

A DETERMINISTIC ALMOST-TIGHT DISTRIBUTED ALGORITHM FOR APPROXIMATING SINGLE-SOURCE SHORTEST PATHS*

MONIKA HENZINGER[†], SEBASTIAN KRINNINGER[‡], AND DANUPON NANONGKAI[§]

Abstract. We present a deterministic $(1 + o(1))$ -approximation $(n^{1/2+o(1)} + D^{1+o(1)})$ -time algorithm for solving the single-source shortest paths problem on distributed weighted networks (the CONGEST model); here n is the number of nodes in the network, D is its (hop) diameter, and edge weights are positive integers from 1 to $\text{poly}(n)$. This is the first non-trivial deterministic algorithm for this problem. It also improves (i) the running time of the randomized $(1 + o(1))$ -approximation $\tilde{O}(n^{1/2}D^{1/4} + D)$ -time¹ algorithm of Nanongkai [STOC 2014] by a factor of as large as $n^{1/8}$, and (ii) the $O(\epsilon^{-1} \log \epsilon^{-1})$ -approximation factor of Lenzen and Patt-Shamir’s $\tilde{O}(n^{1/2+\epsilon} + D)$ -time algorithm [STOC 2013] within the same running time. Our running time matches the known time lower bound of $\Omega(n^{1/2}/\log n + D)$ [Das Sarma et al. STOC 2011] modulo sub-polynomial terms, thus essentially settling the status of this problem which was raised at least a decade ago [Elkin SIGACT News 2004]. It also implies a $(2 + o(1))$ -approximation $(n^{1/2+o(1)} + D^{1+o(1)})$ -time algorithm for approximating a network’s weighted diameter which almost matches the lower bound by Holzer and Pinski [OPODIS 2015].

In achieving this result, we develop two techniques which might be of independent interest and useful in other settings: (i) a deterministic process that replaces the “hitting set argument” commonly used for shortest paths computation in various settings, and (ii) a simple, deterministic, construction of an $(n^{o(1)}, o(1))$ -hop set of size $n^{1+o(1)}$. We combine these techniques with many distributed algorithmic techniques, some of which from problems that are not directly related to shortest paths, e.g. ruling sets [Goldberg et al. STOC 1987], source detection [Lenzen, Peleg PODC 2013], and partial distance estimation [Lenzen, Patt-Shamir PODC 2015]. Our hop set construction also leads to single-source shortest paths algorithms in two other settings: (i) a $(1 + o(1))$ -approximation $n^{o(1)}$ -time algorithm on *congested cliques*, and (ii) a $(1 + o(1))$ -approximation $n^{o(1)}$ -pass $n^{1+o(1)}$ -space *streaming* algorithm. The first result answers an open problem in [Nanongkai, STOC 2014]. The second result partially answers an open problem raised by McGregor in 2006 [sublinear.info, Problem 14].

Key words. distributed algorithms, shortest paths

AMS subject classifications. 68W15, 05C85, 05C12

1. Introduction. In the area of *distributed graph algorithms* we study the complexity for a network to compute its own topological properties, such as minimum spanning tree, maximum matching, or distances between nodes. A fundamental question in this area that has been studied for many years is how much *time complexity* is needed to solve a problem in the so-called *CONGEST model* (e.g. [32, 62, 21, 15, 52]). In this model (see Section 2 for details), a network is modeled by a weighted undirected graph G , where each node represents a processor that initially only knows its adjacent edges and their weight, and nodes must communicate with each other over *bounded-bandwidth* links to discover global topological properties of the network. The

*Submitted to the editors October 2016.

Funding: The research leading to this work has received funding from the European Union’s Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 317532 and from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement number 340506. Support by Swedish Research Council grant 2015-04659 “Algorithms and Complexity for Dynamic Graph Problems.” M. Henzinger’s work was done in part while visiting the Simons Institute for Theory of Computing. S. Krinninger’s work was partially done while visiting the Simons Institute for the Theory of Computing and while at Max Planck Institute for Informatics, Saarland Informatics Campus, Germany.

[†]University of Vienna, Faculty of Computer Science, Austria

[‡]University of Vienna, Faculty of Computer Science, Austria

[§]KTH Royal Institute of Technology, Sweden.

¹Throughout, we use $\tilde{O}(\cdot)$ to hide polylogarithmic factors in n .

communication between nodes is carried out in *rounds*, where in each round each node can send a small, logarithmic-size message to each neighbor. The time complexity is measured as the number of rounds needed to finish the task. It is usually measured by n , the number of nodes in the network, and D , the diameter of the communication network (when edge weights are omitted). Typically, $D \ll n$.

In this paper, we consider the problem of *approximating single-source shortest paths* (SSSP). In this problem, a node s is marked as the *source node*, and the goal is for every node to know how far it is from s . The unweighted version – the *breadth-first search tree* computation – is one of the most basic tools in distributed computing, and is well known to require $\Theta(D)$ time (see e.g. [61]). In contrast, the only available solution for the weighted case is the distributed version of the Bellman-Ford algorithm [8, 30], which takes $O(n)$ time to compute an exact solution. In 2004, Elkin [20] raised the question whether distributed *approximation* algorithms can help improving this time complexity and showed that any α -approximation algorithm requires $\Omega((n/\alpha)^{1/2}/\log n + D)$ time [21]. Das Sarma et al. [15] (building on [62, 48]) later strengthened this lower bound by showing that any poly(n)-approximation (randomized) algorithm requires $\Omega(n^{1/2}/\log n + D)$ time. This lower bound was later shown to hold even for quantum algorithms [24].

Since running times of the form $\tilde{O}(n^{1/2} + D)$ show up in many distributed algorithms (e.g. MST [49, 62], connectivity [70, 63], and minimum cut [60, 34]) it is natural to ask whether the lower bound of [15] can be matched. The first answer to this question is a randomized $O(\epsilon^{-1} \log \epsilon^{-1})$ -approximation $\tilde{O}(n^{1/2+\epsilon} + D)$ -time algorithm by Lenzen and Patt-Shamir [52]². The running time of this algorithm is nearly tight if we are satisfied with a large approximation ratio. For a small approximation ratio, Nanongkai [58] presented a randomized $(1+o(1))$ -approximation $\tilde{O}(n^{1/2}D^{1/4} + D)$ -time algorithm. The running time of this algorithm is nearly tight when D is small, but can be close to $\tilde{\Theta}(n^{2/3})$ even when $D = o(n^{2/3})$. This created a rather unsatisfying situation: First, one has to sacrifice a large approximation factor in order to achieve the near-optimal running time, and to achieve a $(1+o(1))$ approximation factor, one must pay an additional running time of $D^{1/4}$ which could be as far from the lower bound as $n^{1/8}$ when D is large. Because of this, the question whether we can close the gap between upper and lower bounds for the running time of $(1+o(1))$ -approximation algorithms was left as the main open problem in [58, Problem 7.1]. Secondly, and more importantly, both these algorithms are randomized. Given that designing deterministic algorithms is an important issue in distributed computing, this leaves the important open problem whether there is a deterministic algorithm that is faster than Bellman-Ford’s algorithm, i.e., that runs in *sublinear-time*.

Our Results. In this paper, we resolve the two issues above. We present a deterministic $(1+o(1))$ -approximation $(n^{1/2+o(1)} + D^{1+o(1)})$ -time algorithm for this problem (the $o(1)$ term in the approximation ratio hides a $1/\text{polylog } n$ factor and the $o(1)$ term in the running time hides an $O(\sqrt{\log \log n / \log n})$ factor). Our algorithm almost settles the status of this problem as its running time matches the lower bound of Das Sarma et al. [15] up to an $n^{o(1)}$ factor.

Since an α -approximate solution to SSSP gives a 2α -approximate value of the network’s *weighted diameter* (cf. Section 2), our algorithm can $(2+o(1))$ -approximate the weighted diameter within the same running time. Previously, Holzer and Pinsker [41] (building on [42]) showed that for any $\epsilon > 0$, a $(2-\epsilon)$ -approximation algorithm for this problem requires $\tilde{\Omega}(n)$ time. Thus, the approximation ratio provided by our

²Note that the result of Lenzen and Patt-Shamir in fact solves a more general problem.

algorithm cannot be significantly improved without increasing the running time. The running time of our algorithm also cannot be significantly improved because of the lower bound of $\Omega(n^{1/2}/\log n + D)$ [15] for approximate SSSP which holds for any $\text{poly}(n)$ -approximation algorithm.

Using the same techniques, we also obtain a deterministic $(1+o(1))$ -approximation $n^{o(1)}$ -time algorithm for the special case of *congested clique*, where the underlying network is fully connected. This gives a positive answer to Problem 7.5 in [58]. Previous algorithms solved this problem exactly in time $\tilde{O}(n^{1/2})$ [58] and $\tilde{O}(n^{1/3})$ [11], respectively, and $(1+o(1))$ -approximately in time $O(n^{0.158})$ [11]³. We can also compute a $(2+o(1))$ -approximation of the weighted diameter within the same running time. The lower bound of Holzer and Pinski [41] also applies in this setting: computing a $(2-o(1))$ -approximation of the diameter requires $\tilde{\Omega}(n)$ time in the worst case.

Our techniques also lead to a (non-distributed) *streaming algorithm* for $(1+o(1))$ -approximate SSSP where the edges are presented in an arbitrary-order stream, and an algorithm with limited space (preferably $\tilde{O}(n)$) reads the stream in *passes* to determine the answer (see, e.g., [56] for a recent survey). It was known that $\tilde{O}(n)$ space and one pass are enough to compute an $O(\log n/\log \log n)$ -spanner and therefore approximate all distances up to a factor of $O(\log n/\log \log n)$ [29] (see also [28, 5, 27, 22]). This almost matches a lower bound which holds even for the s - t -shortest path problem (stSP), where we just want to compute the distance between two specific nodes s and t [29]. On unweighted graphs one can compute $(1+\epsilon, \beta)$ -spanners in β passes and $O(n^{1+1/k})$ space [27] (for some integer β depending on k and ϵ), and get $(1+\epsilon)$ -approximate SSSP in a total of $O(\beta/\epsilon)$ passes. In 2006, McGregor raised the question whether we can solve stSP better with a larger number of passes (see [1]). Very recently Guruswami and Onak [37] showed that any p -pass algorithm on unweighted graphs requires $\tilde{\Omega}(n^{1+\Omega(1/p)})/O(p)$ space. This does not rule out, for example, an $O(\log n)$ -pass $\tilde{O}(n)$ -space algorithm. Our algorithm, which solves the more general SSSP problem, gets close to this: it takes $n^{o(1)}$ passes and $n^{1+o(1)}$ space.

In all these models, we have formulated our algorithms to compute $(1+o(1))$ -approximate SSSP. More generally, we can, for any $0 < \epsilon \leq 1$ compute a $(1+\epsilon)$ -approximation taking $(\sqrt{n}+D) \cdot 2^{O(\sqrt{\log n \log(\epsilon^{-1} \log n)})}$ rounds in the CONGEST model, $2^{O(\sqrt{\log n \log(\epsilon^{-1} \log n)})}$ rounds in the congested clique model, and $2^{O(\sqrt{\log n \log(\epsilon^{-1} \log n)})}$ passes with $n \cdot 2^{O(\sqrt{\log n \log(\epsilon^{-1} \log n)})}$ space in the streaming model, respectively. We provide the necessary details for deriving these numbers in Section 3, but omit them later on for the sake of succinctness. We assume throughout that the edge weights are positive integers in the range $\{1, \dots, W\}$ where W is polynomial in n . More generally, for W of arbitrary size, all of the above asymptotic bounds need to be multiplied by the factor $\log W$.

Overview of Techniques. Our crucial new technique is a deterministic process that can replace the following “path hitting” argument: For any c , if we pick $\tilde{\Theta}(c)$ nodes uniformly at random as *centers* (typically $c = n^{1/2}$), then a shortest path containing n/c edges will contain a center with high probability. This allows us to create shortcuts between centers – where we replace each path of length n/c between centers by an edge of the same length – and focus on computing shortest paths between centers. This argument has been repetitively used to solve shortest paths problems in various settings (e.g. [72, 38, 18, 6, 65, 66, 16, 17, 57, 10, 52, 58]). In the sequential model a

³With this running time, [11] can in fact solve the all-pairs shortest paths problem. See also [50] for further developments in the direction of [11].

set of centers of size $\tilde{\Theta}(c)$ can be found deterministically with the greedy hitting set heuristic once the shortest paths containing n/c edges are known [73, 46]. We are not aware of any non-trivial deterministic process that can achieve the same effect in the distributed setting. The main challenge is that the greedy process is heavily sequential as the selection of the next node depends on all previous nodes, and is thus hard to implement efficiently in the distributed setting⁴.

In this paper, we develop a new deterministic process to pick $\tilde{\Theta}(c)$ centers. The key new idea is to carefully divide nodes into $O(\log n)$ *types*. Roughly speaking, we associate each type t with a value w_t and make sure that the following properties hold: (i) every path π with $\Omega(n/c)$ edges and weight $\Theta(w_t)$ contains a node of type t , and (ii) there is a set of $O(n/c)$ *centers of type t* such that every node of type t has at least one center at distance $o(w_t)$. We define the set of centers to be the collection of centers of all types. The two properties together guarantee that every long path will be *almost* hit by a center: for every path π containing at least n/c edges, there is a center whose distance to some node in π is $o(w(\pi))$ where, $w(\pi)$ is the total weight of π . This is already sufficient for us to focus on computing shortest paths only between centers as we would have done after picking centers using the path hitting argument. To the best of our knowledge, such a deterministically constructed set of centers that almost hits *every* long path was not known to exist before. The process itself is not constrained to the distributed setting and thus might be useful for derandomizing other algorithms that use the path hitting argument.

To implement the above process in the distributed setting, we use the *source detection* algorithm of Lenzen and Peleg [54] to compute the type of each node. We then use the classic *ruling set* algorithm of Goldberg et al. [36] to compute the set of centers of each type that satisfies the second property above. (A technical note: we also need to compute a bounded-depth shortest-path tree from every center. In [58], this was done using the *random delay* technique. We also derandomize this step by adapting the *partial distance estimation algorithm* of Lenzen and Patt-Shamir [53].)

Another tool, which is the key to the improved running time, is a new *hop set* construction. An (h, ϵ) -hop set of a graph $G = (V, E)$ is a set F of weighted edges such that the distance between any pair of nodes in G can be $(1 + \epsilon)$ -approximated by their h -hop distance (given by a path containing at most h edges) on the graph $H = (V, E \cup F)$ (see Section 2 for details). The notion of hop set was defined by Cohen [13] in the context of parallel computing, although it has been used implicitly earlier, e.g. [72, 47] (see [13] for a detailed discussion). The previous SSSP algorithm [58] was able to construct an $(n/k, 0)$ -hop set of size kn , for any integer $k \geq 1$, as a subroutine (in [58] this was called *shortest paths diameter reduction*⁵). In this paper, we show that this subroutine can be replaced by the construction of an $(n^{o(1)}, o(1))$ -hop set of size $n^{1+o(1)}$.

Our hop set construction is based on computing *clusters* which is the basic subroutine of Thorup and Zwick’s distance oracles [68] and spanners [68, 69]. It builds on a line of work in dynamic graph algorithms. In [9], Bernstein showed that clusters can be used to construct an $(n^{o(1)}, o(1))$ -hop set of size $n^{1+o(1)}$. Later in [39], we showed that the same kind of hop set can be constructed by using a structure similar to clusters while restricting the shortest-path trees involved to some small distance

⁴We note that the algorithm of King [46] for constructing a blocker can be viewed as an efficient way to greedily pick a hitting set by efficiently computing the scores of nodes. The process is as highly sequential like other greedy heuristics.

⁵This follows the notion of shortest paths diameter used earlier in distributed computing [45]

and use such a construction in the dynamic (more precisely, decremental) setting. The construction, however, has to deal with several complications of the dynamic setting and heavily relies on randomization. In this paper, we build on the same idea, i.e., we construct a hop set using bounded-distance clusters. However, our construction is significantly simplified, to the point that we can treat the cluster computation as a black box. This makes it easy to apply on distributed networks and to derandomize. To this end, we derandomize the construction simply by invoking the deterministic clusters construction of Roditty, Thorup, and Zwick [64] and observe that it can be implemented efficiently on distributed networks⁶. A similar type of derandomization by locally computing approximate hitting sets has been done before by Holzer and Pinski [41] when derandomizing Nanongkai’s exact hop set construction [58] on the congested clique. We note that it might be possible to use Cohen’s hop set construction [13] instead of Bernstein’s [9] in our application. However, Cohen’s construction heavily relies on randomness and derandomizing it seems significantly more difficult.

Recent Developments. After the preliminary version of this paper appeared [40], Becker et al. [7] showed that the $n^{o(1)}$ term in our bounds can be eliminated. Elkin and Neiman showed the first construction of sparse hop sets with a constant number of hops [25], removing also the inherent dependence on $\log W$, the logarithm of the largest edge weight, in their construction. The latter carries over to the bounds for approximating single-source shortest paths in the congested clique model and the multipass streaming model. They further showed an application of their hop sets in computing approximate shortest paths from s sources. In particular, using our hop set and a modification of the framework in [58] and this paper⁷, this problem can be solved in $(sn)^{1/2+o(1)} + D^{o(1)}$ rounds. Elkin and Neiman showed a hop set which can be used to reduce the bound to $\tilde{O}((sn)^{1/2} + D)$ when $s = n^{\Omega(1)}$ [25]. In [26], they also showed further applications of hop sets in the distributed construction of routing schemes. It was pointed out by Patt-Shamir (see [71]) that using our algorithm as a black box, one can simplify and obtain improved running time in the construction of compact routing tables in [53]. (On the other hand, we note that our construction is based on many ideas from [53].) Our hop set construction also found applications in metric-tree embeddings [31].

Organization. We start by introducing notation and the main definition in Section 2. Then in Section 3 we explain the deterministic hop set construction in, which is based on a variation of Thorup and Zwick’s clusters [68]. In Section 4, we give our main result, namely the $(1 + o(1))$ -approximation $(n^{1/2+o(1)} + D^{1+o(1)})$ -time algorithm. In that section we explain the deterministic process for selecting centers mentioned above, as well as how to implement the hop set construction in the distributed setting. Finally, our remaining results are proved in Section 5.

2. Preliminaries.

2.1. Notation. In this paper, we consider weighted undirected graphs with positive integer edge weights in the range $\{1, 2, \dots, W\}$. We usually assume in the following that $W = \text{poly}(n)$, i.e., the edge weights are *polynomially bounded*. For a graph $G = (V, E)$, V is the set of nodes and E is the set of edges. We denote by $n := |V|$ and $m := |E|$ the number of nodes and edges of G , respectively. For a set of edges E , the weight of each edge $(u, v) \in E$ is given by a function $w(u, v, E)$. If

⁶We note that Thorup-Zwick’s distance oracles and spanners were considered before in the distributed setting (e.g. [53, 14]).

⁷More precisely, following Elkin and Neiman [25], one constructs an overlay network of size \sqrt{sn} instead of \sqrt{n} as done in this paper.

$(u, v) \notin E$, we set $w(u, v, E) = \infty$. We define $w(u, v, G) := w(u, v, E)$. Whenever we define a set of edges E as the union of two sets of edges $E_1 \cup E_2$, we set the weight of every edge $(u, v) \in E$ to $w(u, v, E) := \min(w(u, v, E_1), w(u, v, E_2))$. We denote the weight of a path π in a graph G by $w(\pi, G) := \sum_{(u,v) \in \pi} w(u, v, G)$ and the number of edges of π by $|\pi|$.

Given a graph $G = (V, E)$ and a set of edges $F \subseteq V^2$, we define $G \cup F$ as the graph that has V as its set of nodes and $E \cup F$ as its set of edges. The weight of each edge (u, v) is given by $w(u, v, G \cup F) = w(u, v, E \cup F) = \min(w(u, v, E), w(u, v, F))$.

We denote the distance between two nodes u and v in G , i.e., the weight of the shortest path between u and v , by $d(u, v, G)$. We define the distance between a node u and a set of nodes $A \subseteq V$ by $d(u, A, G) = \min_{v \in A} d(u, v, G)$. For every pair of nodes u and v we define distance up to range R by

$$d(u, v, R, G) = \begin{cases} d(u, v, G) & \text{if } d(u, v, G) \leq R \\ \infty & \text{otherwise.} \end{cases}$$

and for a node v and set of nodes $A \subseteq V$ by $d(u, A, R, G) = \min_{v \in A} d(u, v, R, G)$.

For any positive integer h and any nodes u and v , we define the h -hop distance between u and v , denoted by $d^h(u, v, G)$, as the weight of the shortest among all u - v paths containing at most h edges. More precisely, let $\Pi^h(u, v)$ be the set of all paths between u and v such that each path $\pi \in \Pi^h(u, v)$ contains at most h edges. Then, $d^h(u, v, G) = \min_{\pi \in \Pi^h(u, v)} w(\pi, G)$ if $\Pi^h(u, v) \neq \emptyset$, and $d^h(u, v, G) = \infty$ otherwise.

We denote the hop-distance between two nodes u and v , i.e., the distance between u and v when we treat G as an unweighted graph, by $\text{hop}(u, v, G)$. The hop diameter of graph G is defined as $D(G) = \max_{u, v \in V} \text{hop}(u, v, G)$. When G is clear from the context, we use D instead of $D(G)$. We note that this is different from the *weighted diameter*, which is defined as $WD(G) = \max_{u, v \in V} d(u, v, G)$. Throughout this paper we use ‘‘diameter’’ to refer to the hop diameter (as it is typically done in the literature, see, e.g., [32, 49, 44, 55, 34]). We do not consider superlogarithmic values for the bandwidth B in this paper.

The following definition formalizes the concept of hop sets introduced by Cohen [13].

DEFINITION 1. *Given any graph $G = (V, E)$, any integer h , and $\epsilon \geq 0$, we say that a set of weighted edges F is an (h, ϵ) -hop set of G if*

$$d(u, v, G) \leq d^h(u, v, H) \leq (1 + \epsilon)d(u, v, G)$$

for every pair of nodes $u, v \in V$, where $H = (V, E \cup F)$.

In this paper we are only interested in $(n^{o(1)}, o(1))$ -hop sets of size $|F| = n^{1+o(1)}$. We refer to them simply as ‘‘hop sets’’ (without specifying parameters).

2.2. CONGEST Model and Problem Formulation. In the CONGEST model, a network of processors is modeled by an undirected weighted graph G , where nodes model the processors and edges model the bounded-bandwidth links between the processors. Nodes are assumed to have unique IDs in the range $\{1, 2, \dots, \text{poly}(n)\}$ and infinite computational power⁸ as the primary focus of this model is communication complexity. We denote by λ the number of bits used to represent each ID, i.e.,

⁸In the algorithms developed in this paper this strong assumption is not necessary as the number of internal computational steps at each node is proportional to the number messages received in all rounds.

$\lambda = O(\log n)$. Each node has limited topological knowledge; in particular, every node u only knows the IDs of each neighbor v and $w(u, v, G)$, the weight of their connecting edge. As in [58], we assume that edge weights are polynomially bounded, i.e., the largest edge weight of the graph is polynomial in the number of nodes. This is a typical assumption as it allows to encode the weight of an edge in one (or a constant number of) messages.

The distributed communication is performed in *rounds*. At the beginning of each round, all nodes wake up simultaneously and then each node u then sends an arbitrary message of $B = O(\log n)$ bits through each edge (u, v) , and the message will arrive at node v at the end of the round. For the algorithms presented in this paper, we consider the weaker *broadcast CONGEST model*, where in every individual round the message sent by each node is the same for all neighbors. The *running time* of a distributed algorithm is the worst-case number of rounds needed to finish a task. It is typically analyzed based on n (the number of nodes) and D (the network diameter) [61].

DEFINITION 2 (Single-Source Shortest Paths (SSSP)). *In the single-source shortest paths problem (SSSP), we are given a weighted network G and a source node s , i.e., each node knows (i) the IDs of its neighbors, (ii) the weight of its incident edges and (iii) whether it is the source s or not. We want to find the distance between s and every node v in G , denoted by $d(s, v, G)$, i.e., we want every node v to know the value of $d(s, v, G)$. In the α -approximate single-source shortest paths problem (SSSP) each node additionally knows the value $\alpha \geq 1$ and the goal is for every node v to know a distance estimate $\hat{d}(s, v)$ such that $d(s, v, G) \leq \hat{d}(s, v) \leq \alpha \cdot d(s, v, G)$.*

Recovering shortest paths. We note that although we define the problem to be computing the distances, we can easily recover the shortest paths in the sense that every node u knows its neighbor v that is in the shortest path between u and s . This is because our algorithm computes a distance estimate that satisfies the following property:

- (1) every node $u \neq s$ has a neighbor v such that $\hat{d}(s, v) + w(u, v, G) \leq \hat{d}(s, u)$,

where $\hat{d}(s, v)$ is the approximate distance between s and v . For any distance approximation $\hat{d}(s, \cdot)$ that satisfies Equation (1), we can recover the approximate shortest paths by assigning v as the intermediate neighbor of u in the approximate shortest path between u and s .

It can be easily checked throughout that the distance estimate that we compute satisfies Equation (1). This is simply because our algorithm always rounds an edge weight $w(u, v, G)$ up to some value $w'(u, v)$, and computes the approximate distances based on this rounded edge weight. For this reason, we can focus only on computing approximate distances in this paper.

2.3. Toolkit. In the following we review, in more detail, known results used for designing our algorithm. The first is a weight-rounding technique [47, 12, 73, 9, 57, 10, 58] for scaling down edge weights at the cost of approximation. Intuitively, we will use this technique to efficiently compute approximate shortest paths up to a fixed number of hops. As we will use this technique repeatedly, we give a proof in Appendix A for completeness

LEMMA 3 ([58]). *Let $h \geq 1$ and let G be a graph with positive integer edge weights in the range $\{1, \dots, W\}$. For every integer $0 \leq i \leq \lfloor \log(nW) \rfloor$, set $\rho_i = \frac{\epsilon 2^i}{h}$ and let G_i be the graph with the same nodes and edges as G and weight $w(u, v, G_i) = \lceil \frac{w(u, v, G)}{\rho_i} \rceil$*

for every edge (u, v) . Then for all pairs of nodes u and v and every $0 \leq i \leq \lfloor \log(nW) \rfloor$

$$(2) \quad \rho_i \cdot d(u, v, G_i) \geq d(u, v, G).$$

Moreover, if $2^i \leq d^h(u, v, G) \leq 2^{i+1}$, then

$$(3) \quad d(u, v, G_i) \leq (1 + 2/\epsilon)h \quad \text{and}$$

$$(4) \quad \rho_i \cdot d(u, v, G_i) \leq (1 + \epsilon) \cdot d^h(u, v, G),$$

where $d^h(u, v, G)$ is the h -hop distance between u and v in G .

An important subroutine in our algorithm is a procedure for solving the *source detection* problem [54] in which we want to find the σ nearest “sources” in a set S for every node u , given that u is at distance at most γ from them. Ties are broken by ID. The following definition is from [53]:

DEFINITION 4 ((S, γ, σ) -detection). *Consider a graph $G = (V, E)$, a set of “sources” $S \subseteq V$, and parameters $\gamma, \sigma \in \mathbb{N}$. For any node v let $\mathcal{L}(v, S, \gamma, \sigma, G)$ denote the proximity list resulting from ordering the set $\{(d(u, v, G), u) \mid u \in S \wedge d(u, v, G) \leq \gamma\}$ lexicographically in ascending order, i.e., where*

$$\begin{aligned} ((d(u, v, G), v) < (d(u', v, G), u')) &\iff \\ &(d(u, v, G) < d(u', v, G)) \vee (d(u, v, G) = d(u', v, G) \wedge u < u'), \end{aligned}$$

and restricting the resulting list to the first σ entries. The goal of the (S, γ, σ) -detection problem is to compute $\mathcal{L}(v, S, \gamma, \sigma, G)$ for every node $v \in V$. In the distributed setting we assume that, as part of the input, each node knows γ, σ , and whether it is in S or not and the goal is that every node $v \in V$ knows its list $\mathcal{L}(v, S, \gamma, \sigma, G)$.

Lenzen and Peleg designed a source detection algorithm for unweighted networks in the CONGEST model [54]. Their algorithm maintains, for each node, a tentative proximity list, where each entry is a pair consisting of a distance value and a source node in S . The list of every node v is initialized with the pair $(0, v)$, and in every round, each node sends, to all of its neighbors, the smallest entry in its list (according to lexicographic order) that it has not transmitted before. Upon receiving a pair (δ_s, s) , a node first checks if there already is some entry (δ'_s, s) with $\delta'_s \leq \delta_s + 1$ in its tentative proximity list and if not, it adds the pair $(\delta_s + 1, s)$ to its list (and marks it as not yet transmitted). Lenzen and Peleg show that after $\min(\gamma, D) + \min(\sigma, |S|)$ rounds, the first σ entries in the list maintained by every node v correspond to $\mathcal{L}(v, S, \gamma, \sigma, G)$. Holzer and Pinsker [41] had two observations about this algorithm. The first observation is that the guarantees of the algorithm directly carry over to the broadcast CONGEST model as in every round each node sends the same message to all of its neighbors. The second observation is that one can also run the algorithm on weighted networks (see also [53, proof of Theorem 3.3]), by replacing each edge of some weight L with an unweighted path of length L where all the nodes added for some weighted edge (u, v) are “simulated” by either u or v . Note that the “simulated” nodes are never sources. Furthermore, the tentative lists of the “simulated” nodes do not have to be maintained explicitly. The following modification of the algorithm for weighted graphs is functionally equivalent to the simulation approach: every time a node u wants to send some entry (δ_s, s) to some neighbor v via an edge of weight $w(u, v, G)$, it delays this message by $w(u, v) - 1$ rounds; upon reception, v first checks if there already is some entry (δ'_s, s) with $\delta'_s \leq \delta_s + w(u, v, G)$ in its tentative proximity list and if not, it adds the pair $(\delta_s + w(u, v, G), s)$ to its list (and marks it as not yet transmitted).

THEOREM 5 (Implicit in [54]). *In the broadcast CONGEST model, there is a deterministic algorithm for solving the (S, γ, σ) -detection problem in $\min(\gamma, WD) + \min(\sigma, |S|)$ rounds on weighted networks, where WD is the weighted diameter.*

We remark that in an earlier version of this paper we have, for example in the streaming model, used an additional source detection algorithm of Roditty, Thorup, and Zwick [64].⁹ Using a second algorithm is however not essential as the algorithm by Lenzen and Peleg [54] provides all necessary guarantees.

Another subproblem arising in our algorithm is the computation of *ruling sets*. The following definition was adapted from the recent survey of Barenboim and Elkin [4, Section 9.2].

DEFINITION 6 (Ruling Set). *For a (possibly weighted) graph $G = (V, E)$, a subset $U \subseteq V$ of nodes, and a pair of positive integers α and β , a set $T \subset U$ is an (α, β) -ruling set for U in G if*

1. *for every pair of distinct nodes $u, v \in T$, it holds that $d(u, v, G) \geq \alpha$, and*
2. *for every node $u \in U \setminus T$, there exists a “ruling” node $v \in T$, such that $d(u, v, G) \leq \beta$.*

The classic result of Goldberg et al. [36] shows that in the distributed setting, for any $c \geq 1$, we can compute a $(c, c\lambda)$ -ruling set deterministically in $O(c \log n)$ rounds, where λ is the number of bits used to represent each ID in the network. Since it was not explicitly stated that this algorithm works in the broadcast CONGEST model, we sketch an implementation of this algorithm in [Appendix C](#) (see [4, Chapter 9.2] and [61, Chapter 22] for more detailed algorithm and analysis).

THEOREM 7 (implicit in [36]). *In the broadcast CONGEST model, there is a deterministic algorithm that, for every $c \geq 1$, computes a $(c, c\lambda)$ -ruling set in $O(c \log n)$ rounds, where λ is the number of bits used to represent each ID in the network.*

3. Deterministic Hop Set Construction. In this section we present a deterministic algorithm for constructing an $(n^{o(1)}, o(1))$ -hop set (see [Definition 1](#)). In [Subsection 3.2](#) we first give an algorithm with a weaker guarantee that computes a set of edges F that reduces the number of hops between all pairs of nodes in the following way for some fixed $\Delta \geq 1$: if the shortest path has weight R , then using the edges of F , we can find a path with $\tilde{O}(R/\Delta)$ edges at the cost of a multiplicative error of $o(1)$ and an additive error of $n^{o(1)}\Delta$. Our algorithm obtains F by computing the *clusters* of the graph. We explain clusters and their computation in [Subsection 3.1](#). In [Subsection 3.3](#), we show how to repeatedly apply the first algorithm for different edge weight modifications to obtain a set of edges F providing the following stronger hop reduction for all pairs of nodes: if the shortest path has h hops, then, using the edges of F , we can find a path with $\tilde{O}(h/\Delta)$ hops at the cost of a multiplicative error of $o(1)$ and no additive error. Finally, in [Subsection 3.4](#), we obtain the hop set by repeatedly applying the hop reduction.

3.1. Deterministic Clusters. The basis of our hop set construction is a structure called *cluster* introduced by Thorup and Zwick [68] who used it, e.g., to construct distance oracles [68] and spanners [69] of small size.

⁹Roditty, Thorup, and Zwick [64] solve a variant of the source detection problem with $\gamma = \infty$ in their centralized algorithm for computing distance oracles and spanners deterministically. They essentially reduce the source detection problem to a sequence of single-source shortest paths computations on graphs with $O(n)$ additional nodes and edges. This reduction can be modified in a straightforward way to generalize their algorithm to arbitrary γ .

DEFINITION 8. Consider an integer p such that $2 \leq p \leq \log n$ and a hierarchy \mathcal{A} of sets of nodes $(A_i)_{0 \leq i \leq p}$ such that $A_0 = V$, $A_p = \emptyset$, and $A_0 \supseteq A_1 \supseteq \dots \supseteq A_p$. We say that a node v has priority i if $v \in A_i \setminus A_{i+1}$ (for $0 \leq i \leq p-1$). For every node $v \in V$ the cluster of v in G is defined as

$$C(v, \mathcal{A}, G) = \{u \in V \mid d(u, v, G) < d(u, A_{i+1}, G)\},$$

where i is the priority of v .

In the non-centralized models of computation considered in this paper, the straightforward way of computing clusters as defined above is not efficient enough for our purposes. We can however afford to compute the following restricted clusters. For any node v of priority i , let

DEFINITION 9. Consider p and \mathcal{A} as in Definition 8 and $R \geq 0$. For every node $v \in V$ the restricted cluster up to distance R of v in G is defined as

$$C(v, \mathcal{A}, G) = \{u \in V \mid d(u, v, G) < d(u, A_{i+1}, G)\},$$

where i is the priority of v .

3.1.1. Computing Priorities \mathcal{A} . The performance of our algorithms relies on the total size of the clusters, which in turn depends on how we compute nodes' priorities. If randomization is allowed, we can use the following algorithm due to Thorup and Zwick [68, 69]: set $A_0 = V$ and $A_p = \emptyset$, and for $1 \leq i \leq p-1$ obtain A_i by picking each node from A_{i-1} with probability $((\ln n)/n)^{1/p}$. It can be argued that for $\mathcal{A} = (A_i)_{0 \leq i \leq p}$ the size of all clusters, i.e., $\sum_{v \in V} |C(v, \mathcal{A}, G)|$, is $O(pn^{1+1/p})$ in expectation [68]. We now explain how to deterministically compute the priorities of nodes (given by a hierarchy of sets of nodes $\mathcal{A} = (A_i)_{0 \leq i \leq p}$) such that the total size of the resulting clusters is $\sum_{v \in V} |C(v, \mathcal{A}, G)| = O(pn^{1+1/p})$.

Thorup and Zwick [68] introduced the notion of *bunches* to analyze the sizes of clusters. For every node $u \in V$, we define the bunch and, for every $0 \leq i \leq p-1$, the i -bunch, both restricted to distance R , as follows:

$$B_i(u, \mathcal{A}, R, G) = \{v \in A_i \setminus A_{i+1} \mid d(u, v, G) < d(u, A_{i+1}, G) \text{ and } d(u, v, G) \leq R\}$$

$$B(u, \mathcal{A}, R, G) = \bigcup_{0 \leq i \leq p-1} B_i(u, \mathcal{A}, R, G).$$

The crucial insight is that $v \in B(u, \mathcal{A}, R, G)$ if and only if $u \in C(v, \mathcal{A}, R, G)$. Thus, it suffices to choose a hierarchy of sets A_i such that $|B_i(u, \mathcal{A}, R, G)| \leq O(n^{1/p})$ for every $u \in V$ and $0 \leq i \leq p-1$.

Our algorithm for deterministically computing this hierarchy of sets of nodes follows the main idea of Roditty, Thorup, and Zwick [64]. Its pseudocode is given in Procedure 1. As a subroutine this algorithm solves a weighted source detection problem, i.e., for suitable parameters q , A , and R , it computes for every node v the proximity list $\mathcal{L}(v, A, R, q, G)$ containing the $q = \tilde{O}(n^{1/p})$ nodes of A that are closest to v – up to distance R ; if there are fewer than q nodes of A in distance R to v , $\mathcal{L}(v, A, R, q, G)$ contains all of them. Our algorithm for constructing the hierarchy of sets $(A_i)_{0 \leq i \leq p}$ is as follows. We set $A_0 = V$ and $A_p = \emptyset$ and to construct the set A_{i+1} given the set A_i for $0 \leq i \leq p-2$ we first find for each node $v \in V$ the set $\mathcal{L}(v, A_i, R, q, G)$ using a source detection algorithm. Then we view the collection of sets $\{\mathcal{L}(v, A_i, R, q, G)\}_{v \in V}$ as an instance of the hitting set problem over the universe A_i , where we want to find

a set $A_{i+1} \subseteq A_i$ of minimum size such that each set $\mathcal{L}(v, A_i, R, q, G)$ contains at least one node of A_{i+1} , i.e., a hitting set. We let A_{i+1} be an approximate hitting set whose size is within a factor of $1 + \ln n$ of the optimum¹⁰ produced by the deterministic greedy heuristic (always adding the element “hitting” the largest number of “un-hit” sets) [43, 3]. Following [68], we explicitly set $A_p = \emptyset$ to avoid the introduction of special notation for clusters of the largest priority. In the following we prove the desired bound on the size of the bunches, which essentially requires us to argue that setting $q = \tilde{O}(n^{1/p})$ is sufficient.

Procedure 1: PRIORITIES(G, p, R)

Input: Weighted graph $G = (V, E)$ with positive integer edge weights, number of priorities $p \geq 2$, distance range $R \geq 1$

Output: Hierarchy of sets $\mathcal{A} = (A_i)_{0 \leq i \leq p}$

```

1  $q \leftarrow \lceil 2n^{1/p} \ln(3n)(1 + \ln n) \rceil$ 
2  $A_0 \leftarrow V$ 
3 for  $i = 0$  to  $p - 2$  do
4   Compute  $\mathcal{L}(v, A_i, R, q, G)$  for every node  $v \in V$  using a source detection
   algorithm
5    $\mathcal{C} \leftarrow \emptyset$ 
6   foreach  $v \in V$  do
7     if  $|\mathcal{L}(v, A_i, q, R, G)| = q$  then  $\mathcal{C} \leftarrow \mathcal{C} \cup \{\mathcal{L}(v, A_i, R, q, G)\}$ 
8   Compute a  $(1 + \ln n)$ -approximate minimum hitting set  $A_{i+1} \subseteq A_i$  using a
   greedy heuristic
9  $A_p \leftarrow \emptyset$ 
10 return  $\mathcal{A} := (A_i)_{0 \leq i \leq p}$ 

```

LEMMA 10 (Implicit in [64]). *Given a finite collection of sets $\mathcal{C} = \{S_1, \dots, S_k\}$ over a universe U and a parameter $x \geq 1$ such that $|S_i| \geq 2x \ln(3k)$ for all $1 \leq i \leq k$, there exists a hitting set $T \subseteq U$ of size at most $|T| \leq |U|/x$ such that $T \cap S_i \neq \emptyset$ for all $1 \leq i \leq k$.*

We give a proof of Lemma 10 in Appendix B for completeness.

LEMMA 11. *Given a weighted graph $G = (V, E)$ with positive integer edge weights and a parameters $p \geq 2$ and $R \geq 1$, Procedure 1 computes a hierarchy $\mathcal{A} = (A_i)_{0 \leq i \leq p}$ of sets of nodes such that*

$$\sum_{u \in V} |C(u, \mathcal{A}, R, G)| = \sum_{u \in V} |B(u, \mathcal{A}, R, G)| = O(pn^{1+1/p} \log n).$$

Proof. We first show by induction that $|A_i| \leq n^{1-i/p}$ for all $0 \leq i \leq p - 1$. If $i = 0$ the claim is trivially true because we set $A_0 = V$. We now assume the induction hypothesis $|A_i| \leq n^{1-i/p}$ and argue that $|A_{i+1}| \leq n^{1-(i+1)/p}$. Our algorithm computes a $(1 + \ln n)$ -approximate hitting set of the collection of sets \mathcal{C} containing each set $\mathcal{L}(v, A_i, q, R, G)$ of size q (i.e., not considering all such sets that have size strictly less than q). Let $k \leq n$ be the number of sets contained in \mathcal{C} . As each set in \mathcal{C} has size

¹⁰In principle, local computation is free in the models considered in this paper and we could thus compute a minimum hitting set exactly. However, we decided to present the algorithm in a way that avoids solving NP-complete problems by local computation.

$q = \lceil 2n^{1/p} \ln(3n)(1 + \ln n) \rceil \geq 2x \ln(3q)$ for $x = n^{1/p}(1 + \ln n)$, we know by [Lemma 10](#) that there is a hitting set A' for \mathcal{C} of size at most $|A_i|/x = |A_i|/(n^{1/p}(1 + \ln n)) \leq n^{1-i/p}/(1 + \ln n)$ and thus the minimum hitting set A_{i+1} computed by the greedy heuristic has size at most $(1 + \ln n)|A'| \leq n^{1-i/p}$. Note that each set $\mathcal{L}(v, A_i, q, R, G)$ might have been empty and in this case the algorithm would have computed $A_{i+1} = \emptyset$, the trivial hitting set.

We now show that for every node $u \in V$ and every $0 \leq i \leq p-1$, $B_i(u, \mathcal{A}, R, G) \leq q = O(n^{1/p} \log n)$, which immediately implies the desired bound on the total size of the bunches and clusters. We argue by a case distinction that $B_i(u, \mathcal{A}, R, G) \subseteq \mathcal{L}(u, A_i, R, q, G)$ and thus, by the definition of the set of the q closest nodes in A_i , $|B_i(v, \mathcal{A}, R, G)| \leq |\mathcal{L}(u, A_i, R, q, G)| \leq q$. If $|\mathcal{L}(u, A_i, R, q, G)| < q$, then clearly $\mathcal{L}(u, A_i, R, q, G) = \{v \in A_i \mid d(u, v, G) \leq R\} \supseteq B_i(v, \mathcal{A}, R, G)$. Otherwise we have $|\mathcal{L}(u, A_i, R, q, G)| = q$ and, as the algorithm computed a suitable hitting set, we have $B_i(u, \mathcal{A}, R, G) \subseteq \mathcal{L}(u, A_i, R, q, G)$. \square

As an alternative to the algorithm proposed above, the hitting sets can also be computed with the deterministic algorithm of Roditty, Thorup, and, Zwick [\[64\]](#) which produces so-called “early hitting sets”. For this algorithm we have to set $q = O(n^{1/p} \log n)$, and obtain slightly smaller clusters of total size $O(pn^{1+1/p})$. However, since the logarithmic factors are negligible for our purpose, we have decided to present the simpler algorithm above.

3.1.2. Computing Clusters. Given the priorities of the nodes, the clusters can be computed by finding, for every priority i and for every node v of priority i , the shortest paths up to nodes whose distance to v is more than (or equal to) their distance to nodes of priority more than i . In the pseudocode of [Procedure 2](#) we formulate this algorithm as a variant of weighted breadth-first search. We will not analyze the performance of this algorithm at this point since it depends on the models of computation that simulate it (see [Subsection 3.1.1](#) and [Section 5](#) for implementations of the algorithm and performance analyses).

We summarize our guarantees with the following theorem.

THEOREM 12. *Given a weighted graph $G = (V, E)$ with positive integer edge weights and parameters $p \geq 2$ and $R \geq 1$, [Procedure 2](#) computes a hierarchy of sets $\mathcal{A} = (A_i)_{0 \leq i \leq p}$, where $V = A_0 \subseteq A_1 \subseteq \dots \subseteq A_p = \emptyset$, such that $\sum_{v \in V} |C(v, \mathcal{A}, R, G)| = \tilde{O}(pn^{1+1/p})$. It also computes for every node v the set $C(v, \mathcal{A}, R, G)$ and for each node $w \in C(v, \mathcal{A}, R, G)$ the value of $d(v, w, G)$.*

3.2. Hop Reduction with Additive Error. Consider the following algorithm for computing a set of edges F . First, deterministically compute clusters with $p = \lfloor \sqrt{\log n / \log(9/\epsilon)} \rfloor$ priorities (determined by a hierarchy of sets $\mathcal{A} = (A_i)_{0 \leq i \leq p}$) up to distance $R = n^{1/p} \Delta$, where Δ is a parameter controlling both the extent of the hop reduction and the quality of the resulting approximation. Let F be the set containing an edge for every pair $(u, v) \in V^2$ such that $v \in C(u, \mathcal{A}, R, G)$ and set the weight of such an edge $(u, v) \in F$ to $w(u, v, F) = d(u, v, G)$, where the distance is returned by the algorithm for computing the clusters. [Procedure 3](#) presents the pseudocode of this algorithm.

LEMMA 13. *Let $F \subseteq V^2$ be the set of edges computed by [Procedure 3](#) for a weighted graph $G = (V, E)$ with positive integer edge weights and parameters $\Delta \geq 1$ and $0 < \epsilon \leq 1$. Then F has size $\tilde{O}(pn^{1+1/p})$, where $p = \lfloor \sqrt{(\log n) / (\log(9/\epsilon))} \rfloor$, and in the*

Procedure 2: CLUSTERS(G, p, R)

Input: Weighted graph $G = (V, E)$ with positive integer edge weights, number of priorities $p \geq 2$, distance range $R \geq 1$

Output: Clusters of G as specified in [Theorem 12](#)

```

1  $\mathcal{A} = (A_i)_{0 \leq i \leq p} \leftarrow \text{PRIORITIES}(G, p, R)$ 
2 For each  $1 \leq i \leq p - 1$  and every node  $v \in V$  compute  $d(v, A_i, R, G)$ 
3 foreach  $u \in V$  do // Compute cluster of every node
  // Initialization
4 Let  $i$  be the priority of  $u$ , i.e.,  $u \in A_i \setminus A_{i+1}$ 
5 foreach  $v \in V$  do  $\delta(u, v) \leftarrow \infty$ 
6  $\delta(u, u) \leftarrow 0$ 
7  $C(u) \leftarrow \emptyset$ 
  // Iteratively add nodes to cluster
8 for  $L = 0$  to  $R$  do
9   foreach node  $v$  with  $\delta(u, v) = L$  do
10     // Check if  $v$  joins cluster of  $u$  at current level
11     if  $\delta(u, v) < d(v, A_{i+1}, R, G)$  then
12        $C(u) \leftarrow C(u) \cup \{v\}$ 
13       foreach  $(v, w) \in E$  do // Update neighbors of  $v$ 
14          $\delta'(u, w) \leftarrow (w(v, w, G) + \delta(u, v))$ 
15         if  $\delta'(u, w) < \delta(u, w)$  then  $\delta(u, w) \leftarrow \delta'(u, w)$ 
15 return  $(C(v), \delta(v, \cdot))_{v \in V}$ 

```

Procedure 3: HOPREDUCTIONADDITIVEERROR(G, Δ, ϵ)

Input: Graph $G = (V, E)$ with positive integer edge weights, $\Delta \geq 1$, $0 < \epsilon \leq 1$

Output: Hop-reducing set of edges $F \subseteq V^2$ as specified in [Lemma 13](#)

```

1  $p \leftarrow \left\lfloor \sqrt{\frac{\log n}{\log(9/\epsilon)}} \right\rfloor$ 
2  $R \leftarrow n^{1/p} \Delta$ 
3  $F \leftarrow \emptyset$ 
4  $(C(v, \mathcal{A}, R, G), \delta(v, \cdot))_{v \in V} \leftarrow \text{CLUSTERS}(G, p, R)$ 
5 foreach  $u \in V$  do
6   foreach  $v \in C(u)$  do
7      $F \leftarrow F \cup \{(u, v)\}$ 
8      $w(u, v, F) = \delta(u, v)$ 
9 return  $F$ 

```

graph $H = G \cup F$, for every pair of nodes u and v , we have

$$d^{(p+1)\lceil d(u,v,G)/\Delta \rceil}(u, v, G) \leq (1 + \epsilon)d(u, v, G) + \epsilon n^{1/p} \Delta / (p + 2),$$

i.e., there is a path π' in H of weight $w(\pi', H) \leq (1 + \epsilon)d(u, v, G) + \epsilon n^{1/p} \Delta / (p + 2)$ consisting of $|\pi'| \leq (p + 1)\lceil d(u, v, G)/\Delta \rceil$ edges.

We devote the rest of this section to proving [Lemma 13](#). The bound on the size of F

immediately follows from [Theorem 12](#). We analyze the hop-reducing properties of F by showing the following. Let π be a shortest path from u to v in G . Then there is a node w on π and a path π' from u to w in $H = G \cup F$ with the following properties:

- (1) The distance from u to w in G is at least Δ .
- (2) The path π' consists of at most p edges of F and at most one edge of G .
- (3) The ratio between the weight of π' in H and the distance from u to w in G is at most $(1 + \epsilon)$ if $w \neq v$ and if $w = v$ then the weight of π' in H is at most β (for some β that we set later).

When we go from u to w using the path π' instead of the subpath of π we are using a “shortcut” of at most $p + 1$ hops that brings us closer to v by a distance of at least Δ at the cost of some approximation. Conditions (1) and (2) guarantee that by repeatedly applying this shortcutting we can find a path π'' from u to v that has at most $(p + 1)\lceil d(u, v, G)/\Delta \rceil$ hops (as we replace subpaths of π with weight at least Δ by paths with at most $p + 1$ hops). Condition (3) guarantees that the multiplicative error introduced by using the shortcut is at most $1 + \epsilon$, except possibly for the last time such a shortcut is used, where we allow an additive error of β . We show in [Lemma 14](#) that we can guarantee a value of β that is bounded by $\epsilon n^{1/p}/(p + 2)$. This type of analysis has been used before by Thorup and Zwick [69] to obtain a spanner for unweighted graphs defined from the partial shortest path trees of the clusters, but without considering the hop-reduction aspect. Bernstein [9] also used a similar analysis to obtain a hop set for weighted graphs using clusters with full distance range. We previously used this type of analysis to obtain a randomized hop set which is not based on clusters, but on a similar notion [39].

To carry out the analysis as explained above we define a value r_i for every $0 \leq i \leq p - 1$ as follows:

$$\begin{aligned} r_0 &= \Delta \\ r_i &= \frac{(4 + 2\epsilon) \sum_{0 \leq j \leq i-1} r_j}{\epsilon}. \end{aligned}$$

The intuition is that a node u of priority i tries to take an edge of F to shortcut the way to v by at least r_i . If this fails it will find an edge in F going to a node v' of higher priority. Thus, to fulfill Condition (3), v' has to try and shortcut even more “aggressively”. Consequently, the values of r_i grow exponentially with the priority i .

We have chosen the range of the clusters large enough such that nodes of the highest priority always find the desired shortcut edge in F . We will show that the additive error incurred by this strategy is at most

$$\beta = \sum_{0 \leq i \leq p-1} 2r_i.$$

This value can in turn be bounded as follows.

LEMMA 14. $\beta \leq \epsilon n^{1/p} \Delta / (p + 2)$.

Proof. We first show that, for all $0 \leq i \leq p - 1$, $\sum_{0 \leq j \leq i} r_j \leq 7^i \Delta / \epsilon^i$. The proof is by induction on i . For $i = 0$ we have $r_0 = \Delta = 7^0 \Delta$ by the definition of r_0 and for $i \geq 1$ we use the inequalities $\epsilon > 0$ and $\epsilon \leq 1$ and the induction hypothesis as follows:

$$\sum_{0 \leq j \leq i} r_j = \sum_{0 \leq j \leq i-1} r_j + r_i = \sum_{0 \leq j \leq i-1} r_j + \frac{(4 + 2\epsilon) \sum_{0 \leq j \leq i-1} r_j}{\epsilon}$$

$$\leq \frac{\sum_{0 \leq j \leq i-1} r_j}{\epsilon} + \frac{6 \sum_{0 \leq j \leq i-1} r_j}{\epsilon} \leq \frac{7 \sum_{0 \leq j \leq i-1} r_j}{\epsilon} \leq \frac{7 \cdot 7^{i-1} \Delta}{\epsilon \cdot \epsilon^{i-1}} \leq \frac{7^i \Delta}{\epsilon^i}.$$

Using this inequality and the fact that $(2p+7)7^{p-1} \leq 9^p$ for all $p \geq 0$ we get

$$\frac{(p+2)\beta}{\epsilon} = \frac{(p+2) \sum_{0 \leq j \leq p-1} 2r_j}{\epsilon} \leq \frac{(2p+4)7^{p-1} \Delta}{\epsilon^p} \leq \frac{9^p \Delta}{\epsilon^p} \leq n^{1/p} \Delta$$

The last inequality holds as by our choice of $p = \lfloor \sqrt{\log n / \log(9/\epsilon)} \rfloor$:

$$\begin{aligned} (9/\epsilon)^p &= 2^{p \cdot \log(9/\epsilon)} \leq 2^{\frac{\sqrt{\log n}}{\sqrt{\log(9/\epsilon)}} \cdot \log(9/\epsilon)} = 2^{\sqrt{\log n} \cdot \sqrt{\log(9/\epsilon)}} \\ &= 2^{\frac{\sqrt{\log(9/\epsilon)}}{\sqrt{\log n}} \cdot \log n} = 2^{\log n \cdot (1/p)} = n^{1/p}. \quad \square \end{aligned}$$

In the following we fix some values of ϵ and Δ and let F denote the set of edges computed by Procedure 3. We now show that F has a certain structural property before we carry out the hop-reduction proof.

LEMMA 15. *Let u and v be nodes such that $u \in A_i \setminus A_{i+1}$ (i.e., u has priority i) and $d(u, v, G) \leq r_i$. Either*

- (1) F contains an edge (u, v) of weight $w(u, v, F) = d(u, v, G)$ or
- (2) F contains an edge (u, v') to a node $v' \in A_{i+1}$ of priority $j \geq i+1$ of weight $w(u, v', F) \leq 2r_i$.

Proof. Consider first the case $v \in C(u, \mathcal{A}, R, G)$. Then F contains the edge (u, v) of weight $w(u, v, F) = d(u, v, G)$.

Consider now the case $v \notin C(u, \mathcal{A}, R, G)$. Note that by the definition of β we have $r_i \leq \beta/2 < \beta$ and by Lemma 14 we have $\beta \leq n^{1/p} \Delta$ (where p is the parameter controlling the construction of clusters). As the algorithm sets $R = n^{1/p} \Delta$ we have $r_i \leq R$ by Lemma 14 and thus $d(u, v, G) \leq R$. From the definition of $C(u, \mathcal{A}, R, G)$ it now follows that $d(v, u, G) \geq d(v, A_{i+1}, G)$. Thus there exists some node $v_1 \in A_{i+1}$ of priority $p_1 \geq i+1$ such that $d(v, v_1, G) \leq d(u, v, G)$. By the triangle inequality we get $d(u, v_1, G) \leq d(u, v, G) + d(v, v_1, G) \leq 2d(u, v, G) \leq 2r_i \leq R$. If $u \in C(v_1, \mathcal{A}, R, G)$ then we are done as F contains the edge (u, v_1) of weight $w(u, v_1, F) = d(u, v_1, G) \leq 2r_i$. Otherwise it follows from the definition of $C(v_1, \mathcal{A}, R, G)$ that there is some node $v_2 \in A_{p_1+1}$ of priority $p_2 \geq p_1+1 \geq i+1$ such that $d(u, v_2, G) \leq d(u, v_1, G) \leq 2r_i \leq R$. By repeating the argument above we want to find some node $v_j \in A_{i+1}$ of priority $p_j \geq i+1$ such that $d(u, v_j, G) \leq d(u, v, G) \leq 2r_i$. As for every node $v' \in A_{p-1}$ of priority $p-1$, $C(v', \mathcal{A}, R, G)$ contains all nodes that are at distance at most R from v' in G this repeated argument stops eventually and we find such a node. \square

To finish the proof of Lemma 13 we show in the next lemma that F has the properties we demanded, i.e., in the shortcut graph H which consists of G and the additional edges of F , we can approximate shortest paths using a reduced number of hops.

LEMMA 16. *For every pair of nodes $u, v \in V$ such that $d(u, v, G) < \infty$, the graph $H = G \cup F$ contains a path π' from u to v of weight $w(\pi', H) \leq (1+\epsilon)d(u, v, G) + \beta$ consisting of $|\pi'| \leq (p+1)\lceil d(u, v, G)/\Delta \rceil$ edges.*

Proof. The proof is by induction on the distance from u to v in G . The claim is trivially true for the base case $d(u, v, G) = 0$ in which $u = v$. Thus, we only need to consider the induction step in which $d(u, v, G) \geq 1$.

Let π denote the shortest path from u to v in G . We now define a sequence of nodes u_0, u_1, \dots, u_l where $l \leq p-1$. For every $0 \leq j \leq l$, we denote by p_j the priority

of u_j . We set $u_0 = u$ and, given u_j , we define u_{j+1} as follows. Let x be the node on π closest to v that is at distance at most r_{p_j} from u_j in G (this node might be v itself). If F contains the edge (u_j, x) we stop (and set $l = j$). Otherwise we know by [Lemma 15](#) that F (and therefore also H) contains an edge (u_j, u') to a node u' of priority at least $p_j + 1$. In that case we set $u_{j+1} = u'$. We know further by [Lemma 15](#) that $d(u_j, u_{j+1}, G) \leq 2r_{p_{j+1}-1}$. Note the definition of this sequence u_0, u_1, \dots, u_l ensures that $l \leq p - 1$ as the priority strictly increases with each node u_j . Having defined the sequence, we denote by x the node on π closest to v that is at distance at most r_{p_l} from u_l in G (again, this node might be v itself). The definition of u_l guarantees that F (and thus H) contains the edge (u_l, x) . [Figure 1](#) illustrates the definition of this sequence.




Fig. 1: Schematic illustration of the definition of the sequence of nodes u_0, u_1, \dots, u_l, x . The bottom line represents the shortest path π from u to v . The thick, blue edges are the edges from F used to shorten the distance to v . The dashed, blue edges are not contained in F and imply the existence of edges to nodes of increasing priority. The node y is the successor of x on π and the thick, red edge (x, y) from G is the last edge on the shortcut path from u_0 to y . The dotted lines indicate repetitions that are omitted in the picture.

Consider first the case that $x = v$. Let π' denote the path $\langle u_0, \dots, u_l, x \rangle$. This path has at most p hops and since $d(u, v, G) \geq 1$ we trivially have $p \leq (p + 1)\lceil d(u, v, G)/\Delta \rceil$. Furthermore we can bound the weight of π' as follows:

$$\begin{aligned}
w(\pi', H) &= \sum_{0 \leq j \leq l-1} w(u_j, u_{j+1}, H) + w(u_l, x, H) \\
&\leq \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) + d(u_l, x, G) \\
&\leq \sum_{0 \leq j \leq l-1} 2r_{p_{j+1}-1} + r_{p_l} \\
&\leq \sum_{0 \leq j \leq p-1} 2r_j \\
&= \beta \leq (1 + \epsilon)d(u, v, G) + \beta.
\end{aligned}$$

Consider now the case $x \neq v$. Let y be the neighbor of x on π (that in G is closer to v than x is). We will define the path π' from u to v as the concatenation of two

paths π_1 and π_2 . Let π_1 be the path $\langle u_0, \dots, u_l, x, y \rangle$. We will define the path π_2 from y to v later on. Note that π_1 consists of $|p_1| \leq p + 1$ hops. We will now show that

$$(5) \quad w(\pi_1, H) \leq (1 + \epsilon)d(u, y, G).$$

In order to get this bound we will need some auxiliary inequalities. By [Lemma 15](#) we have, for all $0 \leq j \leq l - 1$,

$$(6) \quad d(u_j, u_{j+1}, G) \leq 2r_{p_{j+1}-1}$$

and by the definition of r_{p_l} we have

$$(7) \quad \epsilon r_{p_l} = (4 + 2\epsilon) \sum_{0 \leq j \leq p_l-1} r_j.$$

Remember that x is the node on the path π closest to v that is at distance at most r_{p_l} from u_l in G . Since the neighbor y of x is closer to v than x is, this definition of x guarantees that $d(u_l, y, G) > r_{p_j}$. As $d(u_l, y, G) \leq d(u_l, x, G) + d(x, y, G)$ by the triangle inequality, we have

$$(8) \quad d(u_l, x, G) + d(x, y, G) > r_{p_j}.$$

By the triangle inequality we also have

$$d(u_l, x, G) \leq \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) + d(u, x, G)$$

and thus

$$(9) \quad d(u_l, x, G) - \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \leq d(u, x, G).$$

We now obtain Inequality (5) as follows:

$$\begin{aligned} w(\pi_1, H) &= \sum_{0 \leq j \leq l-1} w(u_j, u_{j+1}, H) + w(u_l, x, H) + w(x, y, H) \\ &= \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) + d(u_l, x, G) + d(x, y, G) \\ &= (2 + \epsilon) \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) + d(u_l, x, G) + d(x, y, G) \\ &\quad - (1 + \epsilon) \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \\ &\stackrel{(6)}{\leq} (2 + \epsilon) \sum_{0 \leq j \leq l-1} 2r_{p_{j+1}-1} + d(u_l, x, G) + d(x, y, G) \\ &\quad - (1 + \epsilon) \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \\ &\leq (2 + \epsilon) \sum_{0 \leq j \leq p_l-1} 2r_j + d(u_l, x, G) + d(x, y, G) \\ &\quad - (1 + \epsilon) \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \end{aligned}$$

$$\begin{aligned}
&\stackrel{(7)}{=} \epsilon r_{p_l} + d(u_l, x, G) + d(x, y, G) - (1 + \epsilon) \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \\
&\stackrel{(8)}{<} \epsilon (d(u_l, x, G) + d(x, y, G)) + d(u_l, x, G) + d(x, y, G) \\
&\quad - (1 + \epsilon) \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \\
&= (1 + \epsilon) \left(d(u_l, x, G) - \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \right) + (1 + \epsilon) d(x, y, G) \\
&\stackrel{(9)}{\leq} (1 + \epsilon) d(u, x, G) + (1 + \epsilon) d(x, y, G) \\
&\leq (1 + \epsilon) (d(u, x, G) + d(x, y, G)) = (1 + \epsilon) d(u, y, G)
\end{aligned}$$

Note that $d(y, v, G) < d(u, v, G)$. Therefore we may apply the induction hypothesis on y and get that the graph H contains a path π_2 from y to v of weight $w(\pi_2, H) \leq (1 + \epsilon)d(y, v, G) + \beta$ that has $|\pi_2| \leq (p + 1)\lceil d(y, v, G)/\Delta \rceil$ hops. Let π' denote the concatenation of π_1 and π_2 . Then π' is a path from u to v in H of weight

$$\begin{aligned}
w(\pi', H) &= w(\pi_1, H) + w(\pi_2, H) \\
&\leq (1 + \epsilon)d(u, y, G) + (1 + \epsilon)d(y, v, G) + \beta \\
&= (1 + \epsilon)(d(u, y, G) + d(y, v, G)) + \beta \\
&= (1 + \epsilon)d(u, v, G) + \beta.
\end{aligned}$$

It remains to bound the number of hops of π' . To get the desired bound we first show that $d(u, y, G) \geq \Delta$. By the triangle inequality we have

$$d(u_l, y, G) \leq \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) + d(u, y, G).$$

As argued above, we have $d(u_l, y, G) > r_{p_j}$ and

$$\sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \leq \sum_{0 \leq j \leq p_l-1} 2r_j.$$

By the definition of r_{p_l} we therefore get:

$$\begin{aligned}
d(u, y, G) &\geq d(u_l, y, G) - \sum_{0 \leq j \leq l-1} d(u_j, u_{j+1}, G) \\
&\geq r_{p_l} - \sum_{0 \leq j \leq p_l-1} 2r_j = (4/\epsilon) \sum_{0 \leq j \leq p_l-1} r_j \geq r_0 = \Delta.
\end{aligned}$$

Now that we know that $d(u, y, G) \geq \Delta$, or equivalently $d(u, y, G)/\Delta \geq 1$, we get the following for counting the number of hops of π' by adding the number of hops of π_1 to the number of hops of π_2 :

$$\begin{aligned}
|\pi'| &= |\pi_1| + |\pi_2| \leq p + 1 + (p + 1)\lceil d(y, v, G)/\Delta \rceil \\
&= (p + 1)(1 + \lceil d(y, v, G)/\Delta \rceil) \\
&= (p + 1)\lceil 1 + d(y, v, G)/\Delta \rceil \\
&\leq (p + 1)\lceil d(u, y, G)/\Delta + d(y, v, G)/\Delta \rceil \\
&= (p + 1)\lceil (d(u, y, G) + d(y, v, G))/\Delta \rceil \\
&= (p + 1)\lceil d(u, v, G)/\Delta \rceil.
\end{aligned}$$

Thus, π' has the desired number of edges. \square

3.3. Hop Reduction without Additive Error. Consider a shortest path π from u to v with h hops and weight $R \geq \Delta$. With the hop reduction of Procedure 3 we can compute a set of edges F such that in $G \cup F$ we find a path from u to v with $\tilde{O}(R/\Delta)$ hops of weight approximately R (where we incur an additive error of roughly $\Delta n^{o(1)}$). We now use the weight-rounding technique of Lemma 3 and repeatedly apply this algorithm to obtain a set of edges F such that in $G \cup F$ there is a path from u to v with $O(h/\Delta)$ hops and weight approximately R . As in general R can only be upper-bounded by nW (where W is the maximum edge weight of G) and h can be upper-bounded by n , the second type of hop reduction seems more desirable. Additionally, if h is sufficiently larger than Δ , then the additive error inherent in the hop reduction of Procedure 3 can be counted as a small multiplicative error.¹¹

The second hop reduction algorithm roughly works as follows. For every possible distance range of the form $2^j \dots 2^{j+1}$ we scale down the edge weights of G by a certain factor and run the algorithm of Procedure 3 on the modified graph \hat{G}_j to compute a set of edges \hat{F}_j . We then simply return the union of all these edge sets (with the weights scaled back to normal again). Procedure 4 shows the pseudocode of this algorithm.

Procedure 4: HOPREDUCTION($G, \Delta, h, \epsilon, W$)

Input: Weighted graph $G = (V, E)$ with positive integer edge weights in $\{1, \dots, W\}$, $\Delta \geq 1$, $h \geq 1$, $0 < \epsilon \leq 1$

Output: Hop-reducing set of edges $F \subseteq V^2$ as specified in Lemma 17

```

1  $\epsilon' \leftarrow \frac{\epsilon}{6}$ 
2  $\Delta' \leftarrow \frac{3\Delta}{\epsilon'}$ 
3  $F \leftarrow \emptyset$ 
4 for  $j = 0$  to  $\lfloor \log(nW) \rfloor$  do
5    $\hat{G}_j \leftarrow (V, E)$ 
6    $\rho_j \leftarrow \frac{\epsilon' 2^j}{h}$ 
7   foreach  $(u, v) \in E$  do  $w(u, v, \hat{G}_j) \leftarrow \left\lceil \frac{w(u, v, G)}{\rho_j} \right\rceil$ 
8    $\hat{F}_j \leftarrow \text{HOPREDUCTIONADDITIVEERROR}(\hat{G}_j, \Delta', \epsilon')$ 
9   foreach  $(u, v) \in \hat{F}_j$  do
10     $F \leftarrow F \cup \{(u, v)\}$ 
11     $w(u, v, F) \leftarrow \min(w(u, v, \hat{F}_j) \cdot \rho_j, w(u, v, F))$ 
12 return  $F$ 

```

LEMMA 17. Let $F \subseteq V^2$ be the set of edges computed by Procedure 4 for a weighted graph $G = (V, E)$ with positive integer edge weights in $\{1, \dots, W\}$ and parameters $\Delta \geq 1$, $h \geq 1$, and $0 < \epsilon \leq 1$. Then F has size $\tilde{O}(pn^{1+1/p} \log nW)$, where $p = \lfloor \sqrt{(\log n)/(\log(54/\epsilon))} \rfloor$, and if $h \geq n^{1/p} \Delta / (p + 2)$, then in the graph $H = G \cup F$ we have, for every pair of nodes u and v ,

$$d^{(p+2)h/\Delta}(u, v, H) \leq (1 + \epsilon) d^h(u, v, G).$$

¹¹Note that for smaller values of h , π itself has a small enough number of hops and thus there is no need to find a path in $G \cup F$ with a small number of hops and weight approximately R .

Proof. Let u and v be a pair of nodes and set $j = \lfloor \log d^h(u, v, G) \rfloor$, i.e., $2^j \leq d^h(u, v, G) \leq 2^{j+1}$. Let π be a shortest $\leq h$ hop path in G , i.e., π has weight $w(\pi, G) = d^h(u, v, G)$ and π consists of $|\pi| \leq h$ hops. The algorithm sets $\epsilon' = \epsilon/6$ and uses a graph \hat{G}_j which has the same nodes and edges as G , but in which every edge weight is first scaled down by a factor of $\rho_j = \epsilon'2^j/h$ and then rounded up to the next integer. By [Lemma 3](#) we have $d(u, v, \hat{G}_j) \cdot \rho_j \leq (1 + \epsilon')d^h(u, v, G)$ and $d(u, v, \hat{G}_j) \leq (1 + 2/\epsilon')h \leq 3h/\epsilon'$.

Consider the set of edges \hat{F}_j computed in [Procedure 4](#) (such a set does indeed exist because $d^h(u, v, G) \leq nW$). By [Lemma 13](#), there is a path π' in $\hat{H}_j = \hat{G}_j \cup \hat{F}_j$ of weight at most $(1 + \epsilon')d(u, v, \hat{G}_j) + \epsilon'n^{1/p}\Delta'/(p+2)$ and with at most $|\pi'| \leq (p+1)\lceil d(u, v, \hat{G}_j)/\Delta' \rceil$ hops. Since we have $d(u, v, \hat{G}_j) \leq 3h/\epsilon'$ and the algorithm sets $\Delta' = 3\Delta/\epsilon'$, we get

$$\begin{aligned} |\pi'| &\leq (p+1) \cdot \left\lceil \frac{d(u, v, \hat{G}_j)}{\Delta'} \right\rceil \\ &\leq (p+1) \cdot \left\lceil \frac{3h}{\epsilon'\Delta'} \right\rceil \\ &= (p+1) \cdot \left\lceil \frac{h}{\Delta} \right\rceil \\ &\leq (p+1) \left(\frac{h}{\Delta} + 1 \right) \\ &= \frac{(p+1)h}{\Delta} + (p+1) \\ &\leq \frac{(p+1)h}{\Delta} + 4^p \leq \frac{(p+1)h}{\Delta} + n^{1/p} \leq \frac{(p+1)h}{\Delta} + \frac{h}{\Delta} = \frac{(p+2)h}{\Delta}. \end{aligned}$$

The algorithm ‘‘scales back’’ the edge weights of \hat{F}_j when adding them to F and thus $w(u, v, F) \leq w(u, v, \hat{F}_j) \cdot \rho_j$. We now argue that $d^{(p+1)\lceil h/\Delta \rceil}(u, v, H) \leq (1 + \epsilon')d^h(u, v, G)$ by bounding the weight of π' in $H = G \cup F$. For every edge (u, v) of π' we have $w(u, v, H) \leq w(u, v, F) \leq w(u, v, \hat{F}_j) \cdot \rho_j$ if $(u, v) \in \hat{F}_j$ and $w(u, v, H) \leq w(u, v, G) \leq w(u, v, \hat{G}_j) \cdot \rho_j$ otherwise. Thus, $w(\pi', H) \leq w(\pi', \hat{H}_j) \cdot \rho_j$ and together with the assumption $h \geq n^{1/p}\Delta/(p+2)$ we get

$$\begin{aligned} d^{(p+2)h/\Delta}(u, v, H) &\leq w(\pi', H) \leq w(\pi', \hat{H}_j) \cdot \rho_j \\ &\leq \left((1 + \epsilon')d(u, v, \hat{G}_j) + \frac{\epsilon'n^{1/p}\Delta'}{p+2} \right) \cdot \rho_j \\ &= (1 + \epsilon')d(u, v, \hat{G}_j) \cdot \rho_j + \frac{\epsilon'n^{1/p}\Delta'\rho_j}{p+2} \\ &= (1 + \epsilon')d(u, v, \hat{G}_j) \cdot \rho_j + \frac{3\epsilon'2^j n^{1/p}\Delta}{h(p+2)} \\ &\leq (1 + \epsilon')d(u, v, \hat{G}_j) \cdot \rho_j + 3\epsilon'2^j \\ &\leq (1 + \epsilon')^2 d^h(u, v, G) + 3\epsilon' d^h(u, v, G) \\ &\leq (1 + 6\epsilon')d^h(u, v, G) \\ &= (1 + \epsilon)d^h(u, v, G). \end{aligned} \quad \square$$

3.4. Computing the Hop Set. We finally explain how to repeatedly use the hop reduction of Procedure 4 to obtain an $(n^{o(1)}, o(1))$ -hop set. Procedure 4 computes a set of edges F that reduces the number of hops needed to approximate the distance between any pair of nodes by a factor of $1/\Delta$ (where Δ is a parameter). Intuitively we would now like to use a large value of Δ to compute a hop set. However, we want to avoid large values of Δ for two reasons. The first reason is that F only reduces the number of hops if the shortest path has $h \geq \Delta n^{o(1)}$ hops. Thus, for shortest paths that already have $h < \Delta$ hops the hop reduction is not effective. The second reason is efficiency. The algorithm requires us to compute clusters for distances up to $\Delta n^{o(1)}$ and, in the models of computation we consider later on, we do not know how to do this fast enough for our purposes.

We therefore use the following iterative approach in which we repeatedly apply the hop reduction of Procedure 4 with $\Delta = (p+2)n^{1/p} = n^{o(1)}$. We first compute a set of edges F_1 that reduces the number of hops in G by a factor of $1/\Delta$. We then add all these edges to G and consider the graph $H_1 = G \cup F_1$. We apply the algorithm again on H_1 to compute a set of edges F_2 that reduces the number of hops in H_1 by a factor of $1/\Delta$. Now observe that the set of edges $F_1 \cup F_2$ reduces the number of hops in G by a factor of $1/\Delta^2$. We show that by repeating this process $p = \Theta(\sqrt{\log n / \log(\sqrt{\log n / \epsilon})})$ times we can compute a set F that reduces the number of hops to $n^{1/p}$. Procedure 5 shows the pseudocode of this algorithm.

Procedure 5: HOPSET(G, ϵ, W)

Input: Weighted graph $G = (V, E)$ with positive integer edge weights in $\{1, \dots, W\}$, $0 < \epsilon \leq 1$

Output: $(n^{1/p}, \epsilon)$ -hop set $F \subseteq V^2$ as specified in Theorem 18

```

1  $\epsilon' \leftarrow \frac{\epsilon}{\sqrt{\log n}}$ 
2  $W' \leftarrow (1 + \epsilon)nW$ 
3  $p \leftarrow \lfloor \sqrt{\frac{\log n}{\log(54/\epsilon')}} \rfloor$ 
4  $\Delta \leftarrow (p+2)n^{1/p}$ 
5  $F \leftarrow \emptyset$ 
6  $H_0 \leftarrow G$ 
7 for  $i = 0$  to  $p - 1$  do
8    $h_i \leftarrow n^{1-i/p}$ 
9    $F_{i+1} \leftarrow \text{HOPREDUCTION}(H_i, \Delta, h_i, \epsilon', W')$ 
10   $F \leftarrow F \cup F_{i+1}$ 
11   $H_{i+1} \leftarrow H_i \cup F_{i+1}$ 
12 return  $F$ 

```

THEOREM 18. *Let $F \subseteq V^2$ be the set of edges computed by Procedure 5 for a weighted graph $G = (V, E)$ with positive integer edge weights in $\{1, \dots, W\}$ and a parameter $0 < \epsilon \leq 1$. Then F is an $(n^{1/p}, \epsilon)$ -hop set of size $\tilde{O}(p^2 n^{1+1/p} \log nW)$, where $p = \lfloor \sqrt{(\log n) / (\log(54\sqrt{\log n / \epsilon}))} \rfloor$.*

Proof. The algorithm sets $\epsilon' = \epsilon / \sqrt{\log n}$ and $p = \lfloor \sqrt{(\log n) / (\log(54/\epsilon'))} \rfloor$ and uses a parameter $h_i = n^{1-i/p}$ for each graph H_i . For every $0 \leq i \leq p-2$ we set $h_i = n^{1-i/p} \geq n^{2/p} = n^{1/p} \Delta / (p+2)$ and thus, by Lemma 17, for every pair of nodes

u and v we have

$$\begin{aligned} d^{h_{i+1}}(u, v, H_{i+1}) &= d^{h_i/n^{1/p}}(u, v, H_{i+1}) = d^{(p+2)h_i/\Delta}(u, v, H_{i+1}) \\ &\leq (1 + \epsilon')d^{h_i}(u, v, H_i). \end{aligned}$$

By iterating this argument we get

$$d^{h_i}(u, v, H_i) \leq (1 + \epsilon')^i d^{h_0}(u, v, H_0) = (1 + \epsilon')^i d^n(u, v, G) = (1 + \epsilon')^i d(u, v, G)$$

for every $1 \leq i \leq p-1$ and now in particular for $i = p-1$ we have

$$d^{n^{1/p}}(u, v, H_{p-1}) = d^{h_{p-1}}(u, v, H_{p-1}) \leq (1 + \epsilon')^{(p-1)} d(u, v, G).$$

Finally, since $p-1 \leq \sqrt{\log n}$ we have, by Bernoulli's inequality ($(1+x)^r \leq 1+rx$ for reals $0 \leq r \leq 1$ and $x \geq -1$),

$$(1 + \epsilon')^{p-1} = \left(1 + \frac{\epsilon}{\sqrt{\log n}}\right)^{p-1} \leq \left(1 + \frac{\epsilon}{\sqrt{\log n}}\right)^{\sqrt{\log n}} \leq 1 + \epsilon.$$

As $H_{p-1} = G \cup F$ it follows that $d^{n^{1/p}}(u, v, G \cup F) \leq (1 + \epsilon)d(u, v, G)$ and thus $F = \bigcup_{1 \leq i \leq p-1} F_i$ is an $(n^{1/p}, \epsilon)$ -hop set. \square

The main computational cost for constructing the hop set comes from computing the clusters in Procedure 3, which is used as a subroutine repeatedly. Observe that in total it will perform $O(p \log n W)$ calls to Procedure 2 to compute clusters, each with $p = \Theta(\sqrt{(\log n)/(\log(\sqrt{\log n}/\epsilon))})$ priorities and distance range $R = O(p\sqrt{\log n}n^{2/p}/\epsilon)$ on a weighted graph of size $\tilde{O}(m + p^2 n^{1+1/p} \log(nW))$. Note that if $1/\epsilon \leq \text{polylog } n$, then $n^{1/p} = n^{o(1)}$. Thus, Procedure 5 will then compute an $(n^{o(1)}, o(1))$ -hop set of size $O(n^{1+o(1)} \log W)$ and it will perform $\tilde{O}(\log W)$ cluster computations with $p = \Theta(\sqrt{\log n / \log \log n})$ priorities up to distance range $O(n^{o(1)})$ on graphs of size $O(m^{1+o(1)} \log W)$ each, where m is the number of edges of the input graph.

4. Distributed Approximate Single-Source Shortest Paths Algorithm on Networks with Arbitrary Topology. In this section we describe a deterministic distributed algorithm for computing distances from a source node s . It consists of two parts. The first part is constructing a suitable *overlay network*. A *randomized* construction algorithm was given in [58] such that it was sufficient to solve SSSP on the resulting overlay network in order to solve the same problem on the whole network. We give a *deterministic* version of this result in Subsection 4.1. The second part is a more efficient algorithm for computing SSSP on an overlay network using Procedures 1, 2, 3, 4, and 5 from before (see Subsection 4.2). In Subsection 4.3, we show how to finish the computation after combining the two parts following [58].

4.1. Computing an Overlay Network Deterministically. An *overlay network* (also known as *landmark* or *skeleton* [67, 52]) as defined in [58] is a *virtual* network $G' = (V', E')$ of nodes V' and “virtual edges” E' that is built on top of an underlying *real* network $G = (V, E)$; i.e., $V' \subseteq V$ and $E' = V' \times V'$ such that the weight of an edge in G' is an approximation of the distance of its endpoints in G and is ∞ if no path exists between them in G . The nodes in V' are called *centers*. Computing G' means that after the computation every node in G' knows whether it is a center and knows all virtual edges to its neighbors in G' and the corresponding weights. We show in this subsection that there is a $\tilde{O}(\sqrt{n}/\epsilon + D)$ -time algorithm that constructs an




Fig. 2: An overview of the main steps of our algorithm. The left picture depicts the input graph. Thick edges and nodes (in black) in the middle picture depict a possible overlay network. Dashed edges (in red) in the right picture depict a possible hop set of the overlay network.

overlay network G' of $\tilde{O}(\sqrt{n}/\epsilon)$ nodes such that a $(1 + \epsilon/3)$ -approximation to SSSP in G' , can be converted to a $(1 + \epsilon)$ -approximation to SSSP in G , as stated formally below.

THEOREM 19. *In the broadcast CONGEST model, given any weighted undirected network $G = (V, E)$ with polynomially bounded positive integer edge weights and source node s and a parameter $0 < \epsilon \leq 1$, there is an $O(\sqrt{n}\lambda(\lambda + \log(nW)) \log(nW)/\epsilon + D)$ -time deterministic distributed algorithm that computes an overlay network $G' = (V', E')$ and some additional information for every node with the following properties.*

- Property 1: $|V'| = O(\sqrt{n}\lambda \log(nW)/\epsilon)$ and $s \in V'$.
- Property 2: For every node $u \in V$, as soon as u receives a $(1 + \epsilon/3)$ -approximation $\tilde{d}(s, v)$ of $d(s, v, G')$ for all centers $v \in V'$, it can infer a $(1 + \epsilon)$ -approximation of $d(s, u, G)$ without any additional communication.

Note that in this paper we assume that $\lambda = O(\log n)$ and $W = \text{poly}(n)$ and hence the overlay network has size $|V'| = \tilde{O}(\sqrt{n}/\epsilon)$ and the running time is $\tilde{O}(\sqrt{n}/\epsilon + D)$. Observe that the statement of the theorem makes our algorithm very modular by separating the tasks of (i) constructing the overlay network and (ii) computing a $(1 + \epsilon/3)$ -approximation of $d(s, v, G')$ for all centers $v \in V'$. In [Subsection 4.2](#) we show how to perform the second task by implementing the hop-set algorithm of [Section 3](#). It could however be replaced by any other algorithm providing such a $(1 + \epsilon/3)$ -approximation, as is for example done in the recent approximate SSSP algorithm by Becker et al. [\[7\]](#), which also benefits from our deterministic construction of the overlay network.

Before proving the above theorem, we first recall how similar guarantees were achieved with a randomized algorithm in [\[58\]](#)¹² (see Theorem 4.2 of the arXiv version¹³ of [\[58\]](#) for details).

- In the first step of [\[58\]](#), the algorithm selects each node to be a center with probability $\Theta(1/\sqrt{n})$ and also makes s a center. By a standard “hitting set” argument (e.g. [\[72, 16\]](#)), any shortest path containing \sqrt{n} edges will contain a center with high probability. Also, the number of centers is $\Theta(\sqrt{n})$ with high probability.
- In the second step, the algorithm makes sure that every node v knows $(1 + O(\epsilon))$ -

¹²We note that [\[58\]](#) proved this theorem for general parameters λ and α but we will only need it for $\lambda = \alpha = \sqrt{n}$.

¹³<http://arxiv.org/pdf/1403.5171v2.pdf>

G	Graph defining underlying network, input of algorithm
G'	Graph defining overlay network, output of algorithm
ϵ	Input parameter governing approximation quality
V	Node set of G , $ V = n$
V'	Node set G' called centers, $V' \subseteq V$, $ V' = O(\sqrt{n}\lambda \log(nW)/\epsilon)$
W	Maximum edge weight of G
D	Diameter of G
λ	Number of bits used to represent each ID in the network G
$\tilde{\epsilon}$	Parameter set to $\tilde{\epsilon} = \frac{\epsilon}{10}$
$\hat{\epsilon}$	Parameter set to $\hat{\epsilon} = \frac{\epsilon}{18\lambda}$
h	Parameter set to $h = \lfloor \hat{\epsilon}\sqrt{n} \rfloor$
h'	Parameter set to $h' = (1 + 2/\hat{\epsilon})h$
h^*	Parameter set to $h^* = 9\lambda\sqrt{n}$
k	Parameter set to $k = 2h^* + 2\sqrt{n}$
k'	Parameter set to $k' = (1 + 2/\hat{\epsilon})k$
j	Generic index variable $0 \leq j \leq \lfloor \log(nW) \rfloor$
ρ_j	Rounding factor set to $\rho_j = \frac{\tilde{\epsilon}2^j}{h}$
\hat{G}_j	Graph with edge weights $w(u, v, \hat{G}_j) = \lceil \frac{w(u, v, G)}{\rho_j} \rceil$
φ_j	Rounding factor set to $\varphi_j = \frac{\tilde{\epsilon}2^j}{k}$
\tilde{G}_j	Graph with edge weights $w(u, v, \tilde{G}_j) = \lceil \frac{w(u, v, G)}{\varphi_j} \rceil$
$B(v, \hat{G}_j, h')$	Ball of radius h' around v in \hat{G}_j
$t(v)$	Type of node v , smallest j such that $ B(v, \hat{G}_j, h') \geq h$
U_j	Set of nodes of type j , $U_j \subseteq V$
T_j	$(2h' + 1, (2h' + 1)\lambda)$ -ruling set T_j for \hat{G}_j of base set U_j
$d^{h^*}(u, v, G)$	h^* -hop distance between u and v in G
$d^k(u, v, G)$	k -hop distance between u and v in G
$\hat{d}(u, v)$	$(1 + \tilde{\epsilon})$ -approximation of $d^k(u, v, G)$
$\tilde{d}(s, v)$	$(1 + \epsilon/3)$ -approximation of $d(s, v, G')$

Table 1: Overview of notation used in [Subsection 4.1](#)

approximate $\tilde{\Theta}(\sqrt{n})$ -hop distances between v and all centers using a *light-weight bounded-hop single-source shortest paths algorithm* from all centers in parallel combined with the *random delay* technique to avoid congestion.

Let us now give an overview of our new approach. We derandomize the first step as follows: In [Subsection 4.1.1](#) we assign to each node u a *type*, denoted by $t(u)$. (To compute these types, we invoke the *source detection algorithm* of Lenzen and Peleg [54], as we will explain in [Subsection 4.1.1](#).) The important property of nodes' types is that every path π containing \sqrt{n} edges contains a special node u of a "desired" type, meaning that $t(u)$ is not too big compared to $w(\pi, G)$ (see [Lemma 20](#) for details). This is comparable to the property obtained from the hitting set argument, which would be achieved if we made the special node of *every* path a center. However, this may create too many centers (we want the number of centers to be $\tilde{O}(\sqrt{n}/\epsilon)$). Instead we select some nodes to be centers using the *ruling set* algorithm, as described in [Subsection 4.1.2](#). After this, we get a small set of centers such that every node u of type $t(u)$ is not far from one of the centers. Thus, while we cannot guarantee that the path π *contains* a center, we can guarantee that it contains a node that is *not far*

from a center (see Lemma 21 for details).

To derandomize the second step, we use the recent algorithm of Lenzen and Patt-Shamir [53] for the *Partial Distance Estimation* (PDE) problem together with the above Procedures 1, 2, 3, 4, and 5 as we will explain in Subsection 4.1.3.

Since this part of the paper is particularly dense in notation, we summarize the notation used in this subsection in Table 1.

4.1.1. Types of Nodes. Our algorithm initially spends $O(D)$ rounds to make n and λ (and if necessary ϵ) global knowledge. Every node internally sets $\tilde{\epsilon} = \epsilon/10$, $\hat{\epsilon} = \tilde{\epsilon}/(18\lambda)$, $h = \lceil \hat{\epsilon}\sqrt{n} \rceil$, and $h' = (1 + 2/\hat{\epsilon})h$. Note that $h' \leq 3\sqrt{n}$.

For any integer $0 \leq j \leq \lfloor \log nW \rfloor$, we let $\rho_j = \frac{\hat{\epsilon}2^j}{h}$ and let \hat{G}_j be the graph with the same nodes and edges as G and weight $w(u, v, \hat{G}_j) = \lceil \frac{w(u, v, G)}{\rho_j} \rceil$ for every edge (u, v) . Note that we have chosen h' such that $d(u, v, \hat{G}_j) \leq h'$ for all pairs of nodes u and v such that $2^j \leq d(u, v, G) \leq 2^{j+1}$ by Equation (3) of Lemma 3. For any node u , let the ball of u in \hat{G}_j be $B(u, \hat{G}_j, h') = \{v \in V \mid d(u, v, \hat{G}_j) \leq h'\}$. Note that for any index j and nodes u and v , $d(u, v, \hat{G}_{j+1}) \leq d(u, v, \hat{G}_j)$; thus, $B(u, \hat{G}_j, h') \subseteq B(u, \hat{G}_{j+1}, h')$. Let the type $t(u)$ of u be the smallest index j such that $|B(u, \hat{G}_j, h')| \geq h$. We crucially exploit the following structural property.

LEMMA 20. *For every path π of G consisting of $|\pi| = \sqrt{n}$ edges there is a node u on π such that $2^{t(u)} \leq 2\hat{\epsilon}w(\pi, G)$.*

Proof. Let $\ell = \lceil |\pi|/h \rceil \geq 1/\epsilon'$ and let x and y denote the endpoints of π . Partition π into the path π_x consisting of the $(\ell - 1)h$ edges closest to x and the path π_y consisting of the $|\pi| - (\ell - 1)h$ edges closest to y . Further partition π_x into $\ell - 1$ non-overlapping subpaths of exactly h edges, and expand the path π_y by adding edges of π_x to it until it has h edges. Thus, there are now ℓ paths of exactly h edges each and total weight at most $2w(\pi, G)$. It follows that there exists a subpath π' of π consisting of exactly h edges and weight at most $2w(\pi, G)/\ell \leq 2\hat{\epsilon}w(\pi, G)$. Let u and v be the two endpoints of π' and let j be the index such that $2^j \leq d^h(u, v, G) \leq 2^{j+1}$. By Equation (3) of Lemma 3 it follows that $d(u, v, \hat{G}_j) \leq h'$, which implies that $B(u, \hat{G}_j, h')$ contains π' . Hence $|B(u, \hat{G}_j, h')| \geq h$ and $t(u) \leq j$. This shows that $2^{t(u)} \leq 2^j \leq d^h(u, v, G) \leq w(\pi', G) \leq 2\hat{\epsilon}w(\pi, G)$. \square

Computing Types of Nodes. To compute $t(u)$ for all nodes u , it is sufficient for every node u to know, for each j , whether $|B(u, \hat{G}_j, h')| \geq h$. We do this by solving the (S, γ, σ) -detection problem on \hat{G}_j with $S = V$, $\gamma = h'$ and $\sigma = h$, i.e., we compute the list $\mathcal{L}(u, S, \gamma, \sigma, G)$ for all nodes u , which contains the σ nodes from S that are closest to u , provided their distance is at most γ . By Theorem 5 this requires $O(\gamma + \sigma) = O(h + h') = O(\sqrt{n})$ rounds. For any node u , $|\mathcal{L}(u, V, h', h, G)| = h$ if and only if $|B(u, \hat{G}_j, h')| \geq h$. Thus, after we solve the (S, γ, σ) -detection problem on all \hat{G}_j , using $O(\sqrt{n} \log(nW))$ rounds, every node u can compute its type $t(u)$ without any additional communication.

4.1.2. Selecting Centers via Ruling Sets. Having computed the types of the nodes, we compute ruling sets for the nodes of each type to select a small subset of nodes of each type as centers. Remember the two properties of an (α, β) -ruling set T of a base set U : (1) all nodes of T are at least distance α apart and (2) each node in $U \setminus T$ has at least one “ruling” node of T in distance β . We use the algorithm of Theorem 7 to compute, for every $0 \leq j \leq \lfloor \log nW \rfloor$, a $(2h' + 1, (2h' + 1)\lambda)$ -ruling set T_j for \hat{G}_j where the input set U_j consists of all nodes of type j . The number of rounds for

this computation is $O(h' \log(nW)) = O(\sqrt{n} \log(nW))$. We define the set of centers as $V' = (\bigcup_{0 \leq j \leq \lfloor \log(nW) \rfloor} T_j) \cup \{s\}$. Property (1) allows us to bound the number of centers and by property (2) the centers “almost” hit all paths with \sqrt{n} edges.

LEMMA 21. (1) *The number of centers is $|V'| = O(\sqrt{n} \lambda \log(nW)/\epsilon)$. (2) For any path π containing exactly \sqrt{n} edges, there is a node u in π and a center $v \in V'$ such that $d^{h^*}(u, v, G) \leq \tilde{\epsilon} w(\pi, G)$, where $h^* = 9\sqrt{n} \lambda$.*

Proof. (1) For each j , consider any two nodes u and v in T_j . Since $d(u, v, \hat{G}_j) > 2h'$ by Property (1) of the ruling set, $B(u, \hat{G}_j, h') \cap B(v, \hat{G}_j, h') = \emptyset$. As every node $u \in T_j$ is of type j , $|B(u, \hat{G}_j, h')| \geq h$ for every $u \in T_j$. We can therefore uniquely assign h nodes to every node $u \in T_j$ and thus $|T_j| \leq n/h = O(\sqrt{n} \lambda / \epsilon)$.

(2) By Lemma 20, there is a node u in π such that $2^{t(u)} \leq 2\hat{\epsilon} w(\pi, G)$. Moreover, there is a center v in the ruling set $T_{t(u)}$ such that

$$(10) \quad d(u, v, \hat{G}_{t(u)}) \leq (2h' + 1)\lambda \leq 3h'\lambda \leq h^*,$$

where the second inequality is because $h' \leq 3\sqrt{n}$. Let π' be the shortest path between u and v in $\hat{G}_{t(u)}$. Then $w(\pi', \hat{G}_{t(u)}) = d(u, v, \hat{G}_{t(u)}) \leq h^*$, and as a consequence π' contains at most h^* edges. It follows that

$$\begin{aligned} d^{h^*}(u, v, G) &\leq w(\pi', G) && \text{(since } \pi' \text{ is } u\text{-}v \text{ path with } \leq h^* \text{ edges)} \\ &= \sum_{(x,y) \in E(\pi')} w(x, y, G) \\ &\leq \sum_{(x,y) \in E(\pi')} \rho_{t(u)} \cdot w(x, y, \hat{G}_{t(u)}) && \text{(since } w(x, y, \hat{G}_{t(u)}) = \lceil \frac{w(x,y,G)}{\rho_{t(u)}} \rceil) \\ &= \rho_{t(u)} \cdot w(\pi', \hat{G}_{t(u)}) \\ &= \rho_{t(u)} \cdot d(u, v, \hat{G}_{t(u)}) && \text{(since } \pi' \text{ shortest } u\text{-}v \text{ path in } \hat{G}_{t(u)}) \\ &= \frac{\hat{\epsilon} 2^{t(u)}}{h} \cdot d(u, v, \hat{G}_{t(u)}) && \text{(since } \rho_{t(u)} = \frac{\hat{\epsilon} 2^{t(u)}}{h}) \\ &\leq \frac{\hat{\epsilon} 2^{t(u)}}{h} \cdot 3h'\lambda && \text{(by (10))} \\ &\leq 9\lambda 2^{t(u)} && \text{(since } h' = (1 + \frac{2}{\hat{\epsilon}})h \leq \frac{3h}{\hat{\epsilon}}) \\ &\leq 18\lambda \hat{\epsilon} w(\pi, G) && \text{(by Lemma 20)} \\ &= \tilde{\epsilon} w(\pi, G). && \text{(since } \hat{\epsilon} = \frac{\tilde{\epsilon}}{18\lambda}) \end{aligned}$$

□

4.1.3. Computing Distances to Centers. Let $k = 2h^* + 2\sqrt{n}$, where $h^* = 9\sqrt{n} \lambda$ (as in Lemma 21), and $k' = (1 + 2/\hat{\epsilon})k$. In this step, we compute for every node u and every center v a value $\hat{d}(u, v)$ that is a $(1 + \hat{\epsilon})$ -approximation of $d^k(u, v, G)$ such that each node u knows $\hat{d}(u, v)$ for all centers v . In particular we also compute $\hat{d}(u, v)$ for all pairs of centers u and v . To do this we follow the idea of partial distance estimation [53]. As in Subsection 4.1.1, we do this by solving the source detection problem on a graph with rounded weights.¹⁴ For every integer $0 \leq j \leq \lfloor \log nW \rfloor$,

¹⁴We note that the algorithm and analysis described in this subsection is essentially the same as the proof of [53, Theorem 3.3]. We cannot use the result in [53] directly since we need a slightly stronger guarantee, which can already be achieved by the same proof. (We thank Christoph Lenzen for a communication regarding this.)

let $\varphi_j = \frac{\tilde{\epsilon}2^j}{k}$ and let \tilde{G}_j be the weighted graph such that $w(u, v, \tilde{G}_j) = \lceil \frac{w(u, v, G)}{\varphi_j} \rceil$ for every edge (u, v) in G .

We solve the (S, γ, σ) -detection problem on \tilde{G}_j for all $0 \leq j \leq \lfloor \log nW \rfloor$, with parameters $S = V'$, $\gamma = k' = (1 + 2/\tilde{\epsilon})k = O(\sqrt{n}\lambda^2)$, and $\sigma = |V'|$, where $|V'| = O(\sqrt{n}\lambda \log(nW)/\epsilon)$ by Lemma 21. Using the algorithm of Theorem 5 for each graph \tilde{G}_j this takes $O((\gamma + \sigma) \log(nW)) = O(\sqrt{n}\lambda(\lambda + \log(nW)) \log(nW)/\epsilon)$ rounds. At termination, every node u knows the distances up to distance range k' to all centers in all \tilde{G}_j ; i.e., it knows $d(u, v, k', \tilde{G}_j)$ for all j and all centers v . For every node $u \in V$ and every center $v \in V'$ we set $\hat{d}(u, v) = \min_{0 \leq j \leq \lfloor \log nW \rfloor} \{\varphi_j \cdot d(u, v, k', \tilde{G}_j)\}$. Every node u can compute $\hat{d}(u, v)$ without any additional communication as soon as the source detection algorithm is finished. Now consider the index j^* such that $2^{j^*} \leq d^h(u, v, G) \leq 2^{j^*+1}$. It follows from Equation (3) of Lemma 3 that $d(u, v, \tilde{G}_{j^*}) \leq k'$ which implies that $d(u, v, k', \tilde{G}_{j^*}) = d(u, v, \tilde{G}_{j^*})$. With Equations (2) and (4) we then get

$$(11) \quad \hat{d}(u, v) \leq \varphi_{j^*} \cdot d(u, v, k', \tilde{G}_{j^*}) = \varphi_{j^*} \cdot d(u, v, \tilde{G}_{j^*}) \leq (1 + \tilde{\epsilon})d^k(u, v, G).$$

Hence $\hat{d}(u, v)$ is the desired $(1 + \tilde{\epsilon})$ -approximation of $d^k(u, v, G)$.

4.1.4. Completing the Proof of Theorem 19. We define our final overlay network to be the graph G' where the weight between any two centers $u, v \in V'$ is $\hat{d}(u, v)$ (as computed in Subsection 4.1.3). Additionally, for every node $u \in V$ we store the value of $\hat{d}(u, v)$ to all centers $v \in V'$. We now show that all properties stated in Theorem 19 hold for G' . Since we need $O(\sqrt{n} \log(nW))$ rounds in Subsections 4.1.1 and 4.1.2 and $(\sqrt{n}\lambda(\lambda + \log(nW)) \log(nW)/\epsilon)$ rounds in Subsection 4.1.3, the running time to construct G' is $O(\sqrt{n}\lambda(\lambda + \log(nW)) \log(nW)/\epsilon)$. Moreover, $|V'| = O(\sqrt{n}\lambda \log(nW)/\epsilon)$ as shown in Lemma 21. This is as claimed in the first part of Theorem 19. It is thus left to prove the following statement in Theorem 19: “for every node $u \in V$, as soon as u receives a $(1 + \epsilon/3)$ -approximation $\tilde{d}(s, v)$ of $d(s, v, G')$ for all centers $v \in V'$, it can infer a $(1 + \epsilon)$ -approximate value of $d(s, u, G)$ without any additional communication.” Recall that in Subsection 4.2 we show how compute, and make known to all nodes, the values $\tilde{d}(s, v)$ for all centers $v \in V'$.

Consider any node u and let π be the shortest path between s and u in G . If π contains less than \sqrt{n} edges, then $d^k(s, u, G) = d(s, u, G)$ and thus the value $\hat{d}(s, u)$, which is a $(1 + \tilde{\epsilon})$ -approximation of $d^k(s, u, G)$ known by u , is already a $(1 + \tilde{\epsilon})$ -approximation of $d(s, u, G)$ (and thus, by the choice of $\tilde{\epsilon} = \epsilon/10$, also a $(1 + \epsilon)$ -approximation). If π contains at least \sqrt{n} edges, then partition π into subpaths $\pi_0, \pi_1, \dots, \pi_\ell$ (for some $\ell \geq 0$) where π_0 contains s , π_ℓ contains u , π_0 contains at most \sqrt{n} edges, and every subpath except π_0 contains exactly \sqrt{n} edges. By Lemma 21, for every $1 \leq i \leq \ell$, there is a node x_i and a center y_i such that (i) x_i is in π_i and (ii) $d^{h^*}(x_i, y_i, G) \leq \tilde{\epsilon}w(\pi_i, G)$. Additionally, for $1 \leq i \leq \ell - 1$, since x_i and x_{i+1} lie on π , their shortest path is the subpath of π between them and, thus, it consists of at most $2\sqrt{n}$ edges. It follows that $d^{2\sqrt{n}}(x_i, x_{i+1}, G) = d(x_i, x_{i+1}, G)$. By our choice of $k = 2h^* + 2\sqrt{n}$, the triangle inequality and symmetry (i.e., $d(v, v', G) = d(v', v, G)$) give, for every $1 \leq i \leq \ell - 1$

$$(12) \quad \begin{aligned} d^k(y_i, y_{i+1}, G) &\leq d^{h^*}(y_i, x_i, G) + d^{2\sqrt{n}}(x_i, x_{i+1}, G) + d^{h^*}(x_{i+1}, y_{i+1}, G) \\ &\leq \tilde{\epsilon}w(\pi_i, G) + d(x_i, x_{i+1}, G) + \tilde{\epsilon}w(\pi_{i+1}, G). \end{aligned}$$

By the same argument

$$(13) \quad d^k(s, y_1, G) \leq d^{2\sqrt{n}}(s, x_1, G) + d^{h^*}(x_1, y_1, G) \leq d(s, x_1, G) + \tilde{\epsilon}w(\pi_1, G)$$

and

$$(14) \quad d^k(y_\ell, u, G) \leq d^{h^*}(y_\ell, x_\ell, G) + d^{\sqrt{n}}(x_\ell, u, G) \leq d(x_\ell, u, G) + \tilde{\epsilon}w(\pi_\ell, G).$$

We now argue that $\tilde{d}(s, y_\ell) + \hat{d}(u, y_\ell)$ (the sum of two values known to node u) is a $(1 + \epsilon)$ -approximation of $d(s, u, G)$. First, since $\tilde{d}(s, y_\ell)$ is a $(1 + \epsilon/3)$ -approximation of $d(s, y_\ell, G')$, we get

$$\tilde{d}(s, y_\ell) + \hat{d}(u, y_\ell) \leq (1 + \epsilon/3)d(s, y_\ell, G') + \hat{d}(u, y_\ell) \leq (1 + \epsilon/3)(d(s, y_\ell, G') + \hat{d}(u, y_\ell)).$$

We now apply the triangle inequality, exploit that every edge (x, y) in G' has weight $\hat{d}(x, y)$ (implying $d(x, y, G') \leq \hat{d}(x, y)$), and use Equations (11)–(14) from above to get

$$\begin{aligned} d(s, y_\ell, G') + \hat{d}(u, y_\ell) &\leq \left(d(s, y_1, G') + \sum_{i=1}^{\ell-1} d(y_i, y_{i+1}, G') \right) + \hat{d}(u, y_\ell) \\ &\leq \left(\hat{d}(s, y_1) + \sum_{i=1}^{\ell-1} \hat{d}(y_i, y_{i+1}) \right) + \hat{d}(u, y_\ell) \\ &\stackrel{(11)}{\leq} (1 + \tilde{\epsilon}) \left(d^k(s, y_1, G) + \sum_{i=1}^{\ell-1} d^k(y_i, y_{i+1}, G) \right) + (1 + \tilde{\epsilon})d^k(u, y_\ell) \\ &\leq (1 + \tilde{\epsilon}) \left(d^k(s, y_1, G) + \sum_{i=1}^{\ell-1} d^k(y_i, y_{i+1}, G) + d^k(u, y_\ell, G) \right) \\ &\stackrel{(12)-(14)}{\leq} (1 + \tilde{\epsilon}) \left(d(s, u, G) + 2\tilde{\epsilon} \sum_{i=1}^{\ell} w(\pi_i, G) \right) \\ &\leq (1 + \tilde{\epsilon}) (d(s, u, G) + 2\tilde{\epsilon}d(s, u, G)) \\ &= (1 + \tilde{\epsilon})(1 + 2\tilde{\epsilon})d(s, u, G) \\ &\leq (1 + 5\tilde{\epsilon})d(s, u, G) \\ &= (1 + \epsilon/2)d(s, u, G). \end{aligned}$$

By combining the two derivations above, we get

$$\tilde{d}(s, y_\ell) + \hat{d}(u, y_\ell) \leq (1 + \epsilon/3)(1 + \epsilon/2)d(s, u, G) \leq (1 + \epsilon)d(s, u, G).$$

Thus, when u receives $\tilde{d}(s, v')$ for all centers $v \in V'$, it can compute the value $\min_{v \in V'} (\tilde{d}(s, v) + \hat{d}(u, v))$ and, as $y_\ell \in V'$, the argument above show that this value is a $(1 + \epsilon)$ -approximation of $d(s, u, G)$.

4.2. Computing a Hop Set on an Overlay Network. We now show how to implement the algorithm to compute the hop set on the overlay network G' presented in Section 3 and how to compute approximate shortest paths from s in G' using the hop set presented in Section 3. We let G' be the overlay network obtained from Theorem 19 with $\epsilon = 1/\log n$ (to guarantee a $(1 + o(1))$ -approximation in the end). Throughout the algorithm we will work on overlay networks whose node set is the set of centers V' , but which might have different edge weights as, e.g., Procedure 4 calls Procedure 3 on overlay networks with modified edge weights. Thus, we will use G'' to refer to an overlay network (with set of nodes V') on which Procedures 1, 2, 3, and 4 run to emphasize the fact that they might not equal G' . We let N be the number of centers in G' and G'' . Thus $N = \tilde{O}(n^{1/2})$.

4.2.1. Computing Bounded-Distance Single-Source Shortest Paths. We will repeatedly use an algorithm for computing a shortest-path tree up to distance R rooted at some center s on an overlay network G'' , where $R = N^{o(1)}$. At the end of the algorithm every center knows this tree. We do this in a breadth-first search manner, in $R + 1$ iterations. Like in Dijkstra’s algorithm, every center keeps a tentative distance $\delta(s, u)$ from s and a tentative parent in the shortest-path tree, i.e., a center v such that $\delta(s, u) = \delta(s, v) + w(u, v, G'')$. Initially, $\delta(s, s) = 0$ and $\delta(s, v) = \infty$ for every center $v \neq s$. In the L^{th} iteration, for L from 0 up to R , all centers in G'' whose tentative distance $\delta(s, u)$ is exactly L broadcast¹⁵ to all other centers a message $(u, \delta(s, u), v)$ where v is the parent of u . Using this information, every center u will update (“relax”) its tentative distance $\delta(s, u)$ and its tentative parent.

By a straightforward induction, after the L^{th} iteration, centers that have distance L from s (i.e., that are at level L in the shortest-path tree) will already know their correct distance. Thus, at the end of the last iteration every center knows the shortest-path tree rooted at s up to distance R in G'' . To analyze the running time, note that over R rounds we broadcast N messages in total, and if M_L messages are broadcast in the L^{th} iteration, then this iteration takes $O(M_L + D)$ rounds. (We emphasize that the number of rounds depends on the diameter D of the original network, and not of G'' .) The total number of communication rounds used over all iterations is thus $O(RD + \sum_L M_L) = O(N + RD)$.

4.2.2. Computing Priorities. We implement Procedure 1 on an overlay network G'' . All necessary parameters can be computed beforehand and thus do not require any communication. Initially every center knows that it is contained in $A_0 = V'$. To compute A_{i+1} given that A_i is known (i.e., every center knows whether it is in A_i or not), we compute the proximity list $\mathcal{L}(v, A_i, R, q, G'')$ for every center v using a source detection algorithm and distribute each list to every center, where, by our choice of parameters, $R = N^{o(1)}$ and $q = N^{o(1)}$. Then every center runs the same deterministic greedy hitting set approximation algorithm to compute A_{i+1} .¹⁶ We will obtain $\mathcal{A} = (A_i)_{0 \leq i \leq p}$ by repeating this for p iterations. Thus, we have to solve the (S, γ, σ) -source detection problem with $S = A_i$, $\gamma = R$, and $\sigma = q$ on an overlay network G'' . For this purpose we simulate the source detection algorithm of Lenzen and Peleg [54] (see Theorem 5 and preceded description of the algorithm) as if run on the overlay network.

The simulated source detection algorithm consists of at most $\gamma + \sigma$ iterations. Since in each iteration of the source detection algorithm, each center sends *the same* message to all of its neighbors, we can simulate each iteration by broadcasting at most N messages in the underlying network G . Thus, simulating the source detection algorithm takes $O((\sigma + \gamma)(N + D))$ rounds. To compute the priorities, we repeat this process for all $p \leq \log n$ priorities. The overall running time for implementing Procedure 1 therefore is $O((N + D)p(R + q))$. With our choice of parameters ($N = \tilde{O}(n^{1/2})$, $p \leq \log n$, $q = N^{o(1)}$, and $R = N^{o(1)}$), this becomes

$$(15) \quad n^{1/2+o(1)} + Dn^{o(1)}.$$

¹⁵More precisely, there is a designated node (e.g. the node with lowest ID) that aggregates and distributes the messages (via upcasting and downcasting on the breadth-first search tree of the underlying network G), and tells other center when the iteration starts and ends.

¹⁶Note that the number of internal computation steps of the greedy algorithm at each center is linear in its input, which we can upper bound by the number of messages received in all rounds.

4.2.3. Computing Clusters. We now describe how to compute clusters on an overlay network G'' such that at the end of this computation, every center will know all clusters $C(v, \mathcal{A}, R, G'')$ (i.e., its own cluster and the cluster of every other center). We do this by implementing Procedure 2 on G'' . First, we need to compute $d(v, A_{i+1}, R, G'')$, for every $1 \leq i \leq p$. This can be done in exactly the same way as in the first phase of computing the hierarchy $\mathcal{A} = (A_i)_{0 \leq i \leq p}$; i.e., we add a virtual source s^* and edges of weight zero between s^* and centers in A_i , and compute the shortest-path tree up to distance R rooted at s^* . Since such a tree can be computed in $O(N + RD)$ rounds and we have to compute $p \leq \log n$ such trees, the total time we need here is $\tilde{O}(N + RD)$.

Next, we use the information gained above to compute the cluster up to distance R from every center u in G'' , as described in Procedure 2. That is, in iteration L (starting with $L = 0$ and ending with $L = R$), every center v having (i) $\delta(u, v)$ (the tentative distance from u to v) equal to L and (ii) additionally $\delta(u, v) < d(v, A_{i+1}, R, G'')$ will broadcast¹⁷ its distance to u to all other centers so that every other center, say w , can (i) update its tentative distance $\delta(u, w)$ and (ii) add v and $\delta(u, v)$ to its locally stored copy of $C(u)$. Thus, there are $\sum_{v \in V'} |C(v, \mathcal{A}, R, G'')|$ messages broadcast in total, which is bounded from above by $\tilde{O}(pN^{1/p}) = n^{1/2+o(1)}$ due to Theorem 12.

Note that this procedure computes $C(v, \mathcal{A}, R, G'')$, for all centers v , *in parallel*. Each iteration L requires $O(\sum_{v \in V'} M_{v,L} + D)$ rounds, where $M_{v,L}$ is the number of messages broadcast by node v in iteration L in the above computation. The total number of rounds over all R iterations is thus

$$(16) \quad O\left(\sum_{0 \leq L \leq R} \sum_{v \in V'} M_{v,L} + RD\right) = O\left(\sum_{v \in V'} |C(v, \mathcal{A}, R, G'')| + RD\right) = n^{1/2+o(1)} + Dn^{o(1)}.$$

Note that since the computation is done by broadcasting messages, every center knows the cluster $C(v, \mathcal{A}, R, G'')$ for all v at the end of this computation. Together with the running time bound of Equation (15) for computing the priorities, we arrive at the following guarantees.

LEMMA 22. *For any overlay network $G'' = (V', E'')$ with $N = \tilde{O}(n^{1/2})$ centers, the above algorithm, in $n^{1/2+o(1)} + Dn^{o(1)}$ rounds, deterministically computes a hierarchy of centers $\mathcal{A} = (A_i)_{0 \leq i \leq p}$ and clusters $C(v, \mathcal{A}, R, G'')$ for each center v as specified in Theorem 12 with $p \leq \log n$ priorities up to distance $R = N^{o(1)}$ such that $\sum_{v \in V'} |C(v, \mathcal{A}, R, G'')| = n^{1+o(1)}$ (and every center knows $C(v, \mathcal{A}, R, G'')$ for all centers v as well as the value of $d(v, w, G'')$ for every center v and every center $w \in C(v, \mathcal{A}, R, G'')$).*

4.2.4. Computing the Hop Reduction with Additive Error. We implement Procedure 3 on an overlay network G'' . All necessary parameters can be computed beforehand and thus do not require any communication. We can then execute $\text{CLUSTERS}(G'', p, R)$ using the above algorithm to get $(C(v, \mathcal{A}, R, G''), \delta(v, \cdot))_{v \in V'}$. With this information, the set F , as specified in Procedure 3, can be computed without any additional communication. Thus, executing $\text{CLUSTERS}(G'', p, R)$ is the only part

¹⁷We note again that to do this, there is a designated center that aggregates and distributes the messages (via upcasting and downcasting), and tells other centers when the iteration starts and ends.

of computing F that requires communication. By Lemma 22 the total time needed to execute Procedure 3 is therefore

$$(17) \quad n^{1/2+o(1)} + Dn^{o(1)}.$$

4.2.5. Computing the Hop Reduction without Additive Error. We implement Procedure 4 on an overlay network G'' . All necessary parameters can be computed beforehand and thus do not require any communication. Moreover, every center knows about the edges incident to it and we can thus implicitly compute \hat{G}_j , as specified in Procedure 4, by scaling down edge weights without any communication. We then execute Procedure 3 to compute \hat{F}_j . Knowing \hat{F}_j , we can compute F without any additional communication. Thus, executing Procedure 3 is the only part of computing F that requires communication and it is executed $O(\log(nW))$ times. As our implementation of Procedure 3 takes time $n^{1/2+o(1)} + Dn^{o(1)}$, as argued above (cf. Equation (17)), the total time needed to execute Procedure 4 is

$$(18) \quad n^{1/2+o(1)} \log W + Dn^{o(1)} \log W.$$

4.2.6. Computing the Hop Set. We implement Procedure 5 on the overlay network G' . All necessary parameters can be computed beforehand. Computing F_{i+1} is done by calling Procedure 4 on the graph H_i , as specified in Procedure 5, which, as argued above (cf. Equation (18)), takes time $n^{1/2+o(1)} \log W + Dn^{o(1)} \log W$. As every center knows its incident edges, the graph H_{i+1} can be computed from F_{i+1} without any additional communication. As we execute Procedure 4 $p \leq \log n$ times, the total time needed to implement Procedure 5 is $pn^{1/2+o(1)} \log W + pDn^{o(1)} \log W = n^{1/2+o(1)} \log W + Dn^{o(1)} \log W$. By running this algorithm on G' (which, as pointed out, involves performing hop reductions and computing clusters on some other overlay networks), we obtain the following theorem.

THEOREM 23. *In the broadcast CONGEST model, there is a deterministic algorithm that, for any overlay network G' with $N = \tilde{O}(n^{1/2})$ centers and positive integer weights in the range $\{1, \dots, W\}$ on edges between centers, computes an $(n^{o(1)}, o(1))$ -hop set of G' in $n^{1/2+o(1)} \log W + Dn^{o(1)} \log W$ rounds. When the algorithm has finished, every center knows every edge in the hop set.*

4.2.7. Routing via the Hop Set. Remember that the overlay network is computed using the source detection algorithm of Lenzen and Peleg [54]. If a node x of the overlay network wants to send a message to one of its neighbors y in the overlay network it can do so by routing the message along a path in the original network whose length is upper-bounded by the weight of the overlay edge (x, y) . This routing can be obtained by modifying the source detection algorithm to additionally construct BFS trees rooted at the sources (see [54]), which in our case are the nodes of the overlay network.

When we compute the hop set on the overlay network we broadcast all computed clusters to all nodes in the network. In this way the clusters, the corresponding partial shortest path trees of the clusters, as well as the hop set edges become global knowledge. Therefore every node in the overlay network learns for every hop set edge (x, y) to which path from x to y in the overlay network it corresponds. Thus, also for every hop set edge (x, y) of the overlay network, x can send a message to y by routing the message along a path in the original network whose length is upper-bounded by the weight of the overlay edge (x, y) . This means that the hop set computed by our algorithm has the following *awareness* property, as introduced in the full version of

[26]: A hop set F for a graph G is called *aware* if for every hop-set edge $(x, y) \in F$ of weight b there exists a corresponding path π in G between x and y of length b . Furthermore, every node v on π knows $d_\pi(v, x)$ and $d_\pi(v, y)$, and its neighbors on π .

4.3. Final Steps. Let $H = G' \cup F$ be the graph obtained by adding to G' the edges of the $(n^{o(1)}, o(1))$ -hop set F computed above. To $(1 + o(1))$ -approximate $d(s, v, G')$ for every center v in G' , it is sufficient to $(1 + o(1))$ -approximate the h -hop distance $d^h(s, v, H)$, for some $h = n^{o(1)}$. The latter task can be done in $O(hD + |V'|) = n^{o(1)}D + n^{1/2+o(1)}$ rounds by the same method as in Lemma 4.6 in the full version of [58]. We give a sketch here for completeness. Let $\epsilon = 1/\log n$. For any $0 \leq j \leq \lfloor \log(nW) \rfloor$, let \hat{H}_j be the graph obtained by rounding edge weights in H as in Subsection 4.1.1; i.e., for every edge (u, v) we set $w(u, v, \hat{H}_j) = \lceil \frac{w(u, v, H)}{\rho_j} \rceil$, where $\rho_j = \frac{\epsilon 2^j}{h}$. For each \hat{H}_j , we compute the shortest-path tree rooted at s up to distance $R = O(h/\epsilon)$, which can be done in $RD + n^{1/2+o(1)} = n^{o(1)}D + n^{1/2+o(1)}$ rounds, using the algorithm described in Subsection 4.2.1. This gives $d(s, v, R, \hat{H}_j)$ for every center v . We then use the following value as $(1 + o(1))$ -approximation of $d^h(s, v, H)$ (and thus of $d(s, v, G')$): $\tilde{d}(s, v) = \min_j \rho_j \cdot d(s, v, R, \hat{H}_j)$. The correctness of this algorithm follows from Lemma 3.

Once we have $(1 + o(1))$ -approximate values of $d(s, v, G')$ for every center $v \in V'$, we can broadcast these values to the whole network in $O(n^{1/2} + D)$ rounds. Theorem 19 then implies that we have a $(1 + o(1))$ -approximate solution to the single-source shortest paths problem on the original network. The total time spent is $n^{1/2+o(1)} + Dn^{o(1)}$. By observing that the term $n^{o(1)}D$ will show up in the running time only when $D = \omega(n^{o(1)})$, we can write the running time as $n^{1/2+o(1)} + D^{1+o(1)}$, as claimed in the beginning.

We thus have obtained the following result.

THEOREM 24. *In the broadcast CONGEST model, there is a deterministic algorithm that, on any weighted undirected network with polynomially bounded positive integer edge weights, computes $(1 + o(1))$ -approximate shortest paths between a given source node s and every other node in $n^{1/2+o(1)} + D^{1+o(1)}$ rounds.*

5. Algorithms in Other Settings.

5.1. Congested Clique. In the congested clique model, the underlying communication network is a complete graph. Thus, in each round every node can send a message to every other node. Apart from this topological constraint, the congested clique model is similar to the CONGEST model.

We compute an $(n^{o(1)}, o(1))$ -hop set on a congested clique by implementing the hop set construction algorithm in the same way as on the overlay network, as presented in Subsection 4.2. (However, we do *not* compute an overlay network here.) The only difference is the number of rounds needed for nodes to broadcast messages to all other nodes. Consider the situation that M' messages are to be broadcast by some nodes. On a network of arbitrary topology, we will need $O(D + M')$ rounds. On a congested clique, however, we only need $O(M'/n)$ rounds using the routing scheme of Dolev et al. [19, Lemma 1] (also see [51]): If each node is source and destination of up to n messages of size $O(\log n)$ (initially only the sources know destinations and contents of their messages), we will need $O(1)$ rounds to route the messages to their destinations. In particular, we can broadcast n messages in $O(1)$ rounds, and thus M' messages in $O(M'/n)$ rounds. Using this fact, the number of rounds needed for the algorithm in Subsection 4.2 reduces from $O(\sum_{v \in V'} |C(v, \mathcal{A}, R, G')| + RD)$ on the overlay network G'

(cf. Equation (16)) to $O(\sum_{v \in V} |C(v, \mathcal{A}, R, G)|/n + R) = \tilde{O}(pn^{1/p} + R) = n^{o(1)}$ on a congested clique G .¹⁸

Once we have an $(n^{o(1)}, o(1))$ -hop set, we proceed as in Subsection 4.3. Let $H = G \cup F$ be the graph obtained by adding to the input graph G the edges of the $(n^{o(1)}, o(1))$ -hop set F . We can treat H as a congested clique network with edge weights different from G . (H can be computed without any additional communication since every node already knows the hop set.) To $(1 + o(1))$ -approximate $d(s, v, G)$ for every node v in G , it is sufficient to compute the h -hop distance $d^h(s, v, H)$, where $h = n^{o(1)}$. To do this, we follow the same approach for this problem as in [58, Section 5.1], where we execute the distributed version of the Bellman-Ford algorithm for h rounds. That is, every node u maintains a tentative distance from the source s , denoted by $\delta(s, u)$ and in each round every node u broadcasts $\delta(s, u)$ to all other nodes. It can be shown that after k rounds every node v knows the k -hop distance (i.e., $\delta(s, u) = d^k(s, v, H)$) correctly, and thus after h rounds we will get the h -hop distances as desired.¹⁹

THEOREM 25. *In the congested clique model, there is a deterministic algorithm that, on any weighted undirected clique network with polynomially bounded positive integer edge weights, computes $(1 + o(1))$ -approximate shortest paths between a given source node s and every other node in $n^{o(1)}$ rounds.*

5.2. Streaming Algorithm. In the graph streaming model, the edges of the input graph are presented to the algorithm in an arbitrary order. The goal is to design algorithms that process this “stream” of edges using as little space as possible. In the multipass streaming model we are allowed to read the stream several times and want to keep both the number of passes and the amount of space used as small as possible.

Our streaming algorithm for constructing an $(n^{o(1)}, o(1))$ -hop set proceeds in almost the same way as the distributed algorithm in Subsection 4.2. First, observe that a shortest-path tree up to distance R can be computed in $O(R)$ passes and with $\tilde{O}(n)$ space: We use the space to remember the tentative distances of the nodes to s , and the shortest-path tree computed so far. At the end of the L^{th} pass we add nodes having distances exactly L to the shortest-path tree, and update the distance of their neighbors in the $(L + 1)^{\text{th}}$ pass.

We compute the priorities, as described in Subsection 4.2.2, by solving $p \leq \log n$ instances of a (S, γ, σ) -detection problem with $\gamma = R = N^{o(1)}$ and $\sigma = q = N^{o(1)}$. Observe that the guarantees of the source detection algorithm by Lenzen and Peleg for the broadcast CONGEST model directly carry over to the Streaming model by simulating the algorithm as follows:

- The tentative list of each node is stored using $O(\min(\gamma, WD) + \min(\sigma, |S|))$ space as, at any time, each node only needs to know at most $\min(\gamma, WD) + \min(\sigma, |S|)$ entries in its list (upper bounded by the total number of messages each node will send).

¹⁸Instead of relying on the result of Dolev et al., we can use the following algorithm to broadcast M' messages in $O(M'/n)$ rounds as follows. We assign an order to the messages, where messages sent by a node with smaller ID appears first in the order and messages sent by the same node appear in any order (a node can learn the order of its messages after it knows how many messages other nodes have). We then broadcast the first n messages according to this order, say M_1, \dots, M_n , where message M_i is sent to a node with the i^{th} smallest ID, and such a node sends M_i to all other nodes. This takes only two rounds. The next messages are handled similarly. This algorithm broadcasts each n messages using 2 rounds, and thus the total number of rounds is $O(M'/n)$.

¹⁹Note that instead of the Bellman-Ford algorithm, one can also follow the steps in Subsection 4.3 instead. This gives a $(1 + o(1))$ -approximate value for $d^h(s, v, H)$ for every node v , which is sufficient for computing a $(1 + o(1))$ -approximate value for $d(s, v, G)$. This algorithm is however more complicated.

- The broadcast of one message per node in each round is simulated by writing the $O(n)$ messages to space
- The reception of messages in each round is simulated by making a pass over the graph: every time an edge (u, v) and its corresponding weight are read from the stream the reception of u 's message by v is simulated by reading u 's message from space and then manipulating v 's tentative list accordingly.

We can summarize the guarantees of the source detection algorithm in the Streaming model as follows.

THEOREM 26 (Implicit in [54]). *In the multipass streaming model, there is a deterministic algorithm for solving the (S, γ, σ) -detection problem in $\min(\gamma, WD) + \min(\sigma, |S|)$ passes with $O(n \cdot (\min(\gamma, WD) + \min(\sigma, |S|)))$ space.*

The algorithm for computing the priorities therefore needs $O(p(R + q)) = n^{o(1)}$ passes and $O(n(R + q)) = n^{1+o(1)}$ space.

To compute clusters, we compute n shortest-path trees up to distance R rooted at different nodes in parallel. The number of passes is clearly $O(R)$. The space is bounded by the sum of the sizes of the shortest-path trees. This is $O(\sum_{v \in V} |C(v, \mathcal{A}, R, G)|)$ which, by [Theorem 12](#), is $\tilde{O}(pn^{1+1/p}) = n^{1+o(1)}$. To compute the hop set we only have to compute clusters $\tilde{O}(\log W)$ times. So, we need $n^{o(1)} \log W$ passes and $n^{1+o(1)} \log W$ space in total. By considering the edges of the hop set in addition the edges read from the stream it suffices to compute approximate single-source shortest paths up to $n^{o(1)}$ hops. Using the streaming version of the Bellman-Ford algorithm (one pass per iteration), this can be done in $n^{o(1)} \log W$ additional passes.

THEOREM 27. *In the multipass streaming model, there is a deterministic algorithm that, given any weighted undirected graph with polynomially bounded positive integer edge weights, computes $(1 + o(1))$ -approximate shortest paths between a given source node s and every other node in $n^{o(1)}$ passes with $n^{1+o(1)}$ space.*

6. Conclusion and Open Problems. We present deterministic distributed $(1 + o(1))$ -approximation algorithms for solving the single-source shortest paths problem on distributed weighted networks and other settings. The efficiencies of our algorithms match the known lower bounds up to an $n^{o(1)}$ factor. Important tools are a deterministic hop set construction and a deterministic process that replaces the well-known (randomized) hitting set argument.

In the conference version of this paper, we left as an open problem whether the factor of $n^{o(1)}$ in our bounds could be eliminated, and we in particular asked whether this can be done by constructing a $(\text{polylog } n, o(1))$ -hop set of size $\tilde{O}(n)$. Such a hop set construction however can be ruled out by a recent lower bound of Abboud, Bodwin, and Pettie [2]. Our open problem was solved nonetheless by Becker et al. [7] who, using tools from continuous optimization, showed that, in all the models that we considered above, a $(1 + \epsilon)$ -approximation can be obtained with an overhead of $\epsilon^{-O(1)}$ $\text{polylog } n$ compared to known lower bounds.

Our deterministic replacement of the hitting set argument works only when the input graph is undirected. Our second open problem is thus how to derandomize algorithms on directed graphs (where edge directions do not affect the communication; see [58, 59] for more details). In particular, it is known that single-source shortest paths can be $(1 + \epsilon)$ -approximated on directed weighted graphs in $\tilde{O}(n^{1/2}D^{1/2} + D)$ time [58], and single-source reachability can be computed in $\tilde{O}(n^{1/2}D^{1/4} + D)$ time [35]. However, these results are obtained by *randomized* algorithms, and whether there are sublinear-time *deterministic* algorithms for these problems is still open.

Finally, while our paper essentially settles the running time for computing single-source shortest paths approximately, the best running time for solving this problem *exactly* is $O((n \log n)^{2/3} D^{1/3} + (n \log n)^{5/6})$ [23], a recent result obtained after the conference version of our paper appeared. This leaves a gap to the $\tilde{\Omega}(\sqrt{n} + D)$ lower bound [21] and it is therefore natural to ask for an improved upper or lower bound. In fact, in the past few years we have much better understood how to *approximately* solve basic graph problems, such as minimum cut, single-source shortest paths, all-pairs shortest paths, and maximum flows, on distributed networks (e.g. [60, 34, 33]). However, when it comes to solving these problems *exactly*, almost nothing is known. Understanding the complexity of exact algorithms is an important open problem.

We refer to [58, 59] for further open problems.

Appendix A. Proof of Lemma 3.

To prove (2), let π_i be a shortest path between u and v in G_i . Observe that if we consider this path in G (with the corresponding edge weights), then its total weight is at least the distance between u and v in G , i.e. $w(\pi_i, G) \geq d(u, v, G)$, because no path in G can have weight less than the shortest path in G . We therefore get

$$\begin{aligned} \rho_i \cdot d(u, v, G_i) &= \rho_i \cdot \sum_{(x,y) \in \pi_i} w(x, y, G_i) = \sum_{(x,y) \in \pi_i} \rho_i \cdot \left\lceil \frac{w(x, y, G)}{\rho_i} \right\rceil \\ &\geq \sum_{(x,y) \in \pi_i} w(x, y, G) = w(\pi_i, G) \geq d(u, v, G). \end{aligned}$$

To prove (4), let π be a shortest h -hop path from u to v in G . Observe that $w(\pi, G_i) \geq d(u, v, G_i)$, as again no path has smaller weight than the shortest path in G_i . By additionally exploiting the assumption $d^h(u, v, G) \geq 2^i$, we get

$$\begin{aligned} d(u, v, G_i) \cdot \rho_i &\leq w(\pi, G_i) \cdot \rho_i = \sum_{(x,y) \in \pi} w(x, y, G_i) \cdot \rho_i = \sum_{(x,y) \in \pi} \left\lceil \frac{w(x, y, G)}{\rho_i} \right\rceil \cdot \rho_i \\ &\leq \sum_{(x,y) \in \pi} (w(x, y, G) + \rho_i) = w(\pi, G) + |\pi| \cdot \rho_i = d^h(u, v, G) + |\pi| \cdot \rho_i \\ &\leq d^h(u, v, G) + h \cdot \rho_i = d^h(u, v, G) + \epsilon 2^i \leq d^h(u, v, G) + \epsilon d^h(u, v, G) \\ &= (1 + \epsilon) d^h(u, v, G). \end{aligned}$$

To prove (3), we combine (4) with the assumption $d^h(u, v, G) \leq 2^{i+1}$:

$$d(u, v, G_i) \leq \frac{(1 + \epsilon) d^h(u, v, G)}{\rho_i} = \frac{(1 + \epsilon) h d^h(u, v, G)}{\epsilon 2^i} \leq \frac{(1 + \epsilon) h 2^{i+1}}{\epsilon 2^i} = (1 + 2/\epsilon) h.$$

Appendix B. Proof of Lemma 10.

We prove the claim by the probabilistic method. Consider a sampling process that determines a set $T \subseteq U$ by adding each element of U to T independently with probability $1/(2x)$. Let E_0 denote the event that $|T| > |U|/x$ and for every $1 \leq i \leq k$ let E_i denote the event that $T \cap S_i = \emptyset$. First, observe that the size of T is $|U|/(2x)$ in expectation. By Markov's inequality, we can bound the probability that the size of T is at most twice the expectation by at least $1/2$ and thus $\Pr[E_0] = \Pr[|T| > |U|/x] \leq 1/2$. Furthermore, for every $1 \leq i \leq k$, the probability that S_i contains no node of T is

$$\Pr[E_i] = \left(1 - \frac{1}{2x}\right)^{|S_i|} \leq 1 - \left(1 - \frac{1}{2x}\right)^{2x \ln 3k} \leq \frac{1}{e^{\ln 3k}} = \frac{1}{3k}.$$

The set T fails to have the desired properties of a small hitting set if at least one of the events E_i occurs. By the union bound we have

$$\Pr \left[\bigcup_{0 \leq i \leq k} E_i \right] \leq \sum_{0 \leq i \leq k} \Pr[E_i] \leq \frac{1}{2} + k \cdot \frac{1}{3k} = \frac{1}{2} + \frac{1}{3} < 1.$$

It follows that the sampling process constructed a hitting set T for $\mathcal{C} = \{S_1, \dots, S_k\}$ of size at most $|T| \leq |U|/x$ with non-zero probability. Therefore a set T with these properties must really exist. This finishes the proof of [Lemma 10](#).

Appendix C. Ruling Set Algorithm.

For each node v , we represent its ID by a binary number $v_1 v_2 \dots v_\lambda$. Initially, we set $T_0 = U$. The algorithm proceeds for b iterations.

In the i^{th} iteration, we consider u_i for every node $u \in T_{i-1}$. If $u_i = 0$, v remains in T_i and sends a “beep” message to every node within distance $c - 1$. This takes $c - 1$ rounds as beep messages from different nodes can be combined. If $u_i = 1$, it remains in T_i if there is no node $v \in T_{i-1}$ such that $d(u, v, G) \leq c$ and $v_i = 0$; in other words, it remains in T_i if it does not hear any beep after $c - 1$ rounds. The output is $T = T_\lambda$. The running time of the above algorithm is clearly $O(c\lambda) = O(c \log n)$. Also, the distance between every pair of nodes in T is at least c since for every pair of nodes u and v of distance less than c , there is a i such that $u_i \neq v_i$, and in the i^{th} iteration if both u and v are in T_{i-1} , then one of them will send a beep and the other one will not be in T_i . Finally, it can be shown by induction that after the i^{th} round every node in U is at distance at most i from some node in T_i ; thus it follows that every node in U is at distance at most $c\lambda$ from some node in T .

Acknowledgments. We thank Michael Elkin, Stephan Friedrichs, and Christoph Lenzen for their comments and questions on the previous version of this paper. D. Nanongkai thanks Michael Elkin for bringing the notion of hop set to his attention. We thank the anonymous reviewers of SICOMP for their detailed comments.

REFERENCES

- [1] *List of open problems in sublinear algorithms: Problem 14*. <http://sublinear.info/14>.
- [2] A. ABBOD, G. BODWIN, AND S. PETTIE, *A hierarchy of lower bounds for sublinear additive spanners*, in Symposium on Discrete Algorithms (SODA), 2017, pp. 568–576.
- [3] G. AUSIELLO, A. D’ATRI, AND M. PROTASI, *Structure preserving reductions among convex optimization problems*, Journal of Computer and System Sciences, 21 (1980), pp. 136–153. Announced at ICALP’77.
- [4] L. BARENBOIM AND M. ELKIN, *Distributed Graph Coloring: Fundamentals and Recent Developments*, Synthesis Lectures on Distributed Computing Theory, Morgan & Claypool Publishers, 2013.
- [5] S. BASWANA, *Streaming algorithm for graph spanners—single pass and constant processing time per edge*, Information Processing Letters, 106 (2008), pp. 110–114.
- [6] S. BASWANA, R. HARIHARAN, AND S. SEN, *Improved decremental algorithms for maintaining transitive closure and all-pairs shortest paths*, Journal of Algorithms, 62 (2007), pp. 74–92. Announced at STOC’02.
- [7] R. BECKER, A. KARRENBauer, S. KRINNINGER, AND C. LENZEN, *Near-optimal approximate shortest paths and transshipment in distributed and streaming models*, in International Symposium on Distributed Computing (DISC), 2017. To appear.
- [8] R. BELLMAN, *On a routing problem*, Quarterly of Applied Mathematics, 16 (1958), pp. 87–90.
- [9] A. BERNSTEIN, *Fully dynamic $(2 + \epsilon)$ approximate all-pairs shortest paths with fast query and close to linear update time*, in Symposium on Foundations of Computer Science (FOCS), 2009, pp. 693–702.

- [10] A. BERNSTEIN, *Maintaining shortest paths under deletions in weighted directed graphs*, SIAM Journal on Computing, 45 (2016), pp. 548–574. Announced at STOC’13.
- [11] K. CENSOR-HILLEL, P. KASKI, J. H. KORHONEN, C. LENZEN, A. PAZ, AND J. SUOMELA, *Algebraic methods in the congested clique*, in Symposium on Principles of Distributed Computing (PODC), 2015, pp. 143–152.
- [12] E. COHEN, *Fast algorithms for constructing t -spanners and paths with stretch t* , SIAM Journal on Computing, 28 (1998), pp. 210–236. Announced at FOCS’93.
- [13] E. COHEN, *Polylog-time and near-linear work approximation scheme for undirected shortest paths*, Journal of the ACM, 47 (2000), pp. 132–166. Announced at STOC’94.
- [14] A. DAS SARMA, M. DINITZ, AND G. PANDURANGAN, *Efficient distributed computation of distance sketches in networks*, Distributed Computing, 28 (2015), pp. 309–320. Announced at SPAA’12.
- [15] A. DAS SARMA, S. HOLZER, L. KOR, A. KORMAN, D. NANONGKAI, G. PANDURANGAN, D. PELEG, AND R. WATTENHOFER, *Distributed verification and hardness of distributed approximation*, SIAM Journal on Computing, 41 (2012), pp. 1235–1265. Announced at STOC’11.
- [16] C. DEMETRESCU, I. FINOCCHI, AND G. ITALIANO, *Handbook on Data Structures and Applications*, CRC Press, 2005, ch. 36: Dynamic Graphs.
- [17] C. DEMETRESCU, I. FINOCCHI, AND A. RIBICHINI, *Trading off space for passes in graph streaming problems*, ACM Transactions on Algorithms, 6 (2009). Announced at SODA’06.
- [18] C. DEMETRESCU AND G. F. ITALIANO, *Fully dynamic all pairs shortest paths with real edge weights*, Journal of Computer and System Sciences, 72 (2006), pp. 813–837. Announced at FOCS’01.
- [19] D. DOLEV, C. LENZEN, AND S. PELED, *“Tri, Tri again”: Finding triangles and small subgraphs in a distributed setting*, in International Symposium on Distributed Computing (DISC), 2012, pp. 195–209.
- [20] M. ELKIN, *Distributed approximation: a survey*, SIGACT News, 35 (2004), pp. 40–57.
- [21] M. ELKIN, *An unconditional lower bound on the time-approximation trade-off for the distributed minimum spanning tree problem*, SIAM Journal on Computing, 36 (2006), pp. 433–456. Announced at STOC’04.
- [22] M. ELKIN, *Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners*, ACM Transactions on Algorithms, 7 (2011), pp. 20:1–20:17. Announced at ICALP’07.
- [23] M. ELKIN, *Distributed exact shortest paths in sublinear time*, in Symposium on Theory of Computing (STOC), 2017, pp. 757–770.
- [24] M. ELKIN, H. KLAUCK, D. NANONGKAI, AND G. PANDURANGAN, *Can quantum communication speed up distributed computation?*, in Symposium on Principles of Distributed Computing (PODC), 2014, pp. 166–175.
- [25] M. ELKIN AND O. NEIMAN, *Hopsets with constant hopbound, and applications to approximate shortest paths*, in Symposium on Foundations of Computer Science (FOCS), 2016.
- [26] M. ELKIN AND O. NEIMAN, *On efficient distributed construction of near optimal routing schemes: Extended abstract*, in Symposium on Principles of Distributed Computing (PODC), 2016, pp. 235–244.
- [27] M. ELKIN AND J. ZHANG, *Efficient algorithms for constructing $(1 + \epsilon, \beta)$ -spanners in the distributed and streaming models*, Distributed Computing, 18 (2006), pp. 375–385.
- [28] J. FEIGENBAUM, S. KANNAN, A. MCGREGOR, S. SURI, AND J. ZHANG, *On graph problems in a semi-streaming model*, Theoretical Computer Science, 348 (2005), pp. 207–216. Announced at ICALP’04.
- [29] J. FEIGENBAUM, S. KANNAN, A. MCGREGOR, S. SURI, AND J. ZHANG, *Graph distances in the data-stream model*, SIAM Journal on Computing, 38 (2008), pp. 1709–1727. Announced at SODA’05.
- [30] L. R. FORD, *Network flow theory*, Tech. Report P-923, The Rand Corporation, 1956.
- [31] S. FRIEDRICHS AND C. LENZEN, *Parallel metric tree embedding based on an algebraic view on moore-bellman-ford*, in Symposium on Parallelism in Algorithms and Architectures (SPAA), 2016, pp. 455–466.
- [32] J. A. GARAY, S. KUTTEN, AND D. PELEG, *A sublinear time distributed algorithm for minimum-weight spanning trees*, SIAM Journal on Computing, 27 (1998), pp. 302–316. Announced at FOCS’93.
- [33] M. GHAFFARI, A. KARRENBauer, F. KUHN, C. LENZEN, AND B. PATT-SHAMIR, *Near-optimal distributed maximum flow: Extended abstract*, in Symposium on Principles of Distributed Computing (PODC), 2015, pp. 81–90.
- [34] M. GHAFFARI AND F. KUHN, *Distributed minimum cut approximation*, in International Sympo-

- sium on Distributed Computing (DISC), 2013, pp. 1–15.
- [35] M. GHAFFARI AND R. UDWANI, *Brief announcement: Distributed single-source reachability*, in Symposium on Principles of Distributed Computing (PODC), 2015, pp. 163–165.
 - [36] A. V. GOLDBERG, S. A. PLOTKIN, AND G. E. SHANNON, *Parallel symmetry-breaking in sparse graphs*, SIAM Journal on Discrete Mathematics, 1 (1988), pp. 434–446. Announced at STOC’87.
 - [37] V. GURUSWAMI AND K. ONAK, *Superlinear lower bounds for multipass graph processing*, Algorithmica, 76 (2016), pp. 654–683. Announced at CCC’13.
 - [38] M. HENZINGER AND V. KING, *Fully dynamic biconnectivity and transitive closure*, in Symposium on Foundations of Computer Science (FOCS), 1995, pp. 664–672.
 - [39] M. HENZINGER, S. KRINNINGER, AND D. NANONGKAI, *Decremental single-source shortest paths on undirected graphs in near-linear total update time*, in Symposium on Foundations of Computer Science (FOCS), 2014, pp. 146–155.
 - [40] M. HENZINGER, S. KRINNINGER, AND D. NANONGKAI, *A deterministic almost-tight distributed algorithm for approximating single-source shortest paths*, in Symposium on Theory of Computing, (STOC), 2016, pp. 489–498.
 - [41] S. HOLZER AND N. PINSKER, *Approximation of distances and shortest paths in the broadcast congest clique*, in International Conference on Principles of Distributed Systems (OPODIS), 2015.
 - [42] S. HOLZER AND R. WATTENHOFER, *Optimal distributed all pairs shortest paths and applications*, in Symposium on Principles of Distributed Computing (PODC), 2012, pp. 355–364.
 - [43] D. S. JOHNSON, *Approximation algorithms for combinatorial problems*, Journal of Computer and System Sciences, 9 (1974), pp. 256–278. Announced at STOC’73.
 - [44] M. KHAN, F. KUHN, D. MALKHI, G. PANDURANGAN, AND K. TALWAR, *Efficient distributed approximation algorithms via probabilistic tree embeddings*, Distributed Computing, 25 (2012), pp. 189–205. Announced at PODC 2008.
 - [45] M. KHAN AND G. PANDURANGAN, *A fast distributed approximation algorithm for minimum spanning trees*, Distributed Computing, 20 (2008), pp. 391–402. Announced at DISC’06.
 - [46] V. KING, *Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs*, in Symposium on Foundations of Computer Science (FOCS), 1999, pp. 81–91.
 - [47] P. N. KLEIN AND S. SUBRAMANIAN, *A randomized parallel algorithm for single-source shortest paths*, Journal of Algorithms, 25 (1997), pp. 205–220.
 - [48] L. KOR, A. KORMAN, AND D. PELEG, *Tight bounds for distributed minimum-weight spanning tree verification*, Theory of Computing Systems, 53 (2013), pp. 318–340. Announced at STACS’11.
 - [49] S. KUTTEN AND D. PELEG, *Fast distributed construction of small k -dominating sets and applications*, Journal of Algorithms, 28 (1998), pp. 40–66. Announced at PODC’95.
 - [50] F. LE GALL, *Further algebraic algorithms in the congested clique model and applications to graph-theoretic problems*, in International Symposium on Distributed Computing (DISC), 2016, pp. 57–70.
 - [51] C. LENZEN, *Optimal deterministic routing and sorting on the congested clique*, in Symposium on Principles of Distributed Computing (PODC), 2013, pp. 42–50.
 - [52] C. LENZEN AND B. PATT-SHAMIR, *Fast routing table construction using small messages*, in Symposium on Theory of Computing (STOC), 2013, pp. 381–390.
 - [53] C. LENZEN AND B. PATT-SHAMIR, *Fast partial distance estimation and applications*, in Symposium on Principles of Distributed Computing (PODC), 2015, pp. 153–162.
 - [54] C. LENZEN AND D. PELEG, *Efficient distributed source detection with limited bandwidth*, in Symposium on Principles of Distributed Computing (PODC), 2013, pp. 375–382.
 - [55] Z. LOTKER, B. PATT-SHAMIR, AND A. ROSÉN, *Distributed approximate matching*, SIAM Journal on Comput., 39 (2009), pp. 445–460. Announced at PODC’07.
 - [56] A. MCGREGOR, *Graph stream algorithms: A survey*, SIGMOD Record, 43 (2014), pp. 9–20.
 - [57] A. MADRY, *Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms*, in Symposium on Theory of Computing (STOC), 2010, pp. 121–130.
 - [58] D. NANONGKAI, *Distributed approximation algorithms for weighted shortest paths*, in Symposium on Theory of Computing (STOC), 2014, pp. 565–573.
 - [59] D. NANONGKAI, *Some challenges on distributed shortest paths problems, a survey*, in SIROCCO, 2016. Invited talk.
 - [60] D. NANONGKAI AND H. SU, *Almost-tight distributed minimum cut algorithms*, in International Symposium on Distributed Computing (DISC), 2014, pp. 439–453.
 - [61] D. PELEG, *Distributed Computing: A Locality-Sensitive Approach*, Society for Industrial and

- Applied Mathematics, Philadelphia, PA, USA, 2000.
- [62] D. PELEG AND V. RUBINOVICH, *A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction*, SIAM Journal on Computing, 30 (2000), pp. 1427–1442. Announced at FOCS'99.
 - [63] D. PRITCHARD AND R. THURIMELLA, *Fast computation of small cuts via cycle space sampling*, ACM Transactions on Algorithms, 7 (2011), pp. 46:1–46:30. Announced at ICALP'08.
 - [64] L. RODITY, M. THORUP, AND U. ZWICK, *Deterministic constructions of approximate distance oracles and spanners*, in International Colloquium on Automata, Languages and Programming (ICALP), 2005, pp. 261–272.
 - [65] L. RODITY AND U. ZWICK, *On dynamic shortest paths problems*, Algorithmica, 61 (2011), pp. 389–401. Announced at ESA'04.
 - [66] P. SANKOWSKI, *Subquadratic algorithm for dynamic shortest distances*, in International Computing and Combinatorics Conference (COCOON), 2005, pp. 461–470.
 - [67] C. SOMMER, *Shortest-path queries in static networks*, ACM Computing Surveys, 46 (2014), pp. 45:1–45:31.
 - [68] M. THORUP AND U. ZWICK, *Approximate distance oracles*, Journal of the ACM, 52 (2005), pp. 74–92. Announced at STOC'01.
 - [69] M. THORUP AND U. ZWICK, *Spanners and emulators with sublinear distance errors*, in SODA, 2006, pp. 802–809.
 - [70] R. THURIMELLA, *Sub-linear distributed algorithms for sparse certificates and biconnected components*, Journal of Algorithms, 23 (1997), pp. 160–179. Announced at PODC'95.
 - [71] L. TSENG, *PODC 2015 review*, SIGACT News, 46 (2015), pp. 94–102.
 - [72] J. D. ULLMAN AND M. YANNAKAKIS, *High-probability parallel transitive-closure algorithms*, SIAM Journal on Computing, 20 (1991), pp. 100–125. Announced at SPAA'90.
 - [73] U. ZWICK, *All pairs shortest paths using bridging sets and rectangular matrix multiplication*, Journal of the ACM, 49 (2002), pp. 289–317. Announced at FOCS'98.