2 n)$ and for all $s, t \in S$ that

$$\text{dist}(s, t, G) \leq \text{dist}(s, t, G_S) \leq \text{dist}(s, t, H_S) \in (1 + o(1)) \text{dist}(s, t, G_S), \quad (140)$$

which is bounded from above by $\alpha \text{dist}(s, t, G)$ for some $\alpha \in 1 + o(1)$. Analogously to Equations (8.6)–(8.8), define

$$H := (V, E_H, \omega_H), \text{ where} \quad (141)$$

$$E_H := E \cup \binom{S}{2}, \text{ and} \quad (142)$$

$$\omega_H(e) \mapsto \begin{cases} \omega_{H_S}(e) & \text{if } e \in \binom{S}{2} \text{ and} \\ \alpha \omega_G(e) & \text{otherwise.} \end{cases} \quad (143)$$

By construction, we thus have

$$\forall v, w \in V : \quad \text{dist}(v, w, G) \leq \text{dist}(v, w, H) \leq \alpha \text{dist}(v, w, G) \in (1 + o(1)) \text{dist}(v, w, G) \quad (144)$$

w.h.p.

FRT Trees of H . Analogously to Section 8.2, assume that the node IDs of S are ordered before those of $V \setminus S$; then min-hop shortest paths in H contain a single maximal subpath of edges in E_{H_S} . To determine the LE lists for H , we must therefore compute

$$r^V A_H^{\text{SPD}(H)} x^{(0)} = \left(r^V A_{G, \alpha} \right)^\ell \underbrace{\left(r^V A_{H_S} \right)^{\text{SPD}(H_S)}}_{=: \tilde{x}^{(0)}} x^{(0)}, \quad (145)$$

where $A_{G, \alpha}$ is given by multiplying each entry of A_G by the above-mentioned factor of α , and A_{H_S} is extended to an adjacency matrix on the node set V as in Section 8.2.

The Algorithm. We determine the LE lists of H as follows, adapting the approach from Ghaffari and Lenzen [26] outlined in Section 8.2.

- (1) A node v_0 starts the computation by broadcasting a random choice of β . The broadcast is used to construct a BFS tree, nodes generate distinct random IDs of $O(\log n)$ bits w.h.p., and v_0 figures out the ID threshold of the bottom $c\ell$ nodes S with respect to the induced random ordering. This can be done in $\tilde{O}(D(G))$ rounds.
- (2) Each skeleton nodes $s \in S$ computes $\omega'_S(s, t)$ as above for all $t \in S$, using the $(1 + 1/\log^2 n)$ -approximate $(S, \ell, |S|)$ -detection algorithm given in Lenzen and Patt-Shamir [35]. This takes $\tilde{O}(\ell + \ell) = \tilde{O}(n^{1/2})$ rounds.
- (3) Run the algorithm of Henzinger et al. [29] to compute an $(n^{o(1)}, o(1))$ -hop set of G'_S , in the sense that nodes in S learn their incident weighted edges. This takes $n^{1/2+o(1)} + D(G)^{1+o(1)}$ rounds.

- (4) Next, we (implicitly) construct H_S . To this end, nodes in S locally determine their level and broadcast it over the BFS tree, which takes $O(|S| + D(G)) \subset \tilde{O}(\sqrt{n} + D(G))$ rounds; thus, $s \in S$ knows the level of $\{s, t\} \in E_{H_S}$ for each $t \in S$.
- (5) To determine $\tilde{x}^{(0)}$, we follow the same strategy as in Theorem 5.2; that is, we simulate matrix-vector multiplication with A_{H_S} via matrix-vector multiplications with $A_{G'_S}$. Hence, it suffices to show that we can efficiently perform a matrix-vector multiplication $A_{G'_S}x$ for any x that may occur during the computation—applying r^V is a local operation and thus free—assuming each node $v \in V$ knows x_v and its row of the matrix.

Since multiplications with $A_{G'_S}$ only affects lists at skeleton nodes, this can be done by local computations once all nodes know x_s for each $s \in S$. As before, $|x_s| \in O(\log n)$ w.h.p., so $\sum_{s \in S} |x_s| \in O(|S| \log n) \subset \tilde{O}(\sqrt{n})$ w.h.p. We broadcast these lists over the BFS tree of G , taking $\tilde{O}(\sqrt{n} + D(G))$ rounds per matrix-vector multiplication. Due to $\text{SPD}(H_S) \in \tilde{O}(\log^2 n)$, by Theorem 4.5, this results in a round complexity of $\tilde{O}(n^{1/2+o(1)} + D(G)^{1+o(1)})$.

- (6) Applying $r^V A_{G,\alpha}^\ell$ is analogous to step 3 in Section 8.2 and takes $\tilde{O}(\ell) \subseteq \tilde{O}(n^{1/2})$ rounds.

Altogether, this yields a round complexity of $n^{1/2+o(1)} + D(G)^{1+o(1)}$. Combining this result with the algorithm by Khan et al. [30], which terminates quickly if $\text{SPD}(G)$ is small, yields the following result.

THEOREM 8.1. *There is a randomized distributed algorithm that w.h.p. samples from a metric tree embedding of expected stretch $O(\log n)$ in $\min\{(\sqrt{n} + D(G))n^{o(1)}, \tilde{O}(\text{SPD}(G))\}$ rounds of the Congest model.*

9 k -MEDIAN

In this section, we turn to the k -median problem, an application considered by Blelloch et al. [14], and show how their results are improved by applying our techniques. The contribution is that we work on a weighted graph G that only implicitly provides the distance metric $\text{dist}(\cdot, \cdot, G)$; Blelloch et al. require a metric providing constant-time query access. Our solution is more general, as any finite metric defines a complete graph of $\text{SPD } 1$, whereas determining exact distances in graphs (by known techniques) requires $\Omega(\text{SPD}(G))$ depth. The use of hop sets, however, restricts us to polynomially bounded edge-weight ratios.

Definition 9.1 (k -Median). In the k -median problem we are given a weighted graph $G = (V, E, \omega)$ and an integer $k \in \mathbb{N}$. The task is to determine $F \subseteq V$ with $|F| \leq k$ that minimizes

$$\sum_{v \in V} \text{dist}(v, F, G), \quad (146)$$

where $\text{dist}(v, F, G) := \min\{\text{dist}(v, f, G) \mid f \in F\}$ is the distance of v to the closest member of F .

Blelloch et al. [14] solve the following problem: Given a metric with constant-time query access, determine an expected $O(\log k)$ -approximation of k -median using $O(\log^2 n)$ depth and $\tilde{O}(nk + k^3)$ work for $k \geq \log n$; the special case of $k < \log n$ admits an $\tilde{O}(n)$ -work solution of the same depth [15]. Below, we show how to determine an expected $O(\log k)$ -approximation of k -median on a weighted graph using $\text{polylog } n$ depth and $\tilde{O}(m^{1+\epsilon} + k^3)$ work.

The algorithm of Blelloch et al. [14] essentially comprises three steps:

- (1) Use a parallel version of a sampling technique due to Mettu and Plaxton [38]. It samples candidates Q , such that $|Q| \in O(k)$ and there is $F \subseteq Q$ that $O(1)$ -approximates k -median.

- (2) Sample an FRT tree regarding the submetric spanned by Q . Normalize the tree to a binary tree (required by the next step); this is possible without incurring too much overhead with respect to the depth of the tree [14].
- (3) Run an $O(k^3)$ -work dynamic programming algorithm to solve the tree instance optimally without using any Steiner nodes. This yields an $O(\log k)$ -approximate solution on the original metric due to the expected stretch from the FRT embedding.

We keep the overall structure but modify steps 1–2, resulting in the following algorithm:

- (1) The sampling step generates $O(k)$ candidate points Q .
 It requires $O(\log \frac{n}{k})$ iterations and maintains a candidate set U that initially contains all points. In each iteration, $O(\log n)$ candidates S are sampled, and a constant fraction of vertices in U , those closest to S , is removed [14].
 The key to adapting this procedure to graphs lies in efficiently determining $\text{dist}(u, S, G)$ for all $u \in U$ (this would be trivial with constant-time query access to the metric). We achieve this by sampling after embedding in H from Section 4, which only costs a factor of $(1 + o(1))$ in approximation, regardless of k . By Theorem 4.5, we only require $O(\log^2 n)$ iterations of the MBF-like algorithm from Example 3.7 (for $d = \infty$) to determine each node's distance to the closest vertex in S w.h.p. Hence, we require polylogarithmic depth and $\tilde{O}(m^{1+\varepsilon})$ work for this step.
 Since $|U|$ decreases by a constant factor in each iteration and we have $O(\log n)$ iterations, we require a total of $\tilde{O}(m^{1+\varepsilon})$ work and polylogarithmic depth, including the costs for determining Cohen's hop set [17].
- (2) Sample an FRT tree on the submetric spanned by Q .
 To compute the embedding only on Q , set $x_{vv}^{(0)} = 0$ if $v \in Q$ and $x_{vw}^{(0)} = \infty$ everywhere else. Consider only the LE lists of nodes in Q when constructing the tree.
 As we are limited to polynomially bounded edge-weight ratios, our FRT trees have logarithmic depth. We normalize to a binary tree using the same technique as Blleloch et al. [14].
- (3) The $\tilde{O}(k^3)$ -work polylogarithmic-depth dynamic-programming algorithm of Blleloch et al. can be applied without modification.

W.h.p., we arrive at an expected $O(\log k)$ -approximation of k -median:

THEOREM 9.2. *For any fixed constant $\varepsilon > 0$, w.h.p., an expected $O(\log k)$ -approximation to k -median on a weighted graph can be computed using $\text{polylog } n$ depth and $\tilde{O}(m^{1+\varepsilon} + k^3)$ work.*

We remark that, analogously to Corollary 7.11, one can first compute a sparse spanner to reduce the work to $\tilde{O}(m + n^{1+\varepsilon} + k^3)$.

10 BUY-AT-BULK NETWORK DESIGN

In this section, we reduce the work of the approximation algorithm for the buy-at-bulk network design problem given by Blleloch et al. [14] that requires $O(n^3 \log n)$ work and $O(\log^2 n)$ depth w.h.p. while providing the same asymptotic approximation guarantees. Blleloch et al. transform the input graph G into a metric that allows constant-time query access on which they sample an FRT embedding, hence their work is dominated by solving APSP.

Replacing the APSP routine in the algorithm Blleloch et al. with our $O(1)$ -approximate metric from Theorem 6.3 (and keeping the rest of the algorithm in place) directly reduces the work to $\tilde{O}(n^{2+\varepsilon})$ while incurring $\text{polylog } n$ depth. However, using our result from Section 7 to sample an FRT without the detour over the metric, we can guarantee a stronger work bound of $\tilde{O}(\min\{m^{1+\varepsilon} +$

$kn, n^2\}) \subseteq \tilde{O}(n^2)$, which achieves the same depth. The use of hop sets, however, restricts us to polynomially bounded edge ratios (or our solution loses efficiency).

Definition 10.1 (Buy-at-Bulk Network Design). In the *buy-at-bulk network design problem*, one is given a weighted graph $G = (V, E, \omega)$, demands $(s_i, t_i, d_i) \in V \times V \times \mathbb{R}_{>0}$ for $1 \leq i \leq k$, and a finite set of cable types $(u_i, c_i) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$, $1 \leq i \leq \ell$, where the cable of type i incurs costs $c_i \omega(e)$ when purchased for edge e (multiple cables of the same type can be bought for an edge). The goal is to find an assignment of cable types and multiplicities to edges minimizing the total cost, such that the resulting edge capacities allow us to simultaneously route d_i units of (distinct) flow from s_i to t_i for all $1 \leq i \leq k$.

Andrews showed that the buy-at-bulk network design problem is hard to approximate better than with factor $\log^{1/2-o(1)} n$ [5]. Blleloch et al. [14] give an expected $O(\log n)$ -approximation w.h.p. using $\text{polylog } n$ depth and $O(n^3 \log n)$ work for the buy-at-bulk network design problem. It is a straightforward parallelization of the algorithm by Awerbuch and Azar [6]. Our tools allow for a more work-efficient parallelization of this algorithm, as the work of the implementation by Blleloch et al. is dominated by solving APSP to determine the distance metric of the graph; we achieve the same approximation guarantee as Blleloch et al. using $\text{polylog } n$ depth and $\tilde{O}(n^2)$ work. We propose the following modification of the approach of Blleloch et al.

- (1) Metrically embed G into a tree $T = (V_T, E_T, \omega_T)$ with expected stretch $O(\log n)$. As the objective is linear in the edge weights, an optimal solution in G induces a solution in T whose expected cost is by at most a factor $O(\log n)$ larger.
- (2) $O(1)$ -approximate on T : For $e \in E_T$, pick the cable of type i that minimizes $c_i \lceil d_e / u_i \rceil$, where d_e is the accumulated flow on e , see [14]).
- (3) Map the tree solution back to G , increasing the cost by a factor of $O(1)$.

Combining these steps yields an $O(\log n)$ -approximation. Using Corollary 7.11 (for constant ε and $k = \lceil \varepsilon^{-1} \rceil$), the first step has $\text{polylog } n$ depth and $\tilde{O}(m + n^{1+\varepsilon})$ work; for the second step, Blleloch et al. discuss an algorithm of $\text{polylog } n$ depth and $\tilde{O}(n + k)$ work.

Concerning the third step, recall that each tree edge $\{v, w\}$ maps back to a path p of at most $\text{SPD}(H)$ hops in H with $\omega(p) \leq 3 \omega_T(v, w)$, as argued in Section 7.5. Using this observation, we can map the solution on T back to one in H whose cost is at most by factor 3 larger. Assuming suitable data structures are used, this operation has depth $\text{polylog } n$ and requires $\tilde{O}(\min\{k, n\})$ work w.h.p., where we exploit that $\text{SPD}(H) \in O(\log^2 n)$ w.h.p. by Theorem 4.5 and the fact that T has depth $O(\log n)$, implying that the number of edges in T with non-zero flow is bounded by $O(\min\{k, n\} \log n)$.

Finally, we map back from H to G' (G augmented with hop set edges) and then to G . This can be handled with depth $\text{polylog } n$ and $\tilde{O}(n)$ work for a single edge in H because edges in H and hop set edges in G' correspond to polylogarithmically many edges in G' and at most n edges in G , respectively. The specifics depend on the hop set, and, again, we assume that suitable data structures are in place (see Section 7.5). Since we deal with $\tilde{O}(\min\{k, n\})$ edges in H , mapping back the edges yields $\tilde{O}(\min\{kn, n^2\})$ work in total. Together with the computation of the hop set, we have $\tilde{O}(\min\{m + n^{1+\varepsilon}, n^2\} + \min\{kn, n^2\}) = \tilde{O}(\min\{m + n(k + n^\varepsilon), n^2\}) \subseteq \tilde{O}(n^2)$ work.

THEOREM 10.2. *For any constant $\varepsilon > 0$, w.h.p., an expected $O(\log n)$ -approximation to the buy-at-bulk network design problem can be computed using $\text{polylog } n$ depth and $\tilde{O}(\min\{m + n(k + n^\varepsilon), n^2\}) \subseteq \tilde{O}(n^2)$ work.*

11 CONCLUSION

In this work, we show how to sample from an FRT-style distribution of metric tree embeddings at low depth and near-optimal work, provided that the maximum ratio between edge weights is polynomially bounded. While we consider the polylogarithmic factors too large for our algorithm to be of practical interest, this result motivates the search for solutions that achieve low depth despite having work comparable to the currently best-known sequential bound of $O(m \log n)$ w.h.p. [13]. Concretely, better hop set constructions could readily be plugged into our machinery to yield improved bounds, and one may seek to reduce the number of logarithmic factors incurred by the remaining construction.

Our second main contribution is an algebraic interpretation of MBF-like algorithms, reducing the task of devising and analyzing such algorithms to the following recipe:

- (1) Pick a suitable semiring $\mathcal{S}$ and semimodule $\mathcal{M}$ over $\mathcal{S}$.
- (2) Choose a filter r and initial values $x^{(0)} \in \mathcal{M}^V$ so that $r^V A^h x^{(0)}$ is the desired output.
- (3) Verify that r induces a congruence relation on $\mathcal{M}$.
- (4) Leverage (repeated use of) r^V to ensure that iterations can be implemented efficiently.

As can be seen by the example of our metric tree embedding algorithm, further steps may be required to control the number of iterations h ; concretely, we provide an embedding into a complete graph of small SPD and an oracle allowing for efficient MBF-like queries. Nevertheless, we believe that our framework unifies and simplifies the interpretation and analysis of MBF-like algorithms, as illustrated by the examples listed in Sections 3 and the discussion of distributed tree embeddings in Section 8. Therefore, we hope that our framework will be of use in the design of further efficient MBF-like algorithms in the future.

APPENDIX

A ALGEBRAIC FOUNDATIONS

For the sake of self-containment and unambiguity, we give the algebraic definitions required in this article as well as a standard result. Definitions A.1, A.2, and A.3 are slightly adapted from Chapters 1 and 5 of Hebisch and Weinert [28]. In this section, we refer to the neutral elements of addition and multiplication as 0 and 1. Note, however, that in the min-plus semiring $\mathcal{S}_{\min,+}$ the neutral element of “addition” (min) is ∞ and that of “multiplication” (+) is 0.

Definition A.1 (Semigroup). Let $M \neq \emptyset$ be a set and $\circ : M \times M \rightarrow M$ a binary operation. $(M, \circ)$ is a *semigroup* if and only if $\circ$ is associative; that is,

$$\forall x, y, z \in M : \quad x \circ (y \circ z) = (x \circ y) \circ z. \quad (\text{A.1})$$

A semigroup $(M, \circ)$ is *commutative* if and only if

$$\forall x, y \in M : \quad x \circ y = y \circ x. \quad (\text{A.2})$$

$e \in M$ is a *neutral element* of $(M, \circ)$ if and only if

$$\forall x \in M : \quad e \circ x = x \circ e = x. \quad (\text{A.3})$$

Some authors do not require semirings to have neutral elements or an annihilating 0. We, however, need them and work on semirings (mostly on $\mathcal{S}_{\min,+}$, $\mathcal{S}_{\max,\min}$, and $\mathcal{P}_{\min,+}$) which provide them, anyway.

Definition A.2 (Semiring). Let $M \neq \emptyset$ be a set, and $\oplus, \odot : M \times M \rightarrow M$ binary operations. Then $(M, \oplus, \odot)$ is a *semiring* if and only if

- (1) $(M, \oplus)$ is a commutative semigroup with neutral element 0,
- (2) $(M, \odot)$ is a semigroup with neutral element 1,
- (3) the left- and right-distributive laws hold:

$$\forall x, y, z \in M : \quad x \odot (y \oplus z) = (x \odot y) \oplus (x \odot z), \quad (\text{A.4})$$

$$\forall x, y, z \in M : \quad (y \oplus z) \odot x = (y \odot x) \oplus (z \odot x), \text{ and} \quad (\text{A.5})$$

- (4) 0 annihilates with respect to $\odot$:

$$\forall x \in M : \quad 0 \odot x = x \odot 0 = 0. \quad (\text{A.6})$$

Definition A.3 (Semimodule). Let $\mathcal{S} = (S, \oplus, \odot)$ be a semiring. $\mathcal{M} = (M, \oplus, \odot)$ with binary operations $\oplus : M \times M \rightarrow M$ and $\odot : S \times M \rightarrow M$ is a *semimodule over $\mathcal{S}$* if and only if

- (1) $(M, \oplus)$ is a semigroup and
- (2) for all $s, t \in S$ and all $x, y \in M$:

$$1 \odot x = x, \quad (\text{A.7})$$

$$s \odot (x \oplus y) = (s \odot x) \oplus (s \odot y), \quad (\text{A.8})$$

$$(s \oplus t) \odot x = (s \odot x) \oplus (t \odot x), \text{ and} \quad (\text{A.9})$$

$$(s \odot t) \odot x = s \odot (t \odot x). \quad (\text{A.10})$$

$\mathcal{M}$ is *zero-preserving* if and only if

- (1) $(M, \oplus)$ has the neutral element 0 and
- (2) $0 \in S$ is an annihilator for $\odot$:

$$\forall x \in M : \quad 0 \odot x = 0. \quad (\text{A.11})$$

A frequently used semimodule over the semiring $\mathcal{S}$ is $\mathcal{S}^k$ with coordinate-wise addition, i.e., k -dimensional vectors over $\mathcal{S}$. In particular, $\mathcal{S} = \mathcal{S}^1$ is always a semimodule over itself.

LEMMA A.4. *Let $\mathcal{S} = (S, \oplus, \odot)$ be a semiring and $k \in \mathbb{N}$ an integer. Then $\mathcal{S}^k := (\mathcal{S}^k, \oplus, \odot)$ with, for all $s \in \mathcal{S}$, $x, y \in \mathcal{S}^k$, and $1 \leq i \leq k$,*

$$(x \oplus y)_i := x_i \oplus y_i \text{ and} \quad (\text{A.12})$$

$$(s \odot x)_i := s \odot x_i \quad (\text{A.13})$$

is a zero-preserving semimodule over $\mathcal{S}$ with zero $(0, \dots, 0)$.

PROOF. We check the conditions of Definition A.3 one by one. Throughout the proof, let $s, t \in \mathcal{S}$ and $x, y \in \mathcal{S}^k$ be arbitrary.

- (1) $(\mathcal{S}^k, \oplus)$ is a semigroup because $(S, \oplus)$ is.
- (2) Equations (A.7)–(A.10) hold due to

$$(1 \odot x)_i = 1 \odot x_i = x_i, \quad (\text{A.14})$$

$$(s \odot (x \oplus y))_i = s \odot (x_i \oplus y_i) = (s \odot x_i) \oplus (s \odot y_i) = ((s \odot x) \oplus (s \odot y))_i, \quad (\text{A.15})$$

$$((s \oplus t) \odot x)_i = (s \oplus t) \odot x_i = (s \odot x_i) \oplus (t \odot x_i) = ((s \odot x) \oplus (t \odot x))_i, \text{ and} \quad (\text{A.16})$$

$$((s \odot t) \odot x)_i = (s \odot t) \odot x_i = s \odot (t \odot x_i) = (s \odot (t \odot x))_i. \quad (\text{A.17})$$

- (3) $(0, \dots, 0)$ is the neutral element of $(\mathcal{S}^k, \oplus)$ because 0 is the neutral element of $(S, \oplus)$.

(4) 0 is an annihilator for $\odot$:

$$(0 \odot x)_i = 0 \odot x_i = 0. \quad (\text{A.18})$$

□

B DEFERRED PROOFS

This appendix contains the proofs deferred from Section 3 for the sake of presentation.

Proof of Lemma 3.1.

PROOF. The claim trivially holds for $h = 0$. As induction hypothesis, suppose the claim holds for $h \in \mathbb{N}$. We obtain

$$x_{vw}^{(h+1)} = (Ax^{(h)})_{vw} \quad (\text{B.1})$$

$$= \left(\bigoplus_{u \in V} a_{vu} \odot x_u^{(h)} \right)_w \quad (\text{B.2})$$

$$= \bigoplus_{u \in V} a_{vu} \odot x_{uw}^{(h)} \quad (\text{B.3})$$

$$= \min_{u \in V} \{a_{vu} + x_{uw}^{(h)}\} \quad (\text{B.4})$$

$$= \min \left\{ \omega(v, u) + \text{dist}^h(u, w, G) \mid \{v, u\} \in E \right\} \cup \{0 + \text{dist}^h(v, w, G)\}; \quad (\text{B.5})$$

that is, exactly the definition of $\text{dist}^{h+1}(v, w, G)$, as claimed. □

Proof for Example 3.2.

PROOF. Let $s \in \mathcal{S}_{\min,+}$ be arbitrary and let $x, x', y, y' \in \mathcal{D}$ be such that $x \sim x'$ and $y \sim y'$, where $x \sim y :\Leftrightarrow r(x) = r(y)$. By Lemma 2.8, it suffices to show (1) that $r^2 = r$, (2) that $r(sx) = r(sx')$, and (3) that $r(x \oplus y) = r(x' \oplus y')$.

We show the claims one by one. First, observe that $r(x)_v = \infty$ for all $v \in V \setminus S$; hence, without loss of generality, assume $v \in S$ in the following. (1) $r(x)$ has at most k entries, each at most d , so $r(r(x)) = r(x)$ by (3.4). (2) Since multiplication with s uniformly increases the non- ∞ entries of x and x' , it does not affect their ordering with respect to Equation (48). As the k smallest S -entries of x and x' with respect to Equation (48) are identical, so are those of sx and sx' . Some entry $(sx)_v$ may become larger than d , but that happens for $(sx')_v$ as well; hence, $r(sx) = r(sx')$. (3) We have $r(x \oplus y)_v \leq d$ only if $(x \oplus y)_v = \min\{x_v, y_v\} \leq d$ is among the k smallest entries of $(x \oplus y)$ with respect to (48). If that is the case, there are no k entries smaller than $r(x \oplus y)_v$ in x or in y . Hence, these entries exist in x' and y' as well, form the k smallest entries of $(x' \oplus y')$, and $r(x \oplus y)_v = r(x' \oplus y')_v$ follows.

Proof of Lemma 3.10.

PROOF. We check each of the requirements of Definition A.2 in Appendix A. Throughout the proof, let $x, y, z \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ be arbitrary.

- (1) $(\mathbb{R}_{\geq 0} \cup \{\infty\}, \max)$ is a commutative semigroup because $\max$ is associative and commutative. Since 0 is the minimum of $\mathbb{R}_{\geq 0} \cup \{\infty\}$, it is the neutral element of $(\mathbb{R}_{\geq 0} \cup \{\infty\}, \max)$.
- (2) $(\mathbb{R}_{\geq 0} \cup \{\infty\}, \min)$ is a semigroup because $\min$ is associative. As above, ∞ is its neutral element because it is the maximum of $\mathbb{R}_{\geq 0} \cup \{\infty\}$.
- (3) Regarding the left- and right-distributive laws in Equations (A.4)–(A.5), a case distinction between the cases (a) $x \leq y \leq z$, (b) $y \leq x \leq z$, and (c) $y \leq z \leq x$ is exhaustive due to the

commutativity of min and max and reveals that

$$\min\{x, \max\{y, z\}\} = \max\{\min\{x, y\}, \min\{x, z\}\}; \quad (\text{B.6})$$

that is, that the left-distributive law holds. Since min is commutative,

$$\min\{\max\{y, z\}, x\} = \max\{\min\{y, x\}, \min\{z, x\}\} \quad (\text{B.7})$$

immediately follows; hence, $\mathcal{S}_{\max, \min}$ fulfills both distributive laws.

(4) 0 is an annihilator for min because

$$\min\{0, x\} = \min\{x, 0\} = 0. \quad (\text{B.8})$$

Together, it follows that $\mathcal{S}_{\max, \min}$ is a semiring as claimed. $\square$

Proof of Lemma 3.12.

PROOF. The claim holds for $h = 0$ by Equation (54). As induction hypothesis, suppose the claim holds for some $h \in \mathbb{N}$. We obtain

$$x_v^{(h+1)} \stackrel{(3.11)}{=} \left(Ax^{(h)} \right)_v = \bigoplus_{w \in V} a_{vw} \odot x_w^{(h)} \stackrel{(3.9)}{=} \underbrace{\infty \odot x_v^{(h)}}_{x_v^{(h)}} \oplus \bigoplus_{\{v, w\} \in E} \omega(v, w) \odot x_w^{(h)}. \quad (\text{B.9})$$

Recall that $\oplus$ in $\mathcal{W}$ is the element-wise maximum by Corollary 3.11. Hence, we have

$$x_{vu}^{(h+1)} = \max \left\{ x_{vu}^{(h)} \right\} \cup \left\{ \min\{\omega(v, w), x_{wu}^{(h)}\} \mid \{v, w\} \in E \right\} \quad (\text{B.10})$$

and the induction hypothesis yields

$$x_{vu}^{(h+1)} = \max \left\{ \text{width}^h(v, u, G) \right\} \cup \left\{ \min\{\omega(v, w), \text{width}^h(w, u, G)\} \mid \{v, w\} \in E \right\}, \quad (\text{B.11})$$

which is exactly $\text{width}^{h+1}(v, u, G)$. $\square$

Proof of Lemma 3.18.

PROOF. We check the requirements of Definition A.2 in Appendix A step by step. Throughout the proof, let $\pi \in P$ and $x, y, z \in \mathcal{P}_{\min, +}$ be arbitrary.

- (1) We first show that $((\mathbb{R}_{\geq 0} \cup \{\infty\})^P, \oplus)$ is a commutative semigroup with neutral element 0. The associativity of $\oplus$ (and with it the property of $((\mathbb{R}_{\geq 0} \cup \{\infty\})^P, \oplus)$ being a semigroup) follows from the associativity of min:

$$((x \oplus y) \oplus z)_\pi = \min\{\min\{x_\pi, y_\pi\}, z_\pi\} = \min\{x_\pi, \min\{y_\pi, z_\pi\}\} = (x \oplus (y \oplus z))_\pi. \quad (\text{B.12})$$

Since min is commutative, $\oplus$ is, too, and it is easy to check that $(x \oplus 0)_\pi = (0 \oplus x)_\pi = x_\pi$.

- (2) To see that $((\mathbb{R}_{\geq 0} \cup \{\infty\})^P, \odot)$ is a semigroup with neutral element 1, we first check that $\odot$ is associative (i.e., that it is a semigroup):

$$((x \odot y) \odot z)_\pi = \min\{\min\{x_{\pi^1} + y_{\pi^2} \mid \pi^{12} = \pi^1 \circ \pi^2\} + z_{\pi^3} \mid \pi = \pi^{12} \circ \pi^3\} \quad (\text{B.13})$$

$$= \min\{(x_{\pi^1} + y_{\pi^2}) + z_{\pi^3} \mid \pi = (\pi^1 \circ \pi^2) \circ \pi^3\} \quad (\text{B.14})$$

$$= \min\{x_{\pi^1} + (y_{\pi^2} + z_{\pi^3}) \mid \pi = \pi^1 \circ (\pi^2 \circ \pi^3)\} \quad (\text{B.15})$$

$$= (x \odot (y \odot z))_\pi. \quad (\text{B.16})$$

Furthermore, $(1 \odot x)_\pi = \min\{0 + x_\pi\} = x_\pi = (x \odot 1)_\pi$, hence 1 is the neutral element with respect to $\odot$.

(3) Regarding the distributive laws, we begin with the left-distributive law (A.4):

$$(x \odot (y \oplus z))_\pi = \min\{x_{\pi^1} + \min\{y_{\pi^2}, z_{\pi^2}\} \mid \pi = \pi^1 \circ \pi^2\} \quad (\text{B.17})$$

$$= \min\{\min\{x_{\pi^1} + y_{\pi^2}, x_{\pi^1} + z_{\pi^2}\} \mid \pi = \pi^1 \circ \pi^2\} \quad (\text{B.18})$$

$$= \min\{\min\{x_{\pi^1} + y_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\}, \min\{x_{\pi^1} + z_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\}\} \quad (\text{B.19})$$

$$= ((x \odot y) \oplus (x \odot z))_\pi. \quad (\text{B.20})$$

Regarding the right-distributive law (A.5), we obtain:

$$((y \oplus z) \odot x)_\pi = \min\{\min\{y_{\pi^1}, z_{\pi^1}\} + x_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\} \quad (\text{B.21})$$

$$= \min\{\min\{y_{\pi^1} + x_{\pi^2}, z_{\pi^1} + x_{\pi^2}\} \mid \pi = \pi^1 \circ \pi^2\} \quad (\text{B.22})$$

$$= \min\{\min\{y_{\pi^1} + x_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\}, \min\{z_{\pi^1} + x_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\}\} \quad (\text{B.23})$$

$$= ((y \odot x) \oplus (z \odot x))_\pi. \quad (\text{B.24})$$

(4) It remains to check that 0 is an annihilator for $\odot$. We have

$$(0 \odot x)_\pi = \min\{0_{\pi^1} + x_{\pi^2} \mid \pi = \pi^1 \circ \pi^2\} = \min \emptyset = \infty = 0_\pi \quad (\text{B.25})$$

and, equivalently, $(x \odot 0)_\pi = 0_\pi$.

Hence, $\mathcal{P}_{\min,+}$ is a semiring as claimed. $\square$

Proof of Lemma 3.20.

PROOF. We prove the claim by induction. By Equation (63), the claim holds for $h = 0$. As induction hypothesis, suppose the claim holds for all $0 \leq h' \leq h$. The induction step yields

$$x_v^{(h+1)} \stackrel{(3.20)}{=} (Ax^{(h)})_v = \bigoplus_{w \in V} a_{vw} x_w^{(h)} \stackrel{(3.18)}{=} \underbrace{a_{vv}}_1 x_v^{(h)} \oplus \bigoplus_{\{v,w\} \in E} a_{vw} x_w^{(h)}. \quad (\text{B.26})$$

We have $a_{vv} x_v^{(h)} = 1x_v^{(h)} = x_v^{(h)}$ by construction; that is, $a_{vv} x_v^{(h)}$ contains exactly the properly weighted h -hop paths beginning at v by the induction hypothesis. Next, consider $\{v, w\} \in E$. By induction, $x_w^{(h)}$ contains exactly the h -hop paths beginning in w and a_{vw} contains only the edge $\{v, w\}$ of weight $\omega(v, w)$ by Equation (62). Hence, $a_{vw} x_w^{(h)}$ contains all $(h+1)$ -hop paths beginning with $\{v, w\}$. Due to Equation (B.26) and

$$P^{h+1}(v, \cdot, G) = P^h(v, \cdot, G) \cup \bigcup_{\{v,w\} \in E} \{(v, w) \circ \pi \mid \pi \in P^h(w, \cdot, G)\}, \quad (\text{B.27})$$

$x_v^{(h+1)}$ contains exactly the properly weighted $(h+1)$ -hop paths, as claimed. $\square$

Proof of Lemma 3.23.

PROOF. Fix a graph $G = (V, E, \omega)$. We show that $D(G)$ is closed under $\oplus$ and $\odot$. Consider $x, y \in D(G)$ and let $\pi \in P$ be a path. Hence, we have $x_\pi, y_\pi \in \{\omega(\pi), \infty\}$ if π is valid and $x_\pi = y_\pi = \infty$ if π is invalid in G ; recall that we defined $\omega(\pi) = \infty$ for invalid paths.

- (1) Consider $x \oplus y$. It directly follows from $x_\pi, y_\pi \in \{\omega(\pi), \infty\}$ that $(x \oplus y)_\pi = \min\{x_\pi, y_\pi\} \in \{\omega(\pi), \infty\}$. Hence, $(x \oplus y) \in D(G)$.
- (2) Regarding multiplication, we have $(x \odot y)_\pi = x_{\pi^1} + y_{\pi^2}$ for some two-split of π . Due to $x_{\pi^1} \in \{\omega(\pi^1), \infty\}$ and $y_{\pi^2} \in \{\omega(\pi^2), \infty\}$, we obtain $(x \odot y)_\pi = x_{\pi^1} + y_{\pi^2} \in \{\omega(\pi^1) + \omega(\pi^2), \infty\} = \{\omega(\pi), \infty\}$. If π is invalid in G , π^1 or π^2 must be invalid and $(x \odot y)_\pi = x_{\pi^1} + y_{\pi^2} = \infty$ follows.

Also observe that $0, 1 \in D(G)$ for 0 and 1 from Lemma 3.18. Together, it follows from Lemma 3.18 that $\mathcal{P}_{\min,+}(G)$ is a semiring. $\square$

Proof of Lemma 3.25. Let π be a v - w -path and $y \in \mathcal{P}_{\min,+}(G)$ such that y contains only π (i.e., with $y_\pi = \omega(\pi)$ and $y_{\pi'} = \infty$ for all $\pi' \neq \pi$). We say that π is *dominated in x* if and only if $r(x \oplus y)_\pi = \infty$, meaning that making π available in x does not change the outcome of filtering. Regarding the k -SDP, this is the case where either (1) π does not end in s , or (2) x contains k other v - s -paths that are shorter than π or have the same weight as π but are lexicographically ordered before π . While the notion of domination may seem overly complicated for the matter at hand, it sufficiently generalizes the proof for Lemma 3.25 to cover the k -SDP as well as the k -DSDP from Examples 3.27 and 3.26, respectively.

OBSERVATION B.1. *Note that r discards all paths not ending in s . Furthermore, for all paths π , we have*

$$r(x)_\pi = x_\pi \quad \text{or} \quad r(x)_\pi = \infty; \quad (\text{B.28})$$

that is, $r(x) \in \mathcal{P}_{\min,+}(G)$ for all $x \in \mathcal{P}_{\min,+}(G)$, as r either keeps a path with its original weight or discards it by setting the according entry to ∞ . We also obtain that, for all $v \in V$,

$$P_k(v, s, x) = P_k(v, s, r(x)) \quad (\text{B.29})$$

because $P_k(v, s, x)$ is invariant under discarding dominated paths from x .

PROOF. Clearly, r is a projection. Below, we show in one step each that it fulfills Conditions (2.12) and (2.13) of Lemma 2.8. Throughout the proof, fix a graph $G = (V, E, \omega)$, let $x, x', y, y' \in \mathcal{P}_{\min,+}(G)$ be such that $x \sim x'$ and $y \sim y'$.

- (1) We show below that $r(yx) = r(yr(x))$. Equation (18) then follows using that $r(yx) = r(yr(x)) = r(yr(x')) = r(yx')$. To see that $r(yx) = r(yr(x))$, we argue that, for all v - s -paths π , we either have $(yx)_\pi = (yr(x))_\pi$, or both $\pi \notin P_k(v, s, yx)$ and $\pi \notin P_k(v, s, yr(x))$. In other words, the entries regarding π are either equal in yx and $yr(x)$ or r discards π from yx as well as from $yr(x)$.

Consider a v - s -path π with $(yx)_\pi \neq (yr(x))_\pi$. Observe that this implies $\omega(\pi) = (yx)_\pi < (yr(x))_\pi = \infty$ because the non- ∞ entries of $r(x)$ are a subset of those of x . Hence, π is contained in yx . By definition of $\odot$, it holds that $(yx)_\pi = y_{\pi^1} + x_{\pi^2}$ for some partition $\pi = \pi^1 \odot \pi^2$, where π^1 and π^2 are, for some node w , v - w - and w - s -paths, respectively. It follows from Equation (B.28) and $x \in \mathcal{P}_{\min,+}(G)$ that

$$\omega(\pi^2) = x_{\pi^2} < r(x)_{\pi^2} = \infty. \quad (\text{B.30})$$

Hence, $\pi^2 \notin P_k(w, s, x)$, i.e., π^2 is dominated in x . As $P_k(w, s, r(x)) = P_k(w, s, x)$ by Equation (B.29), the following k v - s -paths are contained in $yr(x)$ and in yx each:

$$\left\{ \pi^1 \odot \bar{\pi}^2 \mid \bar{\pi}^2 \in P_k(w, s, r(x)) \right\}. \quad (\text{B.31})$$

We conclude that π is dominated in—and thus discarded from—both yx and $yr(x)$, as claimed.

- (2) Consider a node $v \in V$ and a v - s -path π . As $(x \oplus y)_\pi = \min\{x_\pi, y_\pi\}$, we have $P(v, s, x \oplus y) = P(v, s, x) \cup P(v, s, y)$. In particular, it holds that $P_k(v, s, x \oplus y) \subseteq P_k(v, s, x) \cup P_k(v, s, y)$ because if $\pi \in P(v, s, x \oplus y)$ is dominated in x and in y , it is dominated in $x \oplus y$ as well. Using Equation (B.29), we obtain that

$$P_k(v, s, x \oplus y) \subseteq P_k(v, s, x) \cup P_k(v, s, y) = P_k(v, s, r(x)) \cup P_k(v, s, r(y)). \quad (\text{B.32})$$

As r discards all paths not ending in s , and $x_\pi = r(x)_\pi$ for all $\pi \in P_k(v, s, r(x))$ and, analogously, $y_\pi = r(y)_\pi$ for all $\pi \in P_k(v, s, r(y))$, we conclude that $r(x \oplus y) = r(r(x) \oplus r(y))$. Hence, $r(x \oplus y) = r(r(x) \oplus r(y)) = r(r(x') \oplus r(y')) = r(x' \oplus y')$; that is, r fulfills Equation (2.13).

Since $x \sim x'$ and $y \sim y'$ are arbitrary, r fulfills the preconditions of Lemma 2.8 and the claim follows. $\square$

REFERENCES

- [1] Amir Abboud, Greg Bodwin, and Seth Pettie. 2017. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*. SIAM, Philadelphia, PA, USA, 568–576. DOI : <http://dx.doi.org/10.1137/1.9781611974782.36>
- [2] Miklós Ajtai, János Komlós, and Endre Szemerédi. 1983. An $O(n \log n)$ sorting network. In *Proceedings of the 15th ACM Symposium on Theory of Computing (STOC)*. ACM, New York, NY, USA, 1–9. DOI : <http://dx.doi.org/10.1145/800061.808726>
- [3] Noga Alon, Zvi Galil, and Oded Margalit. 1997. On the exponent of the all pairs shortest path problem. *Journal of Computer and System Science* 54, 2 (1997), 255–262. DOI : <http://dx.doi.org/10.1006/jcss.1997.1388>
- [4] Noga Alon, Richard M. Karp, David Peleg, and Douglas B. West. 1995. A graph-theoretic game and its application to the k -server problem. *SIAM Journal on Computing* 24, 1 (1995), 78–100. DOI : <http://dx.doi.org/10.1137/S0097539792224474>
- [5] Matthew Andrews. 2004. Hardness of buy-at-bulk network design. In *Proceedings of the 45th Symposium on Foundations of Computer Science (FOCS)*. IEEE, Los Alamitos, CA, USA, 115–124. DOI : <http://dx.doi.org/10.1109/FOCS.2004.32>
- [6] Baruch Awerbuch and Yossi Azar. 1997. Buy-at-bulk network design. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*. 542–547. DOI : <http://dx.doi.org/10.1109/SFCS.1997.646143>
- [7] Yair Bartal. 1996. Probabilistic approximations of metric spaces and its algorithmic applications. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, Los Alamitos, CA, USA, 184–193. DOI : <http://dx.doi.org/10.1109/SFCS.1996.548477>
- [8] Yair Bartal. 1998. On approximating arbitrary metrics by tree metrics. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing (STOC)*. ACM, New York, NY, USA, 161–168. DOI : <http://dx.doi.org/10.1145/276698.276725>
- [9] Surender Baswana and Sandeep Sen. 2007. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Structures & Algorithms* 30, 4 (2007), 532–563. DOI : <http://dx.doi.org/10.1002/rsa.20130>
- [10] Ruben Becker, Andreas Karrenbauer, Sebastian Krinninger, and Christoph Lenzen. 2016. Near-optimal approximate shortest paths and transshipment in distributed and streaming models. In *31st International Symposium on Distributed Computing (DISC'17) (Leibniz International Proceedings in Informatics (LIPIcs))*, Andréa W. Richa (Ed.), Vol. 91. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 7:1–7:16. <https://doi.org/10.4230/LIPIcs.DISC.2017.7>. <http://arxiv.org/abs/1607.05127>
- [11] Richard E. Bellman. 1958. On a routing problem. *Quarterly Applied Mathematics* 16 (1958), 87–90.
- [12] Guy E. Blelloch, Yan Gu, Julian Shun, and Yihan Sun. 2016. Parallelism in randomized incremental algorithms. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. ACM, New York, NY, USA, 467–478. DOI : <http://dx.doi.org/10.1145/2935764.2935766>
- [13] Guy E. Blelloch, Yan Gu, and Yihan Sun. 2016. Efficient construction of probabilistic tree embeddings. *CoRR* abs/1605.04651 (2016). <http://arxiv.org/abs/1605.04651>.
- [14] Guy E. Blelloch, Anupam Gupta, and Kanat Tangwongsan. 2012. Parallel probabilistic tree embeddings, k -median, and buy-at-bulk network design. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. ACM, New York, NY, USA, 205–213. DOI : <http://dx.doi.org/10.1145/2312005.2312045>
- [15] Guy E. Blelloch and Kanat Tangwongsan. 2010. Parallel approximation algorithms for facility-location problems. In *Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. ACM, New York, NY, USA, 315–324. DOI : <http://dx.doi.org/10.1145/1810479.1810535>
- [16] Edith Cohen. 1997. Size-estimation framework with applications to transitive closure and reachability. *Journal of Computer and System Sciences* 55, 3 (1997), 441–453. DOI : <http://dx.doi.org/10.1006/jcss.1997.1534>
- [17] Edith Cohen. 2000. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *Journal of the ACM* 47, 1 (2000), 132–166. DOI : <http://dx.doi.org/10.1145/331605.331610>

- [18] Edith Cohen and Haim Kaplan. 2007. Spatially-decaying aggregation over a network. *Journal of Computer and System Sciences* 73, 3 (2007), 265–288. DOI : <http://dx.doi.org/10.1016/j.jcss.2006.10.016>
- [19] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms* (3rd edition). MIT Press. <http://mitpress.mit.edu/books/introduction-algorithms>.
- [20] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. 2012. Distributed verification and hardness of distributed approximation. *SIAM Journal on Computing* 41, 5 (2012), 1235–1265. DOI : <http://dx.doi.org/10.1137/11085178X>
- [21] Edsger W. Dijkstra. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 1 (1959), 269–271. DOI : <http://dx.doi.org/10.1007/BF01386390>
- [22] Michael Elkin and Ofer Neiman. 2016. Hopsets with constant hopbound, and applications to approximate shortest paths. In *Proceedings of the 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, Los Alamitos, CA, USA, 128–137. DOI : <http://dx.doi.org/10.1109/FOCS.2016.22>
- [23] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. 2004. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences* 69, 3 (2004), 485–497. DOI : <http://dx.doi.org/10.1016/j.jcss.2004.04.011>
- [24] Lester R. Ford. 1956. *Network Flow Theory*. Technical Report. The RAND Corporation.
- [25] Stephan Friedrichs and Christoph Lenzen. 2016. Parallel metric tree embedding based on an algebraic view on Moore-Bellman-Ford. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. ACM, New York, NY, USA, 455–466. DOI : <http://dx.doi.org/10.1145/2935764.2935777>
- [26] Mohsen Ghaffari and Christoph Lenzen. 2014. Near-optimal distributed tree embedding. In *Proceedings of the 28th International Symposium on Distributed Computing (DISC)*. Springer, Berlin, Heidelberg, 197–211. DOI : http://dx.doi.org/10.1007/978-3-662-45174-8_14
- [27] Mathias Hauptmann and Marek Karpinski. A Compendium on Steiner Tree Problems. <http://theory.cs.uni-bonn.de/info5/steinerkompndium/netcompendium.html>. Visited 2016-05-12.
- [28] Udo Hebisch and Hanns J. Weinert. 1998. *Semirings: Algebraic Theory and Applications in Computer Science*. World Scientific.
- [29] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. 2016. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In *Proceedings of the 48th Symposium on Theory of Computing (STOC)*. ACM, New York, NY, USA, 489–498. DOI : <http://dx.doi.org/10.1145/2897518.2897638>
- [30] Maleq Khan, Fabian Kuhn, Dahlia Malkhi, Gopal Pandurangan, and Kunal Talwar. 2012. Efficient distributed approximation algorithms via probabilistic tree embeddings. *Distributed Computing* 25, 3 (2012), 189–205. DOI : <http://dx.doi.org/10.1007/s00446-012-0157-9>
- [31] Philip N. Klein and Sairam Subramanian. 1997. A randomized parallel algorithm for single-source shortest paths. *Journal of Algorithms* 25, 2 (1997), 205–220. DOI : <http://dx.doi.org/10.1006/jagm.1997.0888>
- [32] François Le Gall. 2014. Powers of tensors and fast matrix multiplication. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation (ISSAC)*. ACM, New York, NY, USA, 296–303. DOI : <http://dx.doi.org/10.1145/2608628.2608664>
- [33] Christoph Lenzen and Boaz Patt-Shamir. 2013. Fast routing table construction using small messages: Extended abstract. In *Proceedings of the Symposium on Theory of Computing Conference (STOC)*. ACM, New York, NY, USA, 381–390. DOI : <http://dx.doi.org/10.1145/2488608.2488656>
- [34] Christoph Lenzen and Boaz Patt-Shamir. 2014. Improved distributed Steiner forest construction. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*. ACM, New York, NY, USA, 262–271. DOI : <http://dx.doi.org/10.1145/2611462.2611464>
- [35] Christoph Lenzen and Boaz Patt-Shamir. 2015. Fast partial distance estimation and applications. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*. ACM, New York, NY, USA, 153–162. DOI : <http://dx.doi.org/10.1145/2767386.2767398>
- [36] Christoph Lenzen and David Peleg. 2013. Efficient distributed source detection with limited bandwidth. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*. ACM, New York, NY, USA, 375–382. DOI : <http://dx.doi.org/10.1145/2484239.2484262>
- [37] Manor Mendel and Chaya Schwob. 2009. Fast C-K-R partitions of sparse graphs. *Chicago Journal of Theoretical Computer Science* 2009, 2 (2009), 1–18. <http://cjtc.cs.uchicago.edu/articles/2009/2/contents.html>.
- [38] Ramgopal R. Mettu and C. Greg Plaxton. 2004. Optimal time bounds for approximate clustering. *Machine Learning* 56, 1–3 (2004), 35–60. DOI : <http://dx.doi.org/10.1023/B:MACH.0000033114.18632.e0>
- [39] Mehryar Mohri. 2002. Semiring frameworks and algorithms for shortest-distance problems. *Journal of Automata, Languages and Combinatorics* 7, 3 (2002), 321–350.
- [40] Edward F. Moore. 1959. The shortest path through a maze. In *Proceedings of the Symposium on the Theory of Switching*. 87–90.

- [41] David Peleg. 2000. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia PA, USA.
- [42] Hanmao Shi and Thomas H. Spencer. 1999. Time-work tradeoffs of the single-source shortest paths problem. *Journal of Algorithms* 30, 1 (1999), 19–32. DOI: <http://dx.doi.org/10.1006/jagm.1998.0968>
- [43] Uri Zwick. 2002. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *Journal of the ACM* 49, 3 (2002), 289–317. DOI: <http://dx.doi.org/10.1145/567112.567114>

Received August 2016; revised July 2018; accepted August 2018