



# Fast approximate shortest paths in the congested clique

Keren Censor-Hillel<sup>1</sup> · Michal Dory<sup>1</sup> · Janne H. Korhonen<sup>2</sup> · Dean Leitersdorf<sup>1</sup>

Received: 31 October 2019 / Accepted: 25 April 2020 / Published online: 28 May 2020  
© The Author(s) 2020

## Abstract

We design fast *deterministic* algorithms for distance computation in the CONGESTED CLIQUE model. Our key contributions include:

- A  $(2 + \epsilon)$ -approximation for all-pairs shortest paths in  $O(\log^2 n/\epsilon)$  rounds on unweighted undirected graphs. With a small additional additive factor, this also applies for weighted graphs. This is the first *sub-polynomial* constant-factor approximation for APSP in this model.
- A  $(1 + \epsilon)$ -approximation for multi-source shortest paths from  $O(\sqrt{n})$  sources in  $O(\log^2 n/\epsilon)$  rounds on weighted undirected graphs. This is the first *sub-polynomial* algorithm obtaining this approximation for a set of sources of polynomial size.

Our main techniques are new distance tools that are obtained via improved algorithms for sparse matrix multiplication, which we leverage to construct efficient hopsets and shortest paths. Furthermore, our techniques extend to additional distance problems for which we improve upon the state-of-the-art, including diameter approximation, and an *exact* single-source shortest paths algorithm for weighted undirected graphs in  $\tilde{O}(n^{1/6})$  rounds.

**Keywords** Distributed computing · Congested clique model · All-pairs shortest paths · Matrix multiplication · Hopsets

## 1 Introduction

Computing distances in a graph is a fundamental task widely studied in many computational settings. Notable examples are computation of all-pairs shortest paths (APSP), single-source shortest paths (SSSP), and computing specific parameters such as the diameter of a graph. In this work, we study distance computations in the CONGESTED CLIQUE model of distributed computing.

In the CONGESTED CLIQUE model, we have a *fully-connected* communication network of  $n$  nodes, where nodes communicate by sending  $O(\log n)$ -bit messages to each other node in synchronous rounds. The CONGESTED CLIQUE model has been receiving much attention during the past decade or so, due to both its theoretical interest in focusing on congestion alone as a communication resource, and its relation to practical settings that use fully connected overlays [13,14,21,28,29,31–33,36,40–43,47,50–52]. In particular, there have been many recent papers studying distance problems in CONGESTED CLIQUE [8,13,14,24,36,42,48,52].

### 1.1 Distance computation in the congested clique

Many state-of-the-art results for distance computations in the CONGESTED CLIQUE model exploit the well-known connection between computing distances and matrix multiplication [13,14,42]. Specifically, the  $n$ th power of the adjacency matrix  $A$  of a graph  $G = (V, E)$ , taken over the *min-plus* or *tropical* semiring (see e.g. [13] for details), correspond to shortest-path distances. Hence, iteratively squaring

✉ Janne H. Korhonen  
janne.korhonen@ist.ac.at

Keren Censor-Hillel  
ckeren@cs.technion.ac.il

Michal Dory  
smichald@cs.technion.ac.il

Dean Leitersdorf  
dean.leitersdorf@gmail.com

<sup>1</sup> Department of Computer Science, Technion, Haifa, Israel

<sup>2</sup> IST Austria, Klosterneuburg, Austria

a matrix  $\log n$  times allows computing all the distances in the graph. This approach gives the best known algorithms for APSP in the CONGESTED CLIQUE, including (1) an  $\tilde{O}(n^{1/3})$  round algorithm for exact APSP in weighted directed graphs [13], (2)  $O(n^{0.158})$  round algorithms for exact APSP in unweighted undirected graphs and  $(1+o(1))$ -approximate APSP in weighted directed graphs [13], as well as (3) an  $O(n^{0.2096})$  round algorithm for exact APSP in directed graphs with constant weights [42]. Additionally, in [14], this connection is used to show an improved APSP algorithm for *sparse* graphs.

For approximating the distances, faster approximations for larger constants can be obtained by computing a  $k$ -spanner, which is a sparse graph that preserves distances up to a multiplicative factor of  $k$ , and having all nodes learn the entire spanner. Using the CONGESTED CLIQUE spanner constructions of [52], this approach gives a  $(2k-1)$ -approximation for APSP in  $\tilde{O}(n^{1/k})$  rounds, which is still *polynomial* for any constant  $k$ .

This raises the following fundamental question:

**Question 1** Can we obtain constant-factor approximations for APSP in *sub-polynomial* time?

If we restrict our attention to SSSP, sub-polynomial  $(1+\epsilon)$ -approximation is indeed possible [8,34]; in particular, the state of the art is a gradient-descent-based algorithm that obtains a  $(1+\epsilon)$ -approximation in  $O(\epsilon^{-3}\text{polylog } n)$  rounds even in the more restricted broadcast version of the CONGESTED CLIQUE model [8]. However, these algorithms from prior work provide distances only from a single source.

## 1.2 Our contributions

*All-pairs shortest paths* As our first main result, we address the above fundamental question by providing the first *polylogarithmic* constant approximations for APSP in the CONGESTED CLIQUE model. Specifically, we show the following.

**Theorem 1** *There is a deterministic  $(2+\epsilon)$ -approximation algorithm for unweighted undirected APSP in the CONGESTED CLIQUE model that takes  $O(\frac{\log^2 n}{\epsilon})$  rounds.*

We also obtain a nearly  $(2+\epsilon)$ -approximation in  $O(\frac{\log^2 n}{\epsilon})$  rounds in *weighted* undirected graphs, in the sense that for any distance estimate  $d(u, v)$ , there is further additive  $(1+\epsilon)w_{uv}$  error in the approximation, where  $w_{uv}$  is the weight of the heaviest edge on the shortest  $u$ - $v$  path.

Our approximation is almost tight for *sub-polynomial* algorithms in the following sense. As noted by [41], a  $(2-\epsilon)$ -approximate APSP in unweighted undirected graphs is essentially equivalent to fast matrix multiplication, so obtaining a better approximation in complexity below  $O(n^{0.158})$  would result in a faster algorithm for matrix multiplication in

the CONGESTED CLIQUE. Likewise, a sub-polynomial-time algorithm with *any* approximation ratio for directed graphs would give a faster matrix multiplication algorithm [20], so our results will likely not extend to directed graphs.

*Multi-source shortest paths* As our second main result, we show a fast  $(1+\epsilon)$ -approximation algorithm for the multi-source shortest paths problem (MSSP), which is polylogarithmic as long as the number of sources is  $\tilde{O}(\sqrt{n})$ . Specifically, we show the following.

**Theorem 2** *There is a deterministic  $(1+\epsilon)$ -approximation algorithm for the weighted undirected MSSP that takes*

$$O\left(\left(\frac{|S|^{2/3}}{n^{1/3}} + \log n\right) \cdot \frac{\log n}{\epsilon}\right)$$

*rounds in the CONGESTED CLIQUE, where  $S$  is the set of sources. In particular, the complexity is  $O(\frac{\log^2 n}{\epsilon})$  as long as  $|S| \leq O(\sqrt{n} \cdot (\log n)^{3/2})$ .*

This is the first sub-polynomial algorithm that obtains such approximations for a set  $S$  of polynomial size. Other advantages of our approach, compared to the previous  $(1+\epsilon)$ -approximation SSSP algorithm [8], is that it is based on simple combinatorial techniques. In addition, our complexity improves upon the complexity of [8].

*Exact SSSP and diameter approximation* In addition to the above, our techniques allow us to obtain a near  $3/2$ -approximation for the diameter in  $O(\frac{\log^2 n}{\epsilon})$  rounds as well as an  $\tilde{O}(n^{1/6})$ -round algorithm for *exact* weighted SSSP, improving the previous  $\tilde{O}(n^{1/3})$ -round algorithm [13]. All our algorithms are deterministic.

## 1.3 Our techniques

The main technical tools we develop for our distance computation algorithms are a new *sparse matrix multiplication* algorithm, extending the recent result of [14], and a new deterministic *hopset* construction algorithm for the CONGESTED CLIQUE.

*Distance products* We start from the basic idea of using matrix multiplication to compute distances in graphs. Specifically, if  $A$  is the weighted adjacency matrix of a graph  $G$ , it is well known that distances in  $G$  can be computed by iterating the *distance product*  $A \star A$ , defined as

$$(A \star A)[i, j] = \min_k (A[i, k] + A[k, j]),$$

that is, the matrix multiplication over the min-plus semiring.

A simple idea is to apply the recent sparse matrix multiplication algorithm of [14], with running time that depends on the density of the input matrices. In particular, this allows

us to multiply two sparse matrices with  $O(n^{3/2})$  non-zero entries in  $O(1)$  rounds; note that for the distance product, the zero element is  $\infty$ . However, using this algorithm for computing distances directly is inefficient, as  $A \star A$  can be dense even if  $A$  is sparse (e.g. a star graph), and hence iterative squaring is not guaranteed to be efficient. Moreover, our goal is to compute distances in general graphs, not only in sparse graphs. Nevertheless, we show that while using [14] directly may not be efficient, we can use sparse matrix multiplication as a basic building block for distance computation in the CONGESTED CLIQUE.

*Our distance tools* The key observation is that many building blocks for distance computation are actually based on computations in sparse graphs or only consider a limited number of nodes. Concrete examples of such tasks include:

- $k$ - NEAREST *problem* Compute distances for each node to the  $k$  closest nodes in the graph.
- $(S, d, k)$ - SOURCE DETECTION *problem* Given a set of sources  $S$ , compute the distances for each node to the  $k$  nearest sources using paths of at most  $d$  hops.
- DISTANCE THROUGH SETS *problem* Given a set of nodes  $S$  and distances to all nodes in  $S$ , compute the distances between all nodes using paths through nodes in  $S$ .

For all of these problems, there is a degree of sparsity we can hope to exploit if  $k$  or  $|S|$  are small enough. For example, the  $(S, d, k)$ - SOURCE DETECTION problem, requires the multiplication of a dense adjacency matrix and a possibly sparse matrix, depending on the size of  $S$ . However, for any  $S$  of polynomial size the algorithm in [14] is polynomial. An interesting property in this problem, though, is that the output matrix is also sparse. If we look at the  $k$ - NEAREST problem, both input matrices are sparse, hence we can use the previous sparse matrix multiplication algorithm. However, this does not exploit the sparsity of this problem to the end: in this problem we are interested only in computing the  $k$  nearest nodes to each node, hence there is no need to compute the full output matrix. The challenge in this case is that we do not know the identity of the  $k$  closest nodes before the computation. To exploit this sparsity we design new matrix multiplication algorithms, that in particular have the ability to sparsify the matrix throughout the computation, and get a complexity that depends only on the size of the output we are interested in.

*Sparse matrix multiplication* To compute the above, we design new sparse matrix multiplication algorithms, which differ from [14] by taking into account also the sparsity of the *output matrix*. For matrix  $M$ , let  $\rho_M$  denote the *density* of  $M$ , that is, the average number of non-zero entries on a row. Specifically, for distance product computation  $P = S \star T$ , we obtain two variants:

- One variant assumes that the sparsity of the output matrix is known.
- The other *sparsifies* the output matrix on the fly, keeping only the  $\rho_P$  smallest entries for each row.

For these two scenarios, we obtain running times of

$$O\left(\frac{(\rho_S \rho_T \rho_P)^{1/3}}{n^{2/3}} + 1\right)$$

and

$$O\left(\frac{(\rho_S \rho_T \rho_P)^{1/3}}{n^{2/3}} + \log n\right)$$

rounds, respectively, improving the running time of the prior sparse matrix multiplication for  $\rho_P = o(n)$ .

This allows us to obtain faster distance tools, by taking into account the sparsity of the output:

- We can solve the  $k$ - NEAREST problem in

$$O\left(\left(\frac{k}{n^{2/3}} + \log n\right) \log k\right)$$

rounds.

- We can solve the  $(S, d, k)$ - SOURCE DETECTION problem in

$$O\left(\left(\frac{m^{1/3} |S|^{2/3}}{n} + 1\right) d\right)$$

rounds, where  $m$  is the number of edges in the input graph; note that dependence on  $d$  becomes linear in order to exploit the sparsity.

In concrete terms, with these output-sensitive distance tools we still get subpolynomial running times even when the parameters are polynomial. For example, we can get the distances to the  $\tilde{O}(n^{2/3})$  closest nodes in  $\tilde{O}(1)$  rounds. Note that though our final results are only for undirected graphs, these distance tools work for directed weighted graphs.

*Hopsets* An issue with our  $(S, d, k)$ - SOURCE DETECTION algorithm is that in order to exploit the sparsity of the matrices, we must perform  $d$  multiplications to learn the distances of all the nodes at hop-distance at most  $d$  from  $S$ . Hence, to learn the distances of all the nodes from  $S$ , we need to do  $n$  multiplications, which is no longer efficient. To overcome this challenge, we use *hopsets*, which are a central building block in many distance computations in the distributed setting [22–24,34,48]. A  $(\beta, \epsilon)$ -hopset  $H$  is a sparse graph such that the  $\beta$ -hop distances in  $G \cup H$  give  $(1+\epsilon)$ -approximations of the distances in  $G$ . Since it is enough to look only at  $\beta$ -hop distances in  $G \cup H$ , using our source detection algorithm

together with a hopset allows getting an efficient algorithm for approximating distances, as long as  $\beta$  is small enough.

However, the time complexity of all current hopset constructions depends on the size of the hopset [23,24,34], in the following way. The complexity of building a hopset of size  $n\rho$  is at least  $\Omega(\rho)$ . This is a major obstacle for efficient shortest paths algorithms, since based on recent existential results there are no hopsets where both  $\beta$  and  $\rho$  are polylogarithmic [1] (see Sect. 1.4). Nevertheless, we show that our new distance tools allow to build hopsets in a time that *does not* depend on the hopset size. In particular, we show how to implement a variant of the recent hopset construction of Elkin and Neiman [24] in  $O(\frac{\log^2 n}{\epsilon})$  rounds. The size of our hopset is  $\tilde{O}(n^{3/2})$ , hence constructing it using previous algorithms requires at least  $\tilde{\Omega}(\sqrt{n})$  rounds.

*Applying the distance tools* As a direct application of our source detection and hopset algorithms, we obtain a multi-source shortest paths (MSSP) algorithm, allowing to compute  $(1 + \epsilon)$ -approximate distances to  $\tilde{O}(\sqrt{n})$  sources in polylogarithmic time. Our MSSP algorithm, in turn, forms the basis of a near  $3/2$ -approximation for the diameter in unweighted graphs, and a  $(3 + \epsilon)$ -approximation for weighted APSP. To obtain a  $(2 + \epsilon)$ -approximation for unweighted APSP, the high-level idea is to deal separately with paths that contain a high-degree node and paths with only low-degree nodes. A crucial ingredient in the algorithm is showing that in sparser graphs, we can actually compute distances to a larger set of sources  $S$  efficiently, which is useful for obtaining a better approximation. Our exact SSSP algorithm uses our algorithm for finding distances to the  $k$ -nearest nodes, which allows constructing efficiently the  $k$ -shortcut graph described in [22,48].

## 1.4 Additional related work

*Distance computation in the congested clique* APSP and SSSP are fundamental problems that are studied extensively in various computation models. Apart from the matrix multiplication based algorithms in the CONGESTED CLIQUE [13,14,42], previous results include also  $\tilde{O}(\sqrt{n})$ -round algorithms for exact SSSP and  $(2 + o(1))$ -approximation for APSP [48]. Other distance problems studied in the CONGESTED CLIQUE are construction of hopsets [23,24,34] and spanners [52].

*Matrix multiplication in the congested clique* As shown by [13], matrix multiplication can be done in CONGESTED CLIQUE in  $O(n^{1/3})$  rounds over semirings, and in  $O(n^{1-2/\omega})$  rounds over rings, where  $\omega < 2.3728639$  is the exponent of the matrix multiplication [27]. For rectangular matrix multiplication, [42] gives faster algorithms. The first sparse matrix multiplication algorithms for CONGESTED CLIQUE were given by [14], as discussed above.

*Distance computation in the CONGEST model* The distributed CONGEST model is identical to the CONGESTED CLIQUE model, with the difference that the communication network is identical to the input graph  $G$ , and nodes can communicate only with their neighbours in each round. Distance computation is extensively studied in the CONGEST model. The study of *exact* APSP in weighted graphs has been the focus of many recent papers [3,12,22,38], culminating in a near tight  $\tilde{O}(n)$  algorithm [12]. Such results were previously known in unweighted graphs [37,46,53] or for approximation algorithms [45,48]. Approximate and exact algorithms for SSSP are studied in [8,22,26,30,34,44,48]. While near-tight algorithms exist for approximating SSSP [8], there is currently a lot of interest in understanding the complexity of *exact* SSSP and *directed* SSSP [22,26,30]. The *source detection* problem is studied in [46], demonstrating the applicability of this tool for many distance problems such as APSP and diameter approximation in unweighted graphs. An extension for the weighted case is studied in [45]. Algorithms and lower bounds for approximating the diameter are studied in [2,37,42,46,53].

*Distance computation in the sequential setting* Among the rich line of research in the sequential setting, we focus only on the most related to our work. The pioneering work of Aingworth et al. [4], inspired much research on approximate APSP [6,7,10,18,20] and approximate diameter [5,10,16,54], with the goal of understanding the tradeoffs between the time complexity and approximation ratio. Many of these papers use clustering ideas and hitting set arguments as the basis of their algorithms, and our approximate APSP and diameter algorithms are inspired by such ideas.

*Hopsets* Hopsets are a central building block in distance computation and are studied extensively in various computing models [11,17,22–24,34,35,48,55]. The most related to our work are two recent constructions of Elkin and Neiman [24], and Huang and Pettie [39], which are based on the emulators of Thorup and Zwick [56], and are near optimal due to existential results [1]. Specifically, in [39], there is a construction of  $(\beta, \epsilon)$ -hopsets of size  $O(n^{1+\frac{1}{2k+1}})$  with  $\beta = O(k/\epsilon)^k$ , where recent existential results show that any construction of  $(\beta, \epsilon)$ -hopsets with worst case size  $n^{1+\frac{1}{2k-1}-\delta}$  must have  $\beta = \Omega_k((\frac{1}{\epsilon})^k)$ , where  $k \geq 1$  is an integer and  $\delta > 0$ . For a detailed discussion of hopsets see the introduction in [23,24].

## 1.5 Preliminaries

*Notations* Except when specified otherwise, we assume our graphs are undirected with non-negative integer edge weights at most  $O(n^c)$  for a constant  $c$ . Given a graph  $G = (V, E)$  and  $u, v \in V$ , we denote by  $d_G(u, v)$  the distance between  $u$  and  $v$  in  $G$ , and by  $d_G^\beta(u, v)$  the length of the shortest path of

hop-distance at most  $\beta$  between  $u$  and  $v$  in  $G$ . If  $G$  is clear from the context, we use the notation  $d(u, v)$  for  $d_G(u, v)$ .

*Routing and sorting* As basic primitives, we use standard routing and sorting algorithms for the CONGESTED CLIQUE model. In the *routing* task, each node holds up to  $n$  messages of  $O(\log n)$  bits, and we assume that each node is also the recipient of at most  $n$  messages. In the *sorting* task, each node has a list of  $n$  entries from an ordered set, and we want to sort these entries so that node  $i$  receives the  $i$ th batch of entries according to the global order of all the input entries. Both of these task can be solved in  $O(1)$  rounds [19,43].

*Semirings and matrices* We assume we are operating over a semiring  $(R, +, \cdot, 0, 1)$ , where 0 is the identity element for addition and 1 is the identity element for multiplication. Note that we do not require the multiplication to be commutative. For the CONGESTED CLIQUE algorithms, we generally assume that the semiring elements can be encoded in messages of  $O(\log n)$  bits.

All matrices are assumed to be over the semiring  $R$ . For convenience, we identify  $[n]$  with the node set  $V$ , and use set  $V$  to index the matrix entries. For matrix  $S$ , we denote the matrix entry at position  $(v, u)$  by  $S[v, u]$ . For sets  $U, W \subseteq V$ , we denote by  $S[U, W]$  the submatrix obtained by taking rows and columns restricted to  $U$  and  $W$ , respectively.

*Hitting sets* Let  $S_v \subseteq V$  be a set of size at least  $k$ . We say that  $A \subseteq V$  is a *hitting set* of  $\{S_v\}_{v \in V}$  if in each subset  $S_v$  there is a node from  $A$ . We can construct hitting sets easily by adding each node to  $A$  with probability  $p = \frac{\log n}{k}$ . This gives a hitting set of expected size  $O(\frac{n \log n}{k})$ , such that w.h.p there is a node from  $A$  in each subset  $S_v$ . The same parameters are obtained by a recent *deterministic* construction of hitting sets in the CONGESTED CLIQUE [52], which gives the following (see Corollary 17 in [52]).

**Lemma 1** *Let  $\{S_v \subseteq V\}_{v \in V}$  be a set of subsets of size at least  $k$ , such that  $S_v$  is known to  $v$ . There exists a deterministic algorithm in the CONGESTED CLIQUE model that constructs a hitting set of size  $O(n \log n/k)$  in  $O((\log \log n)^3)$  rounds.*

*Partitions* We will use the following lemmas on partitioning a set of weighted items into equal parts. Note that all the lemmas are constructive, that is, they also imply a deterministic algorithm for constructing the partition in a canonical way.

**Lemma 2** ([14]) *Let  $w_1, w_2, \dots, w_n$  be natural numbers, and let  $W, w$  and  $k$  be natural numbers such that  $\sum_{i=1}^n w_i = W, w_i \leq w$  for all  $i$ , and  $k$  divides  $n$ . Then there is a partition of  $[n]$  into  $k$  sets  $I_1, I_2, \dots, I_k$  of size  $n/k$  such that*

$$\sum_{i \in I_j} w_i \leq W/k + w \text{ for all } j.$$

**Lemma 3** *Let  $w_1, w_2, \dots, w_n$  be natural numbers, and let  $W, w$  and  $k$  be natural numbers such that  $\sum_{i=1}^n w_i = W, w_i \leq w$  for all  $i$ . Then there is a partition of  $[n]$  into  $k$  sets  $I_1, I_2, \dots, I_k$  such that for each  $j$ , the set  $I_j$  consist of consecutive elements, and*

$$\sum_{i \in I_j} w_i \leq W/k + w.$$

**Proof** Construct sets  $I_j$  in the partition by starting from the first element, adding new elements to the current set until  $\sum_{i \in I_j} w_i \geq W/k$ . Since all weights  $w_i$  are at most  $w$ , we have  $\sum_{i \in I_j} w_i \leq W/k + w$  for all  $j$ , and since  $\sum_{i \in I_j} w_i \geq W/k$ , this process generates at most  $k$  sets.  $\square$

**Lemma 4** *Let  $w_1, w_2, \dots, w_n$  and  $u_1, u_2, \dots, u_n$  be natural numbers, and let  $W, U, w, u$  and  $k$  be natural numbers such that  $\sum_{i=1}^n w_i = W$  and  $\sum_{i=1}^n u_i = U$ , and we have  $w_i \leq w$  and  $u_i \leq u$  and for all  $i$ . Then there is a partition of  $[n]$  into  $k$  sets  $I_1, I_2, \dots, I_k$  such that for each  $j$ , the set  $I_j$  consist of consecutive elements, and*

$$\sum_{i \in I_j} w_i \leq 2(W/k + w) \quad \text{and} \quad \sum_{i \in I_j} u_i \leq 2(U/k + u).$$

**Proof** We begin by applying Lemma 3 separately to sequences  $w_1, w_2, \dots, w_n$  and  $u_1, u_2, \dots, u_n$ . That is, by Lemma 3, there exist the following two partitions of set  $[n]$  into sets of consecutive indices:

- Partition  $J_1, \dots, J_k$  such that for any  $j$ , we have  $\sum_{i \in J_j} w_i \leq W/k + w$ .
- Partition  $K_1, \dots, K_k$  such that for any  $j$ , we have  $\sum_{i \in K_j} u_i \leq U/k + u$ .

Let  $F$  be a multiset consisting of the last elements of sets from partitions  $J_1, \dots, J_k$  and  $K_1, \dots, K_k$ . Define a sequence  $i_0, i_1, \dots, i_{2k}$  by setting  $i_0 = 0$ , and letting  $i_1, \dots, i_{2k}$  be the elements of  $F$  in increasing order; note that the sequence may contain duplicates.

Define sets  $I_1, \dots, I_k$  as  $I_j = \{i_{2j-2} + 1, \dots, i_{2j}\}$ ; intuitively, this corresponds to taking the *fenceposts* of the partitions  $J_1, \dots, J_k$  and  $K_1, \dots, K_k$ , and taking every other fencepost to give a new partition into sets of consecutive indices. Clearly  $\{I_j\}$  form a partition of  $[n]$ , and since each  $I_j$  overlaps at most two sets from  $J_1, \dots, J_k$  and at most two sets from  $K_1, \dots, K_k$ , we have  $\sum_{i \in I_j} w_i \leq 2(W/k + w)$  and  $\sum_{i \in I_j} u_i \leq 2(U/k + u)$ .  $\square$

## 2 Matrix multiplication

### 2.1 Output-sensitive sparse matrix multiplication

Our first matrix multiplication result is a output-sensitive variant of sparse matrix multiplication. In the *matrix multiplication problem*, we are given two  $n \times n$  matrices  $S$  and  $T$  over semiring  $R$ , and the task is to compute the product matrix  $P = ST$ ,

$$P[u, v] = \sum_{w \in V} S[v, w]T[w, u].$$

Following [13,42], we assume for concreteness that in the CONGESTED CLIQUE model, each node receives the row  $S[v, V]$  and the column  $T[V, v]$  as a local input, and we want to compute the output so that each node locally knows the row  $P[v, V]$  of the output matrix.

*Matrix densities* For matrix  $S$ , we denote by  $\text{nz}(S)$  the number of non-zero entries in  $S$ . Furthermore, we define the *density*  $\rho_S$  as the smallest positive integer satisfying  $\text{nz}(S) \leq \rho_S n$ .

When discussing the density of a product matrix  $P = ST$ , we would like to simply consider the density  $\rho_P$ ; however, for technical reasons, we want to ignore zero entries that appear due to cancellations. Formally, let  $\hat{S}$  be the binary matrix defined as

$$\hat{S}[i, j] = \begin{cases} 1 & \text{if } S[i, j] \neq 0, \\ 0 & \text{if } S[i, j] = 0, \end{cases}$$

and define  $\hat{T}$  similarly. Let  $\hat{P} = \hat{S}\hat{T}$ , where the product is taken over the Boolean semiring. We define the density of the product  $ST$ , denoted by  $\hat{\rho}_{ST}$ , as the smallest positive integer satisfying  $\text{nz}(\hat{P}) \leq \hat{\rho}_{ST} n$ . Note that for most of our applications, we operate over semirings where no cancellations can occur in additions, in which case  $\hat{\rho}_{ST} = \rho_P$ .

We also note that while we assume that the input matrices are  $n \times n$ , we can also use this framework for rectangular matrix multiplications, simply by padding the matrices with zeroes to make them square.

*Sparse matrix multiplication algorithm* Our main result for sparse matrix multiplication is the following:

**Theorem 3** *Matrix multiplication  $ST = P$  can be computed deterministically in*

$$O\left(\frac{(\rho_S \rho_T \hat{\rho}_{ST})^{1/3}}{n^{2/3}} + 1\right)$$

*rounds in CONGESTED CLIQUE, assuming we know  $\hat{\rho}_{ST}$  beforehand.*

We note that for all of our applications, the requirement that we know  $\hat{\rho}_{ST}$  beforehand is satisfied. However, the algorithm can be modified to work without knowledge of  $\hat{\rho}_{ST}$  with the additional cost of multiplicative  $O(\log n)$  factor; we simply start with estimate  $\hat{\rho}_{ST} = 2$ , restarting the algorithm with doubled estimate if the running time for current estimate is exceeded.

The rest of this section gives the proof of Theorem 3. We start by describing the overall structure of the algorithm, and then detail the different phases of the algorithm individually.

#### 2.1.1 Algorithm description

*Algorithm parameters* We define the algorithm parameters  $a, b$  and  $c$  as

$$a = \frac{(\rho_T \hat{\rho}_{ST} n)^{1/3}}{\rho_S^{2/3}}, \quad b = \frac{(\rho_S \hat{\rho}_{ST} n)^{1/3}}{\rho_T^{2/3}},$$

$$c = \frac{(\rho_S \rho_T n)^{1/3}}{\hat{\rho}_{ST}^{2/3}}.$$

These parameter will control how we split the matrix multiplication task into independent subtasks. To see why these parameters are chosen in this particular way, we note that we will require that  $abc = n$ , and the final running time of the algorithm will be  $O(\rho_S a/n + \rho_T b/n + \hat{\rho}_{ST} c/n + 1)$  rounds, which is optimized by selecting  $a, b$  and  $c$  as above; this gives the running time in Theorem 3.

For simplicity, we assume that  $a, b$  and  $c$  are integers. If not, taking the smallest greater integer will cause at most constant overhead.

*Algorithm overview* Our algorithm follows the basic idea of the classical 3D matrix multiplication algorithm, as presented for CONGESTED CLIQUE by [13], and as adapted for sparse matrix multiplication by [14]. That is, we want to reduce the matrix multiplication task into  $n$  smaller instances of matrix multiplication, and use a single node for each one of these. However, due to the fact that we are working with sparse matrices, we have to make certain considerations in our algorithm:

- Whereas 3D matrix multiplication splits the original multiplication into  $n$  products of  $n^{2/3} \times n^{2/3}$  matrices, we would ideally like to split into  $n$  products of shape  $(n/b \times n/c) \times (n/c \times n/a)$ .
- Unlike in the dense case, we also have to make sure that all of our  $n$  products are equally sparse. Recent distributed triangle detection algorithms [15,49] use randomized partitioning techniques to achieve a similar goal, but these do not immediately give all the properties we need. Instead, we adopt a deterministic approach to ensure optimal balancing of the product sparsity.

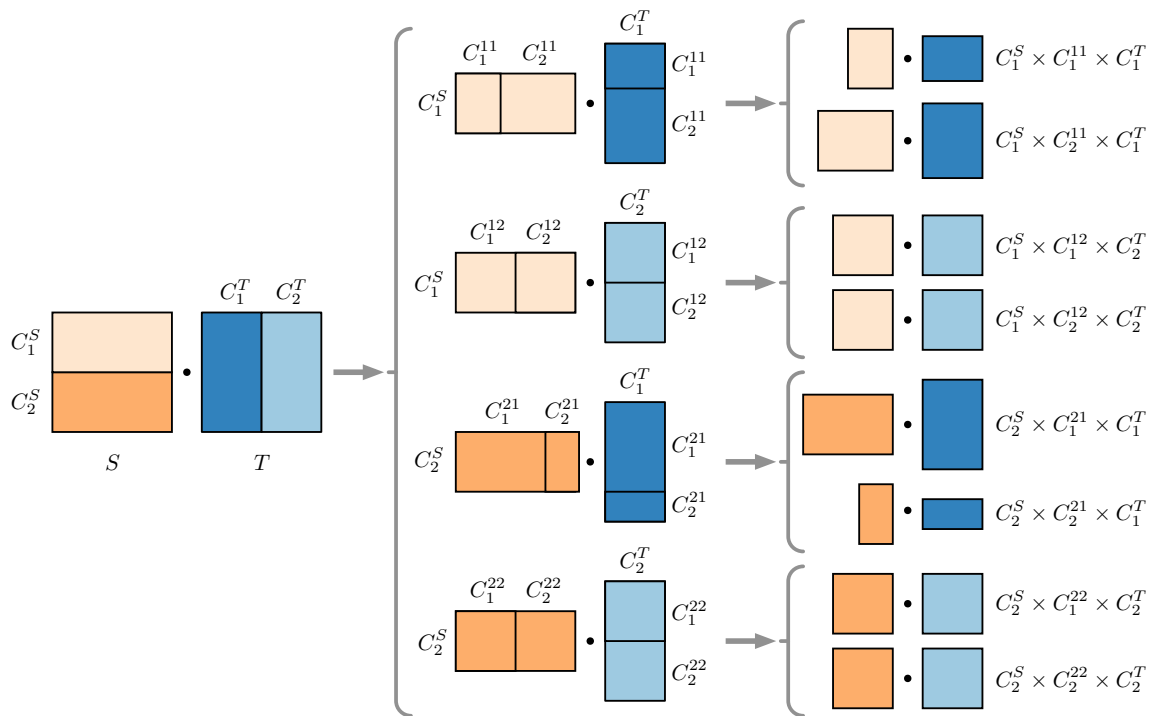


Fig. 1 Matrix multiplication decomposed into subtasks using Lemma 5 for  $a = b = c = 2$  and  $n = 8$ . The notation follows the proof of Lemma 5

With the above considerations in mind, we now give an overview of our sparse matrix multiplication algorithm:

1. We compute a partition of the matrix multiplication task  $P = ST$  into  $n$  subtasks  $P^v = S^v T^v$ , where  $v \in V$ , so that each  $P^v$  is a  $n/b \times n/a$  matrix,  $S^v$  and  $T^v$  are submatrices of  $S$  and  $T$ , respectively, and we have

$$\begin{aligned} \text{nz}(S^v) &= O(\rho_S n/bc) = O(\rho_S a) \quad \text{and} \\ \text{nz}(T^v) &= O(\rho_T n/ac) = O(\rho_T b) \end{aligned}$$

for all  $v \in V$ . This step takes  $O(1)$  rounds (Sect. 2.1.2).

2. Each node  $v$  learns the matrices  $S^v$  and  $T^v$ , and computes the product  $P^v = S^v T^v$ . Note that after this step, some of the matrices  $P^v$  may be very dense. This step takes  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds (Sect. 2.1.3).
3. We balance the output matrices  $P^v$  so that each node holds  $O(\widehat{\rho}_{ST} n/ab) = O(\widehat{\rho}_{ST} c)$  values that need to be summed to obtain the final output matrix  $P$ . This is achieved by duplicating those subtasks where the output is too dense. This step takes  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds (Sect. 2.1.4).
4. The intermediate values obtained in Step 3 are summed together to obtain the output matrix  $P$ . This step takes  $O(\widehat{\rho}_{ST} c/n + 1)$  rounds (Sect. 2.1.5).

Note that the total running time of the above algorithm will be  $O(\rho_S a/n + \rho_T b/n + \widehat{\rho}_{ST} c/n + 1)$  rounds, which

by the choice of  $a, b$  and  $c$  is as stated in Theorem 3. Note that Steps (1) and (2) are essentially streamlined versions of corresponding tools from [14], while Steps (3) and (4) are new.

### 2.1.2 Cube partitioning

We say that a *subcube* of  $V^3$  is a set of form  $V_1 \times V_2 \times V_3$ , where  $V_1, V_2, V_3 \subseteq V$ . Note that such a subcube corresponds to a matrix multiplication task  $S[V_1, V_2]T[V_2, V_3]$ . Thus, a partition of the cube  $V^3$  into subcubes corresponds to a partition of the original matrix multiplication into smaller matrix multiplication tasks, as discussed in the overview; see also Fig. 1.

**Lemma 5** *There is a CONGESTED CLIQUE algorithm running in  $O(1)$  rounds that produces globally known a partition of  $V^3$  into  $n$  disjoint subcubes  $V_i$  such that for each subcube  $V_i = V_i^S \times V_i^P \times V_i^T$ , we have  $|V_i^S| = n/b, |V_i^T| = n/a$  and the total number of non-zero entries is*

1.  $O(\rho_S a + n)$  in the submatrix  $S[V_i^S, V_i^P]$ , and
2.  $O(\rho_T b + n)$  in the submatrix  $T[V_i^P, V_i^T]$ .

**Proof** We start by partitioning the input matrices into equally sparse ‘slices’. Specifically, we do the following:

1. All nodes  $v$  broadcast the number of non-zero entries on row  $v$  of  $S$ . Based on this information, all nodes deterministically compute the same partition of  $V$  into  $b$  sets  $C_1^S, C_2^S, \dots, C_b^S$  of size  $n/b$  such that  $\text{nz}(S[C_i^S, V]) = O(\rho_S n/b + n)$ ; such partition exists by Lemma 2.
2. Using the same procedure as above, the nodes compute a partition of  $V$  into  $a$  sets  $C_1^T, C_2^T, \dots, C_a^T$  of size  $n/a$  such that  $\text{nz}(T[V, C_i^T]) = O(\rho_T n/a + n)$ .

There are now  $ab$  pairs  $(C_i^S, C_j^T)$ , each corresponding to a subcube  $C_i^S \times V \times C_j^T$ . We now partition each of these subcubes in parallel as follows. First, we partition the nodes into  $ab$  sets of  $n/ab = c$  nodes, each such set  $B_{ij}$  corresponding to a pair of indices  $(i, j) \in [b] \times [a]$ . The final partition is now computed as follows:

1. Nodes redistribute the input matrices so that node  $v$  holds column  $v$  of  $S$  and row  $v$  of  $T$ . This can be done in  $O(1)$  rounds.
2. In parallel for each  $i$  and  $j$ , node  $v$  sends the number of non-zero elements in  $S[C_i^S, v]$  and  $T[v, C_j^T]$  to all nodes in  $B_{ij}$ .
3. Nodes in  $B_{ij}$  compute a partition of  $V$  into  $c$  sets  $C_1^{ij}, C_2^{ij}, \dots, C_c^{ij}$  of consecutive indices such that the number of non-zero entries in  $S[C_i^S, C_k^{ij}]$  is  $O(\rho_S n/bc + n)$  and the number of non-zero entries in  $T[C_k^{ij}, C_j^T]$  is  $O(\rho_T n/ac + n)$ ; such partition exists by Lemma 4.
4. For each  $C_k^{ij}$ , the  $k$ th node in  $B_{ij}$  broadcasts the first and last index of  $C_k^{ij}$  to all other nodes, allowing all nodes to reconstruct these sets.

The subcubes  $C_i^S \times C_k^{ij} \times C_j^T$  are now known globally. Furthermore, since  $n/bc = a$  and  $n/ac = b$ , the subcubes satisfy the requirements of the claim by construction.  $\square$

### 2.1.3 Intermediate products

*Balancing* Given the cube partitioning of Lemma 5, the next step is to distribute the matrix entries so that each node can compute a single subproduct given by a subcube in the partition. We start by giving an auxiliary balancing tool that we use to ensure that each node has to send a roughly equal amount of entries, as we will see below.

Concretely, we consider a balancing task where each node  $i$  has  $n$  weighted entries with weights  $w_{i1}, \dots, w_{in} \in \mathbb{N}$  that satisfy

$$\sum_{i,j} w_{ij} \leq W, \quad w_{ij} \leq n,$$

and the goal is to re-distribute the entries so that each node has  $n$  entries with total weight  $O(W/n + n)$ . Concretely,

we assume each weighted entry consists of the weight and  $O(\log n)$  bits of data.

**Lemma 6** *The above balancing task can be solved in  $O(1)$  rounds in CONGESTED CLIQUE.*

**Proof** As a first step, we globally learn the distribution of different weights, and compute a globally known ordering for the weighted entries:

1. All nodes send the number of entries with weight exactly  $i$  to node  $i$ , with node 1 handling entries with both weight 0 and 1. Node  $i$  broadcasts the total to all other nodes.
2. Nodes sort the weighted entries using Lenzen’s sorting algorithm [43].

Since all nodes know the distribution of the weights, all nodes can now locally compute what entries other nodes hold. This gives us a globally consistent numbering for the weighted entries, which we can use to solve the balancing task:

1. Nodes locally compute a partition of weighted entries into  $n$  sets of size  $n$  with total weight  $O(W/n + n)$ ; such partition exists by Lemma 2.
2. Each set is assigned to a separate node. Nodes redistribute the entries so that each node receives their assigned set in  $O(1)$  rounds.

All steps clearly take  $O(1)$  rounds.  $\square$

*Computing intermediate products* We now show how to compute the intermediate products given by the cube partition of Lemma 5. The following lemma is in fact more general; we will use it as a subroutine in subsequent steps of the algorithm.

**Lemma 7** *Let  $V_1, V_2, \dots, V_n$  be a partition of  $V^3$  as in Lemma 5, and let  $\sigma : V \rightarrow [n]$  be a (not necessarily bijective) function that is known to all nodes. There is a CONGESTED CLIQUE algorithm running in  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds such that after the completion of the algorithm, each node  $v$  locally knows the product*

$$P^{\sigma(v)} = S[V_{\sigma(v)}^S, V_{\sigma(v)}^P]T[V_{\sigma(v)}^P, V_{\sigma(v)}^T].$$

**Proof** For each  $i, j \in V$ , let

$$w_{ij} = |\{v \in V : (i, j) \in V_{\sigma(v)}^S \times V_{\sigma(v)}^P\}|$$

if  $S[i, j]$  is non-zero, and  $w_{ij} = 0$  otherwise, that is,  $w_{ij}$  is the number of times  $S[i, j]$  appears in matrices  $S[V_{\sigma(v)}^S, V_{\sigma(v)}^P]$ , or 0 if  $S[i, j]$  is a zero entry. Clearly we have  $w_{ij} \leq n$ , and since the partition satisfies the conditions of Lemma 5, we have that



$$\sum_{i,j} w_{ij} = \sum_{v \in V} \text{nz}(S[V_{\sigma(v)}^S, V_{\sigma(v)}^P]) = O(\rho_S n a + n^2).$$

All nodes can compute the values  $w_{ij}$  locally for their own row, since they depend only on the partition and function  $\sigma$ .

We now distribute the entries of input matrix  $S$  so that each node learns the matrix  $S[V_{\sigma(v)}^S, V_{\sigma(v)}^P]$ :

- Using Lemma 6, we balance the input entries so that each node holds  $n$  entries with total weight  $O(\rho_S a + n)$ . Specifically, for each entry a node receives, it receives the value  $S[i, j]$  along with the index  $(i, j)$ .
- Since the nodes know the partition and function  $\sigma$ , each node computes to which nodes it needs to send each of the entries it received in the first step. Since each entry is duplicated  $w_{ij}$  times, each node needs to send  $O(\rho_S a + n)$  messages, and dually, each node needs to receive a submatrix of  $S$  with  $O(\rho_S a + n)$  entries. These messages can be delivered in  $O(\rho_S a/n + 1)$  rounds.

By identical argument, each node  $v$  can learn the matrix  $T[V_{\sigma(v)}^P, V_{\sigma(v)}^T]$  in  $O(\rho_T b/n + 1)$  rounds and compute  $P^{\sigma(v)}$  locally. □

### 2.1.4 Balanced intermediate products

We say that an *intermediate value* in the matrix multiplication is a value

$$p_v W u = S[v, W]T[W, u] = \sum_{w \in W} S[v, w]T[w, u].$$

That is, an intermediate value is a partial sum of products for a single position of the output matrix. For concreteness, we encode these in the CONGESTED CLIQUE implementation as tuples  $(p_v W u, v, u)$ .

**Lemma 8** *There is a CONGESTED CLIQUE algorithm running in  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds such that after the completion of the algorithm,*

- Each node holds  $O(\widehat{\rho}_{ST} n/ab) = O(\widehat{\rho}_{ST} c)$  non-zero intermediate values, and
- Each non-zero elementary product  $S[v, w]T[w, u]$  in the matrix multiplication is included in exactly one intermediate value held by some node.

**Proof** As the first part of the algorithm, we compute all the intermediate product matrices and learn their densities:

- Compute a partition of  $V^3$  into subcubes  $\{V_v^S \times V_v^P \times V_v^T : v \in V\}$  using Lemma 5.

- Applying Lemma 7 with  $\sigma_1(v) = v$ , each node computes the matrix

$$P^v = S[V_v^S, V_v^P]T[V_v^P, V_v^T].$$

This takes  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds.

- Each node  $v$  broadcasts the number of non-zero entries in  $P^v$  to all other nodes.

Next, we want to balance the dense intermediate product matrices between multiple nodes. We cannot do this directly, but we can instead duplicate the products:

- Construct a function  $\sigma_2$  so that for each  $v$  with  $\text{nz}(P^v) \geq \widehat{\rho}_{ST} c$ , there are at least  $\lfloor \frac{\text{nz}(P^v)}{\widehat{\rho}_{ST} c} \rfloor$  values  $u \in V$  satisfying  $\sigma_2(u) = v$ . To see that this is possible, we observe that by the definition of  $\widehat{\rho}_{ST}$ , there are at most  $\widehat{\rho}_{ST} n$  positions where matrices  $P^v$  can have non-zero entries and each such position is duplicated  $c$  times in the partition of the cube  $V^3$ , implying that  $\sum_{v \in V} \text{nz}(P^v) \leq \widehat{\rho}_{ST} n c$ . Thus, we have

$$\begin{aligned} \sum_{v \in V} \lfloor \frac{\text{nz}(P^v)}{\widehat{\rho}_{ST} c} \rfloor &\leq \sum_{v \in V} \frac{\text{nz}(P^v)}{\widehat{\rho}_{ST} c} = \frac{1}{\widehat{\rho}_{ST} c} \sum_{v \in V} \text{nz}(P^v) \\ &\leq \frac{\widehat{\rho}_{ST} n c}{\widehat{\rho}_{ST} c} = n. \end{aligned}$$

This step can be done locally using information obtained in the first part of the algorithm.

- Apply Lemma 7 with  $\sigma_2$ . This takes  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds.
- For each  $u$ , each node  $v$  with  $\sigma_1(v) = u$  or  $\sigma_2(v) = u$  assumes responsibility for  $O(\widehat{\rho}_{ST} c)$  entries of the matrix  $P^u$  and discards the rest. More specifically, node  $v$  determines  $i$  such that  $v$  is the  $i$ th node responsible for  $P^u$ , splits the non-zero entries of  $P^u$  into  $\lceil \frac{\text{nz}(P^u)}{\widehat{\rho}_{ST} c} \rceil$  parts and selects the  $i$ th part; if both  $\sigma_1(v) = u$  and  $\sigma_2(v) = u$ , then  $v$  selects two parts. This step can be done locally based on information obtained earlier.

After the completion of the algorithm, each node has  $O(\widehat{\rho}_{ST} c)$  intermediate values from at most two matrices  $P^v$ . The total running time is  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds. □

### 2.1.5 Balanced summation

**Lemma 9** *Assume that the non-zero intermediate values of the matrix multiplication  $P = ST$  have been computed as in Lemma 8. Then there is a CONGESTED CLIQUE algorithm running in  $O(\widehat{\rho}_{ST} c/n + 1)$  rounds that computes the output matrix  $P$ .*

**Proof** We start by initializing each row of the output matrix to all zero values. Our objective is to accumulate the intermediate values to this initial matrix, with each node  $v$  being responsible for the row  $v$  of the output matrix.

All nodes split their intermediate values into  $O(\widehat{\rho}_{STc}/n)$  sets of at most  $n$  intermediate values. We then repeat the algorithm  $O(\widehat{\rho}_{STc}/n + 1)$  times, each repetition accumulating one set of  $n$  intermediate values from each node:

1. Nodes sort the  $n^2$  intermediate values being processes globally by matrix position. This takes  $O(1)$  rounds using Lenzen’s sorting algorithm.
2. Each node locally sums all intermediate products it received corresponding to the same position.
3. Each node broadcasts the minimum and maximum matrix position it currently holds. Nodes use this information to deduce if the same matrix position occurs in multiple nodes. If so, all sums corresponding to that position are sent to the smallest id node having that position; each node sends at most one sum and receives at most one sum from each other node, so this step takes  $O(1)$  rounds.
4. If a node received new values from other nodes, these are now added to the appropriate sum.
5. All nodes now hold sums corresponding to at most  $n$  matrix positions. Using Lenzen’s routing algorithm, we redistribute these so that node  $v$  obtains sums corresponding to positions on row  $v$ ; this takes  $O(1)$  rounds. Node  $v$  then accumulates these sums to the output matrix.

Clearly, after completion of all repeats, we have obtained the output matrix  $P$ . Since each repeat of the accumulation algorithm takes  $O(1)$  rounds, and there are  $O(\widehat{\rho}_{STc}/n + 1)$  repeats, the whole process takes  $O(\widehat{\rho}_{STc}/n + 1)$  rounds.  $\square$

## 2.2 Matrix multiplication with sparsification

Our second matrix multiplication result is a variant of sparse matrix multiplication where we control the density of the output matrix.

*Problem definition* In this section, we assume our semiring  $(R, +, \cdot, 0, 1)$  satisfies the following conditions:

1. There is a total order  $<$  on  $R$ , and
2. The addition operation  $+$  satisfies  $x + y = \min(x, y)$ , where  $\min$  is taken in terms of order  $<$ .

Let  $P$  be a matrix and let  $0 \leq \rho \leq n$  be an integer. We define the  $\rho$ -filtered version of  $P$  as a matrix  $\overline{P}$  such that each row of  $\overline{P}$  contains  $\rho$  smallest entries of  $P$ , that is,

1. Either  $\overline{P}[v, u] = 0$  or  $\overline{P}[v, u] = P[v, u]$ ,

2. If row  $v$  of  $P$  has  $\sigma$  non-zero entries, then row  $v$  for  $\overline{P}$  has  $\min(\sigma, \rho)$  non-zero entries, and
3. If  $\overline{P}[v, u] = 0$  and  $P[v, u] \neq 0$ , then  $\max_w \overline{P}[v, w] \leq P[v, u]$ .

When  $\rho$  is clear from the context, we call  $\overline{P}$  a filtered version of  $P$ .

Let  $S$  and  $T$  be the input matrices over the semiring  $R$  and denote  $P = ST$ . In the *filtered matrix multiplication problem*, we are given the input matrices  $S$  and  $T$  along with an output density parameter  $\rho$ , and the task is to compute a  $\rho$ -filtered output matrix  $\overline{P}$ .

*Filtered matrix multiplication algorithm* We now prove an analogue of Theorem 3 for the filtered matrix multiplication problem. Our result is the following:

**Theorem 4** *Assume we know beforehand a set  $R' \subseteq R$  of semiring elements that can appear during the computation of the product  $ST$ , and  $|R'| = W$ . Then  $\rho$ -filtered matrix multiplication can be computed in*

$$O\left(\frac{(\rho_S \rho_T \rho)^{1/3}}{n^{2/3}} + \log W\right)$$

*rounds in the CONGESTED CLIQUE. In particular, when the underlying semiring is the min-plus semiring and the matrix entries are integers of absolute value at most  $O(n^c)$  for a constant  $c \geq 1$ , the running time is*


$$O\left(\frac{(\rho_S \rho_T \rho)^{1/3}}{n^{2/3}} + \log n\right)$$

*rounds.*

The high-level proof idea is largely the same as for Theorem 3. However, since we now cannot guarantee that the output matrix, and by extension the results of the subtasks, are sparse, we will have to perform part of the filtering before we sum the intermediate results together—in essence, between Steps (2) and (3) of the sparse matrix multiplication algorithm.

For this first filtering step, we group the nodes into sets responsible for subsets of intermediate results on the same row of the output matrix, and perform multiple parallel binary searches to identify which intermediate results can be safely filtered out. This allows us to discard all but  $O(\rho nc)$  intermediate entries, which gives us the additive  $O(\log W)$  overhead in the algorithm. After filtering, we follow the same strategy as in the sparse matrix multiplication algorithm—namely, Steps (3) and (4)—to sum together the intermediate entries. Each node will then hold a *partially filtered* row of the output matrix  $P = ST$ . We can then perform a second filtering step locally to obtain  $\overline{P}$ .

**Fig. 2** Example of combining the intermediate product matrices to form the matrices  $P_k$  for  $a = b = c = 2$  and  $n = 8$



**2.2.1 Algorithm description**

*Cube partitioning in detail* For the filtered version of the matrix multiplication, we need to keep more detailed track of the partitioning given by Lemma 5. Specifically, Lemma 5 gives us a partitioning of the cube  $V^3$  into  $n$  subcubes of form

$$C_i^S \times C_k^{ij} \times C_j^T, \quad (i, j, k) \in [b] \times [a] \times [c],$$

which defines the matrix multiplication sub-tasks  $P^v = S^v T^v$ . In this section, we write the products  $S^v T^v$  instead as  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$ .

*Filtering* To precisely formulate our filtering steps, we define  $n \times n$  matrices  $P_k$  for  $k \in [c]$  by setting

$$P_k[C_i^S, C_j^T] = S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T].$$

That is, each  $P_k$  is obtained by combining the smaller  $n/b \times n/a$  matrices  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$  into an  $n \times n$  square matrix; see Fig. 2. In particular, the final unfiltered product matrix  $P = ST$  is obtained by summing the  $P_k$  matrices as  $P = \sum_{k=1}^c P_k$ .

We can obtain the filtered version  $\bar{P}$  of the final output matrix  $P$  by first applying the filtering to the matrices  $P_k$  and then filtering the sum of those matrices again, that is,

$$Q = \sum_{k=1}^c \bar{P}_k, \quad \bar{P} = \bar{Q}.$$

Moreover, by the definition of filtering, each matrix  $\bar{P}_k$  has density at most  $\rho$ , so the total number of entries in matrices  $\bar{P}_k$  is  $O(\rho nc)$ . The high-level idea of our algorithm is now that after computation of the matrix multiplication subtasks  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$ , we identify which intermediate values are entries of the matrices  $\bar{P}_k$  and discard the rest, which allows us to complete the summation step within the desired time budget.

*Algorithm overview* More precisely, the filtered matrix multiplication algorithm proceeds in the following steps:

1. We compute a partition of the matrix multiplication task  $P = ST$  into  $n$  sparse subtasks

$$P_k[C_i^S, C_j^T] = S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T],$$

where  $S[C_i^S, C_k^{ij}]$  and  $T[C_k^{ij}, C_j^T]$  are submatrices of  $S$  and  $T$ , respectively, and we have

$$\text{nz}(S[C_i^S, C_k^{ij}]) = O(\rho sa) \quad \text{and}$$

$$\text{nz}(T[C_k^{ij}, C_j^T]) = O(\rho Tb).$$

This step takes  $O(1)$  rounds (Identical to Sect. 2.1).

2. Each node  $v$  learns the matrices  $S[C_i^S, C_k^{ij}]$  and  $T[C_k^{ij}, C_j^T]$  for a single  $(i, j, k) \in [b] \times [a] \times [c]$ , and computes their product  $P_k[C_i^S, C_j^T]$ . This step takes  $O(\rho sa/n + \rho Tb/n + 1)$  rounds (Identical to Sect. 2.1).
3. For each  $k \in [c]$  and row  $\ell$  of  $P_k$ , the nodes that hold entries from row  $\ell$  of  $P_k$  perform a distributed binary search to find the  $\rho$ th largest value on row  $\ell$ . This allows us to identify which intermediate entries will appear in matrices  $\bar{P}_k$  and to discard the rest. This step takes  $O(\log W)$  rounds (Lemma 10).
4. We balance the entries of matrices  $\bar{P}_k$  so that each node holds  $O(\rho n/ab) = O(\rho c)$  values that need to be summed to obtain the matrix  $Q = \sum_k \bar{P}_k$ . This is achieved by duplicating those subtasks where the output contains too many entries from matrices  $\bar{P}_k$ . This step takes  $O(\rho sa/n + \rho Tb/n + 1)$  rounds (Lemma 11).
5. The intermediate values obtained in Step 5 are summed together to obtain the matrix  $Q$  so that each node holds a single row of  $Q$ . This step takes  $O(\rho c/n + 1)$  rounds. (Identical to Sect. 2.1).

- Each node discards all but the  $\rho$  smallest entries on their row of  $Q$  to obtain final output  $\overline{P}$ . This step requires only local computation.

### 2.2.2 Filtering

*Node grouping for filtering* We identify in an arbitrary fashion each node  $v \in V$  with a triple  $(i, j, k) \in [b] \times [a] \times [c]$ . This node  $v$  is responsible for computing the intermediate product

$$S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T].$$

Furthermore, for  $i \in [b]$  and  $k \in [c]$ , let  $B_{ik} \subseteq V$  be the set of nodes corresponding to triples  $(i, j, k)$  for  $j \in [a]$ . In particular, for fixed  $i$  and  $k$ , the nodes in  $B_{ik}$  are responsible for computing the products that give the rows  $\ell \in C_i^S$  of the matrix  $P_k$ . Finally, we observe that  $|B_{ik}| = a$ .

*Filtering intermediate products* Recall that we assume that we know a set  $R' \subseteq R$  of size  $W$  of semiring elements that can appear during the computation of the product  $ST$ . We now show how the nodes compute the *cutoff value* for each row of each matrix  $\overline{P}_k$ . Formally, to define the cutoff value for row  $\ell$  of matrix  $P_k$ , consider the set  $\{(P_k[\ell, i], i) : P_k[\ell, i] \neq 0\}$  equipped with the natural ordering, that is,  $(r, s) < (r', s')$  if  $r < r'$  or if  $r = r'$  and  $s < s'$ . The cutoff value is the  $\rho$ th largest element in the set, or the largest if the set has size less than  $\rho$ .

**Lemma 10** *Assume that each node has computed the corresponding product  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$ . Then there is a CONGESTED CLIQUE algorithm running in  $O(\log W)$  rounds such that after the completion of the algorithm, each node  $v \in B_{ik}$  knows the cutoff value for row  $\ell$  in  $P_k$  for all rows  $\ell \in C_i^S$ .*

**Proof** Fix  $i \in [b], k \in [c]$  and a row  $\ell \in C_i^S$ . We use binary search to find the smallest value  $r$  such that there are at least  $\rho$  non-zero entries at most  $r$  on row  $\ell$  of  $P_k$ . We fix a node  $u \in B_{ik}$  the coordinator for row  $\ell$  of  $P_k$ . The binary search now proceeds as follows:

- In the first round of the search, all nodes  $v \in B_{ik}$  send the number of non-zero entries they have on row  $\ell$  to the coordinator  $u$ . If the total number of non-zero entries is at most  $\rho$ , we are done. Otherwise, the coordinator sets  $r_1 = \min R'$  and  $r_2 = \max R'$ , and we proceed with the binary search.
- In the subsequent rounds, the coordinator selects the value  $r' \in R'$  halfway between  $r_1$  and  $r_2$  and broadcasts it to all nodes  $v \in B_{ik}$ . The nodes  $v \in B_{ik}$  then send the coordinator the number of values at most  $r'$  they have on row  $\ell$ . If the total number is less than  $\rho$ , the coordinator sets  $r_1 = r'$ , and otherwise the coordinator sets  $r_2 = r'$ .

We repeat this step until there are no values in  $R'$  between  $r_1$  and  $r_2$ .

The binary search finishes after  $O(\log W)$  iterations. Clearly  $r = r_2$  is the desired value. To obtain the final cutoff value, all nodes  $v \in B_{ik}$  tell the coordinator (1) how many values strictly less than  $r$  they have, and (2) how many values equal to  $r$  they have. The coordinator then determines how many values equal to  $r$  should be kept, and determines which node in  $B_{ik}$  holds the entry corresponding to the cutoff value. The coordinator then queries that node to obtain the final cutoff value  $(r, s)$ , and broadcasts it to  $B_{ik}$ . The post-processing after the binary search takes  $O(1)$  rounds.

It remains to show that we can execute all the binary searches in parallel. For each set  $B_{ik}$  of  $a$  nodes, the nodes have to execute  $O(n/b)$  binary searches, and we assign each node  $v \in B_{ik}$  as a leader for  $O(n/ab)$  binary searches. Thus, for each iteration of the binary search, each node  $v \in B_{ik}$  needs to send and receive  $O(an/ab) = O(n/b) = O(n)$  messages in the coordinator role. Dually, each node participates in  $O(n/b)$  binary searches, so it needs to send and receive  $O(n/b) = O(n)$  messages for each iteration of the binary search in the participant role. Thus, one iteration of all binary searches can be completed in  $O(1)$  rounds.  $\square$

Given the cutoff values computed by Lemma 10, all nodes can filter their local products to discard all entries that don't appear in matrices  $\overline{P}_k$ . However, some nodes may still hold too many entries from matrices  $\overline{P}_k$ , so we will balance the entries in a similar way as in Lemma 8:

**Lemma 11** *Assume that the cutoff values of rows of matrices  $P_k$  are computed as per Lemma 10. Then there is a CONGESTED CLIQUE algorithm running in  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds such that after the completion of the algorithm,*

- Each node holds  $O(\rho c)$  non-zero entries from matrices  $\overline{P}_k$ , and
- Each non-zero entry of matrices  $\overline{P}_k$  is held by exactly one node.

**Proof** For node  $v \in V$ , let  $w_v$  be the number of entries from the matrices  $\overline{P}_k$  the node  $v$  holds after the product  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$  is filtered using Lemma 10. We now want to apply Lemma 7; to this end, we construct a helper assignment function  $\sigma$  so that

- For Each  $v$  with  $w_v \geq \rho c$ , there are at least  $\lfloor \frac{w_v}{\rho c} \rfloor$  values  $u \in V$  satisfying  $\sigma(u) = v$ , and
- For Each  $v \in B_{ik}$ , all  $u$  such that  $\sigma(u) = v$  also satisfy  $u \in B_{ik}$ .

To see that this is possible, consider a set  $B_{ik}$ ; since these nodes together hold  $|C_i^S|$  rows of the matrix  $\overline{P}_k$ , we have that

$\sum_{v \in B_{ik}} w_v \leq \rho |C_i^S| = \rho n/b$ . Thus, we need to assign a total of

$$\begin{aligned} \sum_{v \in B_{ij}} \lfloor \frac{w_v}{\rho c} \rfloor &\leq \sum_{v \in B_{ij}} \frac{w_v}{\rho c} = \frac{1}{\rho c} \sum_{v \in B_{ij}} w_v \\ &\leq \frac{\rho n/b}{\rho c} = n/bc = a \end{aligned}$$

nodes with  $\sigma(v) \in B_{ik}$ . Since all nodes  $v \in V$  can broadcast  $w_v$  and the assignments to node sets  $B_{ik}$  are known globally, all nodes can construct the function  $\sigma$  locally.

The rest of the algorithm now proceeds as follows:

1. Apply Lemma 7 with  $\sigma$ . This takes  $O(\rho_S a/n + \rho_T b/n + 1)$  rounds.
2. For each  $(i, j, k) \in [b] \times [a] \times [c]$ , each node that either originally computed the product  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$  or was assigned it by  $\sigma$  computes which entries in the product are non-zero entries of  $\overline{P}_k$ ; this is possible since all nodes that were assigned the product  $S[C_i^S, C_k^{ij}]T[C_k^{ij}, C_j^T]$  are in  $B_{ik}$ , and thus by Lemma 10 know the cutoff values for rows  $r \in C_i^S$  of matrix  $\overline{P}_k$ . Each of these nodes then assumes responsibility for  $O(\rho c)$  entries of  $\overline{P}_k$  from that product and discards the rest. This step can be done locally based on information obtained earlier.  $\square$

### 3 Distance tools

In this section, we use our matrix multiplication algorithms to construct basic distance computation tools that will be used for our final distance computation algorithms. Though we only use the distance tools for undirected graphs, we note that they work also for directed graphs. As noted before, we assume the edge weights are non-negative integers at most  $O(n^c)$  for a constant  $c$ .

#### 3.1 Distance products

*Augmented min-plus semiring* The general algorithmic idea for our distance tools is to apply the matrix multiplication algorithms over the min-plus semiring. However, to ensure that we get consistent results in terms of hop distances, a property we require for our distance  $k$ -NEAREST and  $(S, d, k)$ -SOURCE DETECTION distance tools, we augment the basic min-plus semiring to keep track of the number of hops.

We define the *augmented min-plus semiring*  $R$  to encode paths in distance computations as follows. The elements of  $R$  are tuples  $(w, h)$ , where

1.  $w$  is either the weight of an edge or a path, or  $\infty$ , and
2.  $h$  is a non-negative integer or  $\infty$ , representing the number of hops.

Let  $<$  be the lexicographical order on tuples  $(w, h)$ , and define the addition operator  $\min$  as the minimum over the total order given by  $<$ . The multiplication operation  $+$  is defined as  $(w_1, h_1) + (w_2, h_2) = (w_1 + w_2, h_1 + h_2)$ . It is easy to verify that  $(R, \min, +, (\infty, \infty), (0, 0))$  is a semiring with idempotent addition, that is,  $\min(r, r) = r$  for all  $r$ . Moreover, the structure satisfies the conditions of Theorem 4.

*Distance products* We call the product of  $S$  and  $T$  over the augmented min-plus semiring the *augmented distance product* of  $S$  and  $T$ , and denote it by  $S \star T$ . In particular, for a graph  $G = (V, E)$ , we define the *augmented weight matrix*  $W$  by setting

$$W[u, v] = \begin{cases} (0, 0) & \text{if } u = v, \\ (w(u, v), 1) & \text{if } (u, v) \in E, \text{ and} \\ (\infty, \infty) & \text{otherwise.} \end{cases}$$

As with the regular distance product, the  $d$ th augmented distance product power  $W^d$  gives the distances for all pairs of nodes  $u, v \in V$  using paths of at most  $d$  hops, as well as the associated number of hops.

Finally, we observe that the augmented distance product gives a consistent ordering in terms of distance from  $v$ , in the following sense:

**Lemma 12** *Let  $s, t \in V$ , and let  $P$  be the shortest path of at most  $d$  hops from  $s$  to  $t$ . Then for every node  $v$  on the path  $P$ , we have  $W^d[s, v] < W^d[s, t]$ .*

*Proof* Let  $W^d[s, t] = (w, h)$ . Since we assume that weights are non-negative, the subpath  $P'$  of  $P$  from  $s$  to  $v$  has weight at most  $w$  and uses at most  $h - 1$  edges. If  $P'$  is not a shortest path of at most  $d$  hops from  $s$  to  $v$ , the actual shortest path  $P^*$  of at most  $d$  hops has either weight  $w^* < w$ , or has weight  $w$  and uses  $h^* < h - 1$  hops. Thus we have

$$W^d[s, v] \leq (w, h - 1) < (w, h) = W^d[s, t],$$

proving the claim.  $\square$

*Recovering paths* While we do not explicitly consider recovering the actual shortest paths in our algorithms, we mention that in addition to computing powers of the augmented weight matrix, the distance product algorithm for computing shortest path distances can be extended to recover *routing tables*, as discussed in [13]. That is, a *routing table*  $R$  corresponding to  $W^d$  is an  $n \times n$  matrix, where each entry  $R[u, v]$  is a node  $w \in N(v)$  that lies on the shortest path of at most  $d$  hops from  $u$  to  $v$ , if  $W^d[u, v]$  is finite.

To construct the routing table using distance products, we require that the matrix multiplication over the augmented min-plus semiring also produces a *witness* for each non-zero entry in a product  $P = ST$ ; for any non-zero  $P[u, v]$ , a witness is a node  $w_{uv} \in V$  such that  $P[u, v] = S[u, w_{uv}] + T[w_{uv}, v]$ . Since the matrix multiplication algorithms explicitly compute the non-zero products, they can be modified to provide witnesses over augmented min-plus semiring by keeping track of the entry indices during minimum operations. Given matrices  $W^{d_1}$  and  $W^{d_2}$ , a routing table  $R$  for  $W^{d_1}$  and witnesses for the product  $W^{d_1} \star W^{d_2}$ , one can easily verify that a routing table  $R'$  for  $W^{d_1+d_2}$  can be obtained by setting  $R'[u, v] = R[u, w_{uv}]$ .

### 3.2 k-nearest neighbors

In the  $k$ -NEAREST problem, we are given an integer  $k$ , and the task is to compute for each node  $v$  the set of  $k$  nearest nodes and distances to those nodes, breaking ties first by hop distance and then arbitrarily. More formally, we want each node to compute a set  $N_k(v)$  of  $k$  nodes and distances  $d(v, u)$  for all  $u \in N_k(v)$ , such that the values  $W^n[v, u]$  for  $u \in N_k(v)$  are the  $k$  smallest on row  $v$  of  $W^n$  in terms of the order  $\prec$  on the augmented min-sum semiring  $R$ .

Note that it follows immediately from Lemma 12 that all nodes  $u \in N_k(v)$  are at most  $k$  hops away from  $v$ , and all nodes on the shortest path from  $v$  to  $u$  are also in  $N_k(v)$ .

**Theorem 5** *The  $k$ -NEAREST problem can be solved in*

$$O\left(\left(\frac{k}{n^{2/3}} + \log n\right) \log k\right)$$

rounds in CONGESTED CLIQUE.

**Proof** For matrix  $M$ , let  $\overline{M}$  denote the matrix obtained by discarding all but the  $k$  smallest values on each row of  $M$ . To solve the  $k$ -NEAREST problem, we compute the filtered version  $\overline{W^k}$  of the  $k$ th power of the augmented weight matrix  $W$  as follows:

1. All nodes  $v$  discard all but the  $k$  smallest values on row  $v$  to obtain  $\overline{W}$ .
2. We now observe that  $\overline{W^2} = \overline{W \star W}$ ,  $\overline{W^4} = \overline{\overline{W^2} \star \overline{W^2}}$ ,  $\dots$ ,  $\overline{W^k} = \overline{W^{\frac{k}{2}} \star W^{\frac{k}{2}}}$ . Thus, by applying Theorem 4 with  $\rho = k$  iteratively  $O(\log k)$  times, we can compute the matrix  $\overline{W^k}$  in  $O\left(\left(\frac{k}{n^{2/3}} + \log n\right) \log k\right)$  rounds.

To see that this allows us to recover  $N_k(v)$ , we first observe that by Lemma 12, the hop distance from  $v$  to any node in  $N_k(v)$  is at most  $k$ . Moreover, by a simple induction argument using Lemma 12, we have that for all non-zero entries of  $\overline{W^k}$ , we have  $\overline{W^k}[u, v] = W^k[u, v]$ . Thus, the non-zero

entries on row  $v$  give us the set  $N_k(v)$  and the distances to those nodes. □

### 3.3 Source detection

In the  $(S, d, k)$ -SOURCE DETECTION problem, we are given a set of *sources*  $S \subseteq V$  and integers  $k$  and  $d$ , and the task is to compute for each node  $v$  the set of  $k$  nearest sources within  $d$  hops, as well as the distances to those sources using paths of at most  $d$  hops.

**Theorem 6** *The  $(S, d, k)$ -SOURCE DETECTION problem can be solved in*

$$O\left(\left(\frac{m^{1/3}k^{2/3}}{n} + \log n\right)d\right)$$

or

$$O\left(\left(\frac{m^{1/3}|S|^{2/3}}{n} + 1\right)d\right)$$

rounds in CONGESTED CLIQUE, where  $m$  is the number of edges in the input graph.

**Proof** To obtain the first running time, we solve the  $(S, d, k)$ -SOURCE DETECTION problem as follows:

1. All nodes, including nodes in  $S$ , select the  $k$  lightest edges to nodes in  $S$ , breaking ties arbitrarily. Let  $W_1$  be the augmented weight matrix restricted to these edges; since a total of at most  $nk$  edges were selected, we have  $\text{nz}(W_1) \leq nk$ .
2. We now apply Theorem 4 iteratively  $d - 1$  times to compute the products  $W_2 = \overline{W \star W_1}$ ,  $W_3 = \overline{W \star W_2}$ ,  $\dots$ ,  $W_d = \overline{W \star W_{d-1}}$ . We have  $\rho_W = m/n$  and  $\rho_{W_i} = k$ , and use  $\rho = k$  as the output density, so computing all the products takes  $O\left(\left(\frac{m^{1/3}k^{2/3}}{n} + \log n\right)d\right)$  rounds.

By a simple induction using Lemma 12, we see that the non-zero entries on row  $v$  of  $W_i$  correspond to the  $k$  nearest sources within  $i$  hops of  $v$ .

For the second running time, we instead compute the  $d$ -hop distances from set  $S$  to all other nodes:

1. Let  $U_1$  be the  $n \times |S|$  matrix obtained by restricting the augmented weight matrix  $W$  to edges with at least one endpoint in  $S$ . By padding matrix  $U_1$  with zero entries, we can view it as a square matrix with density  $|S|$ .
2. We now apply Theorem 3 iteratively  $d - 1$  times to compute the products  $U_2 = W \star U_1$ ,  $U_3 = W \star U_2$ ,  $\dots$ ,  $U_d = W \star U_{d-1}$ , where the density of  $W$  is  $m/n$  and the density of all matrices  $U_i$  is  $|S|$ , giving a total running time of  $O\left(\left(\frac{m^{1/3}|S|^{2/3}}{n} + 1\right)d\right)$  rounds.

The matrix  $U_d$  gives the  $d$ -hop distances between nodes in  $S$  and all other nodes, so each node can select the  $k$  closest sources.  $\square$

### 3.4 Distance through node set

In the DISTANCE THROUGH SETS problem, we assume that each node  $v$  has a set  $W_v$  and distance estimates  $\delta(v, w)$  and  $\delta(w, v)$  for all  $w \in W_v$ . The task is for all nodes  $v$  to compute distance estimates  $\min_{w \in W_v \cap W_u} \{\delta(v, w) + \delta(w, u)\}$  for all other nodes  $u \in V$ .

**Theorem 7** *The DISTANCE THROUGH SETS problem can be solved*

$$O\left(\frac{\rho^{2/3}}{n^{1/3}} + 1\right)$$

rounds in CONGESTED CLIQUE, where  $\rho = \sum_{v \in V} |W_v| / n$ .

**Proof** Define matrices  $W_1$  and  $W_2$  as

$$W_1[v, w] = \begin{cases} \delta(v, w) & \text{if } w \in W_v, \text{ and} \\ \infty & \text{if } w \notin W_v, \end{cases}$$

$$W_2[w, v] = \begin{cases} \delta(w, v) & \text{if } w \in W_v \text{ and} \\ \infty & \text{if } w \notin W_v. \end{cases}$$

The distance product  $W_1 \star W_2$  over the standard min-sum semiring clearly gives the desired estimates, and since  $\rho_{W_1} = \rho_{W_2} = \rho$ , it can be computed in  $O\left(\frac{\rho^{2/3}}{n^{1/3}} + 1\right)$  rounds by Theorem 3 (or by [14]), using  $n$  as the density estimate for the output matrix.  $\square$

## 4 Hopsets

In this section, we describe a construction of hopsets in polylogarithmic time in the CONGESTED CLIQUE model. Given a graph  $G = (V, E)$ , a  $(\beta, \epsilon)$ -hopset  $H = (V, E')$  is a graph on the same set of nodes such that the  $\beta$ -hop distances in  $G \cup H$  give  $(1 + \epsilon)$ -approximations for the distances in  $G$ . Formally, for any pair of nodes  $u, v \in V$  it holds that  $d_G(u, v) \leq d_{G \cup H}(u, v)$  and

$$d_G(u, v) \leq d_{G \cup H}^\beta(u, v) \leq (1 + \epsilon)d_G(u, v).$$

Usually the goal is to find a sparse hopset with small  $\beta$ . In our case, we are interested in optimizing both  $\beta$  and the running time for constructing the hopset, but not necessarily the hopset size.

Our construction is based on the recent construction of hopsets in the CONGESTED CLIQUE of Elkin and Neiman [24]. However, the time complexity in [24] depends on  $\beta$

and on the hopset size, and is super-polylogarithmic for any choice of parameters. We show that using our new distance tools we can implement the same construction in polylogarithmic time, regardless of the hopset size, as long as  $\beta$  is polylogarithmic. In particular, we focus on a simple variant with  $\tilde{O}(n^{3/2})$  edges which is enough for all our applications.<sup>1</sup>

### 4.1 Construction overview

We follow the hopset construction of [24], which is based on the emulators of Thorup and Zwick [56]. A similar construction appears also in [39] without a distributed implementation. We focus only on a simple variant of [24,39,56], with slightly different parameters.

Given a graph  $G = (V, E)$ , let  $V = A_0 \supseteq A_1 \supseteq A_2 = \emptyset$ , where  $A_1$  is a hitting set of size  $O(\sqrt{n})$  of all the sets  $N_k(v)$  for  $k = O(\sqrt{n} \log n)$ . I.e., for any node  $v \in V$ , there is a node from  $A_1$  among the closest  $k$  nodes to  $v$ . We can construct  $A_1$  using Lemma 1.

For a given subset  $A \subseteq V$ , we denote by  $d_G(v, A)$  the distance from  $v$  to the closest node in  $A$ . For a node  $v \in V$ , let  $p(v) \in A_1$  be a node of distance  $d_G(v, A_1)$  from  $v$ . In general there may be many possible nodes of distance  $d_G(v, A_1)$  from  $v$ . The node  $p(v)$  would be determined by  $v$ , as follows. During our algorithm, each node computes the set  $N_k(v)$ . Since  $A_1$  is a hitting set there is a node from  $A_1$  in the set  $N_k(v)$  computed, and  $v$  defines  $p(v)$  to be the closest such node (breaking ties arbitrarily). For a node  $v \in A_0 \setminus A_1$ , we define the *bunch*

$$B(v) = \{u \in A_0 : d_G(v, u) < d_G(v, A_1)\} \cup p(v),$$

and for a node  $v \in A_1$ , we define the bunch  $B(v) = A_1$ .

The hopset is the set of edges  $H = \{\{v, u\} : v \in V, u \in B(v)\}$ . We would like to set the length of an edge  $\{v, u\}$  to be  $d_G(v, u)$ . However, we cannot necessarily compute these values. During our algorithm we add exactly all the edges in  $H$  to the hopset, but their weights are not necessarily  $d_G(v, u)$ , but rather an approximation for  $d_G(v, u)$ .

**Claim 13** *The number of edges in  $H$  is  $O(n^{3/2} \log n)$ .*

**Proof** For every node in  $A_1$  we add only edges to nodes in  $A_1$ , and the size of  $A_1$  is  $O(\sqrt{n})$ . In addition, for nodes  $v \in A_0 \setminus A_1$  we add edges only to nodes closer than  $p(v)$ . Since  $A_1$  is a hitting set,  $p(v)$  is among the closest  $k$  nodes to  $v$ , which means that we add at most  $k = O(\sqrt{n} \log n)$  edges for each node, and  $O(n^{3/2} \log n)$  edges in total.  $\square$

<sup>1</sup> Using similar ideas, it is possible to implement also the more general hopset from [24].

## 4.2 Congested clique implementation

Our goal is to build  $H$  efficiently in the CONGESTED CLIQUE model. In [24], the authors suggest an iterative algorithm for computing  $H$ , where in iteration  $1 \leq \ell \leq \log n$  they compute a  $2^\ell$ -bounded hopset, which approximates the distances between all pairs of nodes that have a shortest path between them with at most  $2^\ell$  edges in  $G$ . Given a  $2^\ell$ -bounded hopset  $H^\ell$ , they show that it is enough to run Bellman-Ford explorations up to hop-distance  $O(\beta)$  in the graph  $G \cup H^\ell$  to compute a  $2^{\ell+1}$ -hopset  $H^{\ell+1}$ . For the stretch analysis to carry on, it is also important to add edges to  $H^{\ell+1}$  in a certain order. First, the edges that correspond to bunches of nodes in  $A_0 \setminus A_1$  are added, and only then the edges that correspond to bunches of nodes in  $A_1$  are added. They show that the total time complexity for implementing all the Bellman-Ford explorations is proportional to the hopset size. In particular, for our hopset of size  $\tilde{O}(n^{3/2})$ , the complexity required is at least  $\tilde{O}(\sqrt{n})$  rounds.

We follow the general approach from [24], with two changes. First, we show that using our distance tools we can compute directly all the bunches of nodes in  $A_0 \setminus A_1$ , which allows simplifying slightly the algorithm and analysis. The second and more significant difference is that we use our  $(S, d, k)$ -algorithm to replace the heavy Bellman-Ford explorations, which results in a polylogarithmic complexity instead of a polynomial one. We next describe the algorithm in detail.

### 4.2.1 Algorithm description

*Computing the bunches* We start by computing the bunches for all the nodes  $v \in A_0 \setminus A_1$ . This can be done easily using our algorithm for finding the distances to the  $k$ -nearest nodes. This follows since  $p(v)$  is among the  $k$ -nearest nodes to  $v$ , hence all the nodes in the bunch of  $v$  are among the  $k$ -nearest nodes to  $v$ . Note that all the nodes  $v \in A_0 \setminus A_1$  learn the *exact* distances to all the nodes in their bunch. However, it is not clear how to compute a bunch for a node  $v \in A_1$ , since  $v$  should learn the distances to all the nodes from  $A_1$ . Such nodes may be at hop-distance  $d = \Omega(n)$  from  $v$ , and our  $(S, d, k)$ -algorithm depends linearly on the hop-distance  $d$ . To overcome this, we follow the approach in [24], and show how to implement it efficiently using our distance tools.

*Bounded hopsets* We say that  $H$  is a  $(\beta, \epsilon, t)$ -hopset if for all  $x, y \in G$  it holds that  $d_G(x, y) \leq d_{G \cup H}(x, y)$ , and for all pairs of nodes  $x, y \in G$  that have a shortest path between them with at most  $t$  hops, i.e.,  $d_G(x, y) = d_G^t(x, y)$ , it holds that

$$d_G(x, y) \leq d_{G \cup H}^\beta(x, y) \leq (1 + \epsilon)d_G(x, y).$$

Note that the empty set is a  $(1, 0, 1)$ -hopset, and thus also a  $(\beta, \epsilon, 1)$ -hopset for any  $\beta \geq 1$  and  $\epsilon > 0$ . The algorithm builds hopsets iteratively, where in iteration  $\ell$  it builds a  $(\beta, \epsilon_\ell, 2^\ell)$ -hopset  $H^\ell$ , where  $\epsilon_\ell = \epsilon \cdot \ell$  for some  $0 < \epsilon < 1/\log n$ . The final hopset  $H$  is  $H^{\log n}$ . Let  $H_0 = \{\{u, v\} : u \in A_0 \setminus A_1, v \in B(u)\}$ . We already computed the edges  $H_0$  and include them as part of the hopset for all the hopsets  $H^\ell$ .

*Building  $H^\ell$ .* We define  $H^0 = H_0$ . To construct  $H^\ell$  for  $1 \leq \ell \leq \log n$  we work as follows. Let  $G' = G \cup H^{\ell-1}$ . All the nodes in  $A_1$  compute the distances to all the nodes from  $A_1$  at hop-distance at most  $4\beta$  in the graph  $G'$ . We implement it using our  $(S, d, k')$ -algorithm with  $S = A_1, d = 4\beta, k' = |A_1| = O(\sqrt{n})$ . Each node  $v \in A_1$ , adds to the hopset  $H^\ell$  all the edges  $\{v, u\}$  where  $u \in A_1$  at hop-distance at most  $4\beta$  from  $v$  in  $G'$ . The weight of the edge is the weight that  $v$  learned during the  $(S, d, k)$ -algorithm. The hopset  $H^\ell$  includes all these edges, as well as all the edges of  $H_0$ . Note that if  $v$  adds an edge  $\{v, u\}$  to the hopset, in the next round it can let  $u$  learn about it, so we can assume that both the endpoints know about the edge.

This completes the description of the algorithm. We next analyze the time complexity and prove the correctness of the algorithm.

### 4.2.2 Complexity

**Claim 14** *The complexity of the algorithm is  $O(\beta \log n + \log^2 n)$  rounds.*

*Proof* From Lemma 1, we can construct deterministically in  $O((\log \log n)^3)$  rounds a hitting set  $A_1$  of size  $O(\sqrt{n})$  that hits all the sets  $N_k(v)$ . For the algorithm, each node  $v$  should learn the set  $N_k(v)$  which is computed using our algorithm for finding the distances to the  $k$ -nearest nodes. From Theorem 5, this takes  $O\left(\left(\frac{k}{n^{2/3}} + \log n\right) \log k\right) = O\left(\left(\frac{\sqrt{n} \cdot \log n}{n^{2/3}} + \log n\right) \log n\right) = O(\log^2 n)$  rounds.

The algorithm for constructing  $H$  starts by computing the bunches of nodes  $v \in A_0 \setminus A_1$ , this only requires each node  $v$  to learn the distances to the nodes  $N_k(v)$ , which was already computed. Next, for  $\log n$  iterations, we compute for each node the  $O(\sqrt{n})$ -nearest nodes from  $A_1$  at hop-distance at most  $4\beta$ . Since  $A_1$  is of size  $O(\sqrt{n})$ , by Theorem 6, each iteration takes  $O\left(\left(\frac{n^{2/3}(\sqrt{n})^{2/3}}{n} + 1\right) \cdot \beta\right) = O(\beta)$  rounds. Overall, all the iterations take  $O(\beta \log n)$  rounds.  $\square$

### 4.2.3 Correctness

We next prove that  $H^\ell$  is indeed a  $(\beta, \epsilon_\ell, 2^\ell)$ -hopset, the proof follows the proof in [24], and is included here for completeness. There are slight changes to adapt it to the specific



variant we consider. We start with a simple claim regarding the edges  $H_0$ .

**Claim 15** *For any  $x, y \in V$ , either  $d_{G \cup H_0}^1(x, y) = d_G(x, y)$  or there exists  $z \in A_1$  such that  $d_{G \cup H_0}^1(x, z) \leq d_G(x, y)$ .*

**Proof** If  $x \in A_1$ , the claim clearly holds by setting  $z = x$ . We next focus on  $x \in A_0 \setminus A_1$ . For each  $x \in A_0 \setminus A_1$ ,  $H_0$  includes the edge  $\{x, p(x)\}$ , as well as all the edges of the form  $\{x, y\}$  where  $d_G(x, y) < d_G(x, p(x))$ . Hence, if  $d_G(x, y) < d_G(x, p(x))$ , there is an edge  $\{x, y\}$  in  $H_0$  of weight  $d_G(x, y)$  and we are done. Otherwise,  $d_G(x, p(x)) \leq d_G(x, y)$  and the edge  $\{x, p(x)\}$  of weight  $d_G(x, p(x))$  is in  $H_0$  where  $p(x) \in A_1$  which completes the proof.  $\square$

The next lemma shows that  $H^\ell$  is a  $(\beta, \epsilon_\ell, 2^\ell)$ -hopset.

**Lemma 16** *Let  $0 < \epsilon < 1/\log n$ . Set  $\delta = \epsilon/4$ , and  $\beta = 3/4$  and let  $x, y \in V$  be such that  $d_G(x, y) = d_G^{2^\ell}(x, y)$ . Then,*

$$d_{G \cup H^\ell}^\beta(x, y) \leq (1 + \epsilon_\ell)d_G(x, y),$$

where  $\epsilon_\ell = \epsilon \cdot \ell$ .

**Proof** The proof is by induction on  $\ell$ . For  $\ell = 0$ , the claim holds since the empty set is a  $(1, 0, 1)$ -hopset. Assume it holds for  $\ell - 1$ , and we prove it for  $\ell$ . Let  $x, y \in V$  be such that  $d_G(x, y) = d_G^{2^\ell}(x, y)$ , and let  $\pi(x, y)$  be a shortest path between  $x$  and  $y$  in  $G$  with at most  $2^\ell$  edges. We partition  $\pi(x, y)$  into  $J \leq 1/\delta$  segments  $\{L_j = [u_j, v_j]\}_{1 \leq j \leq J}$ , each of length at most  $\delta \cdot d_G(x, y)$ , and additionally at most  $1/\delta$  edges  $\{v_j, u_{j+1}\}_{1 \leq j \leq J}$  between these segments. This is done as follows.

Define  $u_1 = x$ , and walk from  $u_1$  on  $\pi(x, y)$  towards  $y$  until the first node  $u_2$  such that  $d_G(u_1, u_2) > \delta \cdot d_G(x, y)$ , if such exists or until  $y$ . The last node in  $\pi(x, y)$  before  $u_2$  is  $v_1$ , and from the definition of  $u_2$  it holds that  $d_G(u_1, v_1) \leq \delta \cdot d_G(x, y)$ . We continue in the same manner (walk from  $u_2$  on  $\pi(x, y)$  towards  $y$  to define  $u_3$ , and so on) and define the rest of the nodes  $u_j, v_j$ . If we walk from  $u_j$  towards  $y$  and there is no node  $u_{j+1}$  where  $d_G(u_j, u_{j+1}) > \delta \cdot d_G(x, y)$ , we define  $v_j = y$ , and  $J = j$ . Since  $d_G(u_i, u_{i+1}) > \delta \cdot d_G(x, y)$ , the number of segments is at most  $1/\delta$ .

Now, we use Claim 15 on the pairs  $\{u_j, v_j\}_{j=1}^J$ . One case is that for all  $1 \leq j \leq J$ , it holds that  $d_{G \cup H_0}^1(u_j, v_j) = d_G(u_j, v_j)$ , which means that the edges  $\{u_j, v_j\}$  exist in  $G \cup H_0$  with weights  $d_G(u_j, v_j)$ . This means that in the graph  $G \cup H_0$  there is a path of hop-distance at most  $2/\delta$  between  $x$  and  $y$  that includes all the edges of the form  $\{u_j, v_j\}_{j=1}^J$  and the edges  $\{v_j, u_{j+1}\}$  between segments. The total weight of the path is exactly  $d_G(x, y)$ . Since  $\beta = 3/4 > 2/\delta$ , and  $H_0 \subseteq H^\ell$ , we get that  $d_{G \cup H^\ell}^\beta(x, y) = d_G(x, y) \leq (1 + \epsilon_\ell)d_G(x, y)$ .

We now analyse the case where there is at least one pair  $\{u_j, v_j\}$  that does not satisfy  $d_{G \cup H_0}^1(u_j, v_j) = d_G(u_j, v_j)$ .

Let  $i, j$  be the indexes of the first and last pairs  $\{u_i, v_i\}, \{u_j, v_j\}$  that do not satisfy the above. From Claim 15, there are nodes  $z_i, z_j \in A_1$  such that

$$d_{G \cup H_0}^1(u_i, z_i) \leq d_G(u_i, v_i) \leq \delta \cdot d_G(x, y),$$

$$d_{G \cup H_0}^1(v_j, z_j) \leq d_G(u_j, v_j) \leq \delta \cdot d_G(x, y).$$

In our algorithm, when constructing  $H^\ell$ , the nodes  $z_i, z_j$  add edges to all the nodes of  $A_1$  at hop-distance at most  $4\beta$  in the graph  $G' = G \cup H^{\ell-1}$ . We next show that in  $G'$  there is a path of hop-distance at most  $4\beta$  between  $z_i, z_j$  and bound its weight. From the induction hypothesis,  $H^{\ell-1}$  is a  $(\beta, \epsilon_{\ell-1}, 2^{\ell-1})$ -hopset. Since  $u_i, v_j$  are in the shortest path between  $x$  and  $y$ , it holds that  $d_G(u_i, v_j) = d_G^{2^\ell}(u_i, v_j)$ . Since any path of hop-distance at most  $2^{\ell-1}$  can be partitioned to two paths of hop-distance at most  $2^{\ell-2}$ , it holds that a  $(\beta, \epsilon_{\ell-1}, 2^{\ell-1})$ -hopset is also a  $(2\beta, \epsilon_{\ell-1}, 2^\ell)$  hopset. Hence, there is a path of hop-distance at most  $2\beta$  between  $u_i$  and  $v_j$  in  $G \cup H^{\ell-1}$  of weight at most  $(1 + \epsilon_{\ell-1})d_G(u_i, v_j)$ . Since the edges  $\{u_i, z_i\}, \{v_j, z_j\}$  of weight at most  $\delta \cdot d_G(x, y)$  are in  $H_0$ , we get that there is a path of hop-distance at most  $2\beta + 2$  between  $z_i$  and  $z_j$  in  $G'$  of total weight at most  $(1 + \epsilon_{\ell-1})d_G(u_i, v_j) + 2\delta \cdot d_G(x, y)$ . Since  $2\beta + 2 \leq 4\beta$ ,  $z_i$  and  $z_j$  add an edge between them to  $H^\ell$ , which gives

$$d_{G \cup H^\ell}^1(z_i, z_j) \leq (1 + \epsilon_{\ell-1})d_G(u_i, v_j) + 2\delta \cdot d_G(x, y).$$

We next bound  $d_{G \cup H^\ell}^\beta(x, y)$ . From the definition of  $i, j$ , for any  $i' < i$  or  $j < j'$ , there is an edge  $\{u_{i'}, v_{j'}\}$  in  $H_0$  (and hence in  $H^\ell$ ) of weight  $d_G(u_{i'}, v_{j'})$ . These edges together with the edges in  $\pi(x, y)$  before or after these segments sum up to at most  $2/\delta$  edges. Between  $u_i$  and  $v_j$  we have a 3-hop path  $\{u_i, z_i, z_j, v_j\}$  of total weight at most  $\delta \cdot d_G(x, y) + d_{G \cup H^\ell}^1(z_i, z_j) + \delta \cdot d_G(x, y) \leq (1 + \epsilon_{\ell-1})d_G(u_i, v_j) + 4\delta \cdot d_G(x, y)$ . To conclude,  $G \cup H^\ell$  has a path of hop-distance at most  $2/\delta + 3 \leq \beta$  of total weight  $d_G(x, u_i) + (1 + \epsilon_{\ell-1})d_G(u_i, v_j) + d_G(v_j, y) + 4\delta \cdot d_G(x, y) \leq (1 + \epsilon_{\ell-1} + 4\delta)d_G(x, y)$ . By the choice of parameters  $4\delta = \epsilon$ , which gives  $d_{G \cup H^\ell}^\beta(x, y) \leq (1 + \epsilon_\ell)d_G(x, y)$ . This completes the proof.  $\square$

### 4.2.4 Conclusion

Using Claims 13 and 14 and Lemma 16, we obtain the following:

**Theorem 8** *Let  $0 < \epsilon < 1$ . There is a deterministic construction of a  $(\beta, \epsilon)$ -hopset with  $O(n^{3/2} \log n)$  edges and  $\beta = O((\log n)/\epsilon)$  that takes  $O((\log^2 n)/\epsilon)$  rounds in the CONGESTED CLIQUE model.*

**Proof** By Lemma 16,  $H^{\log n}$  is a  $(\beta, \epsilon_{\log n}, n)$ -hopset, and hence also a  $(\beta, \epsilon_{\log n})$ -hopset. From the choice of parameters

$\epsilon_{\log n} = \epsilon \log n$ , where  $0 < \epsilon < 1/\log n$  and  $\beta = O(1/\epsilon)$ . Let  $\epsilon' = \epsilon_{\log n} = \epsilon \log n$ . Then,  $\epsilon = \frac{\epsilon'}{\log n}$ . Hence,  $H^{\log n}$  is a  $(\beta, \epsilon')$ -hopset where  $0 < \epsilon' < 1$  and  $\beta = O(\frac{\log n}{\epsilon'})$ . By Claim 14, the time complexity for constructing the hopset is  $O(\beta \log n + \log^2 n)$ . Since  $\beta = O(\frac{\log n}{\epsilon'})$ , the complexity is  $O(\frac{\log^2 n}{\epsilon'})$  rounds. The number of edges is  $O(n^{3/2} \log n)$  by Claim 13.  $\square$

### 5 Multi-source shortest paths

As a direct application of our hopsets and source detection algorithms, we get an efficient  $(1 + \epsilon)$ -approximation for the multi-source shortest paths problem (MSSP) with polylogarithmic complexity as long as the number of sources is  $\tilde{O}(\sqrt{n})$ .

**Theorem 2** *There is a deterministic  $(1 + \epsilon)$ - approximation algorithm for the weighted undirected MSSP that takes*

$$O\left(\left(\frac{|S|^{2/3}}{n^{1/3}} + \log n\right) \cdot \frac{\log n}{\epsilon}\right)$$

*rounds in the CONGESTED CLIQUE, where  $S$  is the set of sources. In particular, the complexity is  $O(\frac{\log^2 n}{\epsilon})$  as long as  $|S| \leq O(\sqrt{n} \cdot (\log n)^{3/2})$ .*

**Proof** By Theorem 8, we can construct a  $(\beta, \epsilon)$ -hopset  $H$  with  $\beta = O(\frac{\log n}{\epsilon})$  in  $O(\frac{\log^2 n}{\epsilon})$  rounds. By the definition of hopsets, the  $\beta$ -hop distances in  $G \cup H$  give a  $(1 + \epsilon)$ -approximation to all the distances in  $G$ . Hence, to approximate the distances of all the nodes to the set of sources  $S$ , we run our  $(S, d, k)$ -algorithm in the graph  $G \cup H$  with  $d = \beta = O(\frac{\log n}{\epsilon})$ . By Theorem 6, the complexity is  $O\left(\left(\frac{n^{2/3}|S|^{2/3}}{n} + 1\right) \cdot \beta\right) = O\left(\left(\frac{|S|^{2/3}}{n^{1/3}} + 1\right) \cdot \frac{\log n}{\epsilon}\right)$  rounds. The overall complexity is

$$\begin{aligned} &O\left(\left(\frac{|S|^{2/3}}{n^{1/3}} + 1\right) \cdot \frac{\log n}{\epsilon} + \frac{\log^2 n}{\epsilon}\right) \\ &= O\left(\left(\frac{|S|^{2/3}}{n^{1/3}} + \log n\right) \cdot \frac{\log n}{\epsilon}\right) \end{aligned}$$

rounds.  $\square$

### 6 All-pairs shortest paths

In this section, we use our multi-source shortest paths algorithm from  $\tilde{O}(\sqrt{n})$  sources to construct algorithms that provide constant approximations for *all* the distances in the graph in *polylogarithmic* time. However, the approximation

guarantee increases to  $(2 + \epsilon)$  for the unweighted case, and at most  $(3 + \epsilon)$  for the weighted case.

Next, we present our approximation algorithms for the APSP problem. For simplicity of presentation, we start by sketching a simple  $(3 + \epsilon)$ -approximation for weighted APSP. Using a more careful analysis we show that a variant of this algorithm actually obtains a  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation. Here, we use the notation  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation for an algorithm that for any pair of nodes  $u, v$  provides an estimate of the distance between  $u$  and  $v$  which is at most  $(2 + \epsilon)d(u, v) + (1 + \epsilon)W$  where  $W$  is the weight of the heaviest edge in a shortest path between  $u$  and  $v$ . Note that  $W$  is always at most  $d(u, v)$ , and hence a  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation is always better than a  $(3 + \epsilon')$ -approximation (where  $\epsilon' = 2\epsilon$ ). Finally, we extend the algorithm to obtain a  $(2 + \epsilon)$ -approximation for APSP in unweighted graphs. All our algorithms are deterministic.

#### 6.1 $(3 + \epsilon)$ -approximation for weighted APSP

Our algorithm starts by computing a hitting set  $A$  of size  $\tilde{O}(\sqrt{n})$  that hits all the sets  $N_k(v)$  for  $k = \tilde{O}(\sqrt{n})$ . We compute for each node  $v \in V$ , the distances to all the nodes in  $N_k(v)$ , and we use our MSSP algorithm to compute  $(1 + \epsilon)$ -approximate distances from  $A$  to all the nodes. Since  $k = \tilde{O}(\sqrt{n})$  and  $|A| = \tilde{O}(\sqrt{n})$  the complexity is polylogarithmic. To compute an estimate  $\delta(u, v)$  of the distance  $d(u, v)$ , we work as follows. If  $v \in N_k(u)$ , then  $u$  already computed the value  $d(u, v)$  and sends it to  $v$ . Otherwise, let  $p(u) \in N_k(u) \cap A$  be a closest node to  $u$ . We estimate  $d(u, v)$  with  $d(u, p(u)) + d(p(u), v)$ . Since  $v \notin N_k(u)$ , it follows that  $d(u, p(u)) \leq d(u, v)$ . Hence, using the triangle inequality, we have

$$\begin{aligned} &d(u, p(u)) + d(p(u), v) \\ &\leq d(u, p(u)) + d(u, p(u)) + d(u, v) \leq 3d(u, v). \end{aligned}$$

Since our algorithm computes a  $(1 + \epsilon)$ -approximation for the value  $d(p(u), v)$  and computes exactly the value  $d(u, p(u))$ , we get a  $(3 + \epsilon)$ -approximation.

#### 6.2 $(2 + \epsilon, (1 + \epsilon)W)$ -approximation for weighted APSP

We next describe our  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation. In our  $(3 + \epsilon)$ -approximation, we used the fact that if  $v \notin N_k(u)$ , then  $d(u, p(u)) \leq d(u, v)$ . To obtain a better approximation, we would like to get a better bound on  $d(u, p(u))$ . If, for example,  $d(u, p(u)) \leq \frac{1}{2}d(u, v)$ , the same analysis shows a  $(2 + \epsilon)$ -approximation. We show that either  $\min\{d(u, p(u)), d(v, p(v))\} \leq \frac{1}{2}(d(u, v) + W)$  which proves a  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation, or the shortest path

between  $u$  and  $v$  has a node  $w \in N_k(u) \cap N_k(v)$ . In the latter case,  $d(u, w) + d(w, v) = d(u, v)$ , and since  $|N_k(u)| = |N_k(v)| = \tilde{O}(\sqrt{n})$  we can use matrix multiplication to compute the distances through all nodes  $w \in N_k(u) \cap N_k(v)$ , and find the actual distance  $d(u, v)$ .

### 6.2.1 Algorithm description

We next describe the algorithm in detail. We start by describing the general structure of the algorithm, and then explain how we implement it efficiently. In the algorithm, each node  $v \in V$ , maintains an estimate  $\delta(u, v)$  of the distance  $d(u, v)$  for all  $u \in V$ . Each time a node updates the estimate  $\delta(u, v)$  it takes the minimum value between the previous value of  $\delta(u, v)$  and the current estimate computed. Also, if  $v$  updates  $\delta(u, v)$  it sends the new estimate also to  $u$  to update its estimate accordingly. We use the notation  $w(u, v)$  to denote the weight of the edge  $\{u, v\}$  if exists. The distances  $d(u, v)$  and weights  $w(u, v)$  are with respect to the input graph  $G$ . The algorithm works as follows.

1. Set  $\delta(u, v) = w(u, v)$  if  $\{u, v\} \in E$ , and  $\delta(u, v) = \infty$  otherwise.
2. For each  $v$ , compute the distances to the set  $N_k(v)$  for  $k = \sqrt{n}$ , and update the estimates  $\{\delta(u, v)\}_{u \in N_k(v)}$  accordingly.
3. Set  $\delta(u, v) = \min\{\delta(u, v), \min_{w \in N_k(u) \cap N_k(v)} \{\delta(u, w) + \delta(w, v)\}\}$ .
4. Compute a hitting set  $A$  of size  $O(\sqrt{n} \log n)$  that hits all the sets  $N_k(v)$ .
5. Use the MSSP algorithm to compute  $(1 + \epsilon)$ -approximate distances from  $V$  to  $A$ . For each  $v \in V$  update the estimates  $\{\delta(v, u)\}_{u \in A}$  accordingly.
6. Let  $p(v) \in N_k(v) \cap A$  be a node from  $A$  closest to  $v$  that is chosen by  $v$ . The node  $v$  sends  $p(v)$  to all the nodes.
7. Set  $\delta(u, v) = \min\{\delta(u, v), \delta(u, p(u)) + \delta(p(u), v), \delta(v, p(v)) + \delta(p(v), u)\}$ .

### 6.2.2 Complexity

**Lemma 17** *The complexity of the algorithm is  $O(\frac{\log^2 n}{\epsilon})$  rounds.*

**Proof** Implementing Line 1 is immediate. We implement Line 2 using our  $k$ -nearest algorithm, which takes  $O\left(\left(\frac{k}{n^{2/3}} + \log n\right) \log k\right) = O(\log^2 n)$  rounds by Theorem 5. We implement Line 3 with our distance through sets algorithm, with respect to the sets  $N_k(v)$ . Since  $k = \sqrt{n}$ , it takes  $O\left(\frac{k^{2/3}}{n^{1/3}} + 1\right) = O(1)$  rounds by Theorem 7. We use Lemma 1 to compute the hitting set  $A$  in  $O((\log \log n)^3)$  rounds (note that all the nodes already computed the sets  $N_k(v)$ ). Each node updates all the nodes if it is in  $A$  which

takes one round. We implement Line 5 using our MSSP algorithm from Theorem 2, which takes  $O(\frac{\log^2 n}{\epsilon})$  rounds since  $|A| = O(\sqrt{n} \log n)$ . Implementing Line 6 is immediate since each node  $v$  knows  $N_k(v)$  and  $A$ . Then, all the nodes know all the values  $\{p(v)\}_{v \in V}$ . Now, each node  $u$  sends to each node  $v$  the values  $\{d(u, p(u)), d(u, p(v))\}$ . This allows  $u$  and  $v$  to compute the values  $\delta(u, p(u)) + \delta(p(u), v)$  and  $\delta(v, p(v)) + \delta(p(v), u)$ , which allows them to compute the final estimate  $\delta(u, v)$ . The overall complexity is  $O(\frac{\log^2 n}{\epsilon})$  rounds.  $\square$

### 6.2.3 Correctness

**Lemma 18** *The algorithm gives a  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation for weighted APSP.*

**Proof** Let  $u, v \in V$ . If  $v \in N_k(u)$  or  $u \in N_k(v)$ , then after Line 2,  $\delta(u, v) = d(u, v)$  and we are done. Otherwise, let  $P$  be a shortest path between  $u$  and  $v$ . Let  $u'$  be the furthest node from  $u$  in  $P \cap N_k(u)$ , and let  $v'$  be the furthest node from  $v$  in  $P \cap N_k(v)$ . We divide the analysis to 3 cases. See Fig. 3 for an illustration.

**Case 1:** Either  $u' \in N_k(v)$  or  $v' \in N_k(u)$ . In this case,  $\min_{w \in N_k(u) \cap N_k(v)} \{\delta(u, w) + \delta(w, v)\} = d(u, v)$ , hence after Line 3,  $\delta(u, v) = d(u, v)$  and we are done.


**Case 2:** There is a node  $w \in P$  where  $w \notin N_k(u)$  and  $w \notin N_k(v)$ . In this case, it follows that  $d(u, p(u)) \leq d(u, w)$  and  $d(v, p(v)) \leq d(v, w)$ . Since  $w \in P$ , we also have  $d(u, v) = d(u, w) + d(v, w)$ . Hence

$$\begin{aligned} \min\{d(u, p(u)), d(v, p(v))\} &\leq \\ \min\{d(u, w), d(v, w)\} &\leq d(u, v)/2. \end{aligned}$$

Assume w.l.o.g that  $d(u, p(u)) \leq d(v, p(v))$ , then  $d(u, p(u)) \leq d(u, v)/2$  and  $d(v, p(v)) \leq d(v, u) + d(u, p(u)) \leq \frac{3}{2}d(u, v)$ , which gives  $d(u, p(u)) + d(p(u), v) \leq 2d(u, v)$ . The value  $d(u, p(u))$  is computed exactly in Line 2, and the value  $d(v, p(v))$  is approximated within a  $(1 + \epsilon)$  factor in Line 5. Hence, the final estimate is  $\delta(u, v) \leq (2 + \epsilon')d(u, v)$  for  $\epsilon' = 2\epsilon$ .

**Case 3:** All the nodes in  $P$  are either in  $N_k(u)$  or  $N_k(v)$ , but  $N_k(u) \cap N_k(v) = \emptyset$ . In this case, the path  $P$  is composed of a path between  $u$  and  $u'$ , an edge between  $u'$  and  $v'$  and a path between  $v'$  and  $v$ . Since  $v' \notin N_k(u)$  and  $u' \notin N_k(v)$ , we have  $d(u, p(u)) \leq d(u, u')$  and  $d(v, p(v)) \leq d(v, v')$ . In addition, since  $d(u, v) = d(u, u') + w(u', v') + d(v', v)$ , we have

$$\begin{aligned} d(u, v') + d(v, u') &= \\ &= d(u, u') + w(u', v') + d(v, v') + w(u', v') \\ &= d(u, v) + w(u', v'). \end{aligned}$$



**Fig. 3** There are 3 cases, the first is that  $u'$  or  $v'$  are in  $N_k(u) \cap N_k(v)$ . The second is that there is a node  $w \in P \setminus (N_k(u) \cup N_k(v))$  and the last where there is an edge of  $P$  between  $u'$  and  $v'$

This implies that

$$\begin{aligned} & \min\{d(u, p(u)), d(v, p(v))\} \\ & \leq \min\{d(u, v'), d(v, u')\} \leq \frac{1}{2}(d(u, v) + w(u', v')). \end{aligned}$$

Assume w.l.o.g that  $d(u, p(u)) \leq d(v, p(v))$ , then

$$\begin{aligned} & d(u, p(u)) + d(p(u), v) \\ & \leq d(u, p(u)) + d(p(u), u) + d(u, v) \\ & \leq 2d(u, v) + w(u', v') \leq 2d(u, v) + W. \end{aligned}$$

Since we compute the value  $d(u, p(u))$  exactly, and approximate the value  $d(p(u), v)$  within a  $(1 + \epsilon)$  factor, the final estimate is  $\delta(u, v) \leq (2 + 2\epsilon)d(u, v) + (1 + \epsilon)W$ . Choosing  $\epsilon' = 2\epsilon$  completes the proof.  $\square$

### 6.2.4 Conclusion

From Lemmas 17 and 18, we get the following.

**Theorem 9** *There is a deterministic  $(2 + \epsilon, (1 + \epsilon)W)$ -approximation algorithm for weighted APSP in the CONGESTED CLIQUE model that takes  $O(\frac{\log^2 n}{\epsilon})$  rounds.*

### 6.3 $(2 + \epsilon)$ -approximation for unweighted APSP

In this section, we present our  $(2 + \epsilon)$ -approximation for unweighted APSP.

#### 6.3.1 High-level idea

In our  $(2 + \epsilon, 1 + \epsilon)$ -approximation, the only case where we do not obtain a  $(2 + \epsilon)$ -approximation is when the shortest path between  $u$  and  $v$  is composed of a path between  $u$  and  $u' \in N_k(u) \setminus N_k(v)$ , an edge between  $u'$  and a node  $v' \in N_k(v) \setminus N_k(u)$  and a path between  $v'$  and  $v$ . Note that  $u$  knows the distance to  $u'$ ,  $v$  knows the distance to  $v'$  and both  $u'$  and  $v'$  know about the edge between them. If we add to the graph all the edges between  $u$  and nodes in  $N_k(u)$  and between  $v$  and nodes in  $N_k(v)$ , we have a shortest path of length 3 between  $u$  and  $v$  in the new graph. Intuitively, we would like to use matrix multiplication to learn about this path, but we cannot

do this directly since this graph may be dense. In particular, we have no bound on the degrees of  $u', v'$ .

To overcome this, we present a new algorithm that takes care separately of shortest paths that include a high-degree node, and shortest paths that have only low-degree nodes, where a node has a high-degree if its degree is  $\tilde{O}(\sqrt{n})$ . This is done as follows. We compute a hitting set  $A$  of size  $\tilde{O}(\sqrt{n})$ , such that each high-degree node has a neighbour in  $A$ . Then, we use the fact that  $|A| = \tilde{O}(\sqrt{n})$ , to show that we can compute efficiently  $(1 + \epsilon)$ -approximations of all the shortest paths that go through a node in  $A$ . Then, if the shortest path between  $u$  and  $v$  has a high-degree node, it has a neighbour  $w \in A$ , and the shortest path from  $u$  to  $v$  that goes through  $w$  already gives a good approximation for the shortest  $u - v$  path.

Then, we are left only with paths that do not contain any high-degree node. The graph induced on low-degree nodes is sparse, since the degrees are at most  $\tilde{O}(\sqrt{n})$ , and we would like to exploit it. However, although we can multiply two sparse matrices with  $\tilde{O}(n^{3/2})$  elements efficiently, to compute paths of length 3 we need to multiply 3 matrices of this size, which is no longer efficient. To overcome this, we would like to reduce the degrees further. In particular, we would like to follow the analysis in Lemma 18, but replace  $N_k(v)$  with  $N_{k'}(v)$  for  $k' < k$  which allows focusing on sparser matrices. The obstacle is that we want to compute a hitting set  $A'$  of the sets  $\{N_{k'}(v)\}_{v \in V}$ , and then use our MSSP algorithm to compute distances to all the nodes in  $A'$ , but if  $k' < k$ , we need to compute a larger hitting set  $A'$ , and our MSSP algorithm is efficient only as long as  $|A'| = \tilde{O}(\sqrt{n})$ .

Nevertheless, since our graph is sparse, we show that we can actually compute MSSP from a set  $A'$  of size  $\tilde{O}(n^{3/4})$  which allows us to choose  $k' = \tilde{O}(n^{1/4})$ . To deal with the problematic case in the proof of Lemma 18, we should compute paths of length 3 of the form  $\{u, u', v', v\}$  where  $u' \in N_{k'}(u)$ ,  $v' \in N_{k'}(v)$  and there is an edge between  $u'$  and  $v'$ . Since  $|N_{k'}(u)| = |N_{k'}(v)| = \tilde{O}(n^{1/4})$  and  $u', v'$  are in the sparse graph and have maximum degree  $\tilde{O}(\sqrt{n})$ , we can detect paths of this form efficiently, which allows us to show a  $(2 + \epsilon)$ -approximation.

### 6.3.2 Algorithm description

We next present a pseudo-code of the algorithm, then we show how to implement it efficiently and provide a correctness proof. Again, each node  $v \in V$ , maintains an estimate  $\delta(u, v)$  of the distance  $d(u, v)$  for all  $u \in V$ . Each time a node updates the estimate  $\delta(u, v)$  it takes the minimum value between the previous value of  $\delta(u, v)$  and the current estimate computed. Also, if  $v$  updates  $\delta(u, v)$  it sends the new estimate also to  $u$  to update its estimate accordingly. Let  $N(v) = \{u \in V \mid \{v, u\} \in E\} \cup \{v\}$ . Our algorithm works as follows.

1. Set  $\delta(u, v) = w(u, v)$  if  $\{u, v\} \in E$  and  $\delta(u, v) = \infty$  otherwise.

**First phase - handling shortest paths with a high-degree node:**

2. Let  $A$  be a hitting set of size  $\tilde{O}(\sqrt{n})$  that hits all the neighbourhoods  $N(v)$  of high-degree nodes with  $|N(v)| \geq k$  for  $k = \tilde{O}(\sqrt{n})$ .
3. Use the MSSP algorithm to compute  $(1+\epsilon)$ -approximate distances from  $V$  to  $A$ . For each  $v \in V$  update the estimates  $\{\delta(v, u)\}_{u \in A}$  accordingly.
4. Set  $\delta(u, v) = \min\{\delta(u, v), \min_{w \in A} \{\delta(u, w) + \delta(w, v)\}\}$ .

**Second phase - handling shortest paths that contain only low-degree nodes:**

Let  $G' = (V', E')$  be the graph induced on low-degree nodes with degree at most  $k$ , from now on our algorithm works only in the graph  $G'$ . In particular, the set  $N_{k'}(v)$  is the set of  $k'$  closest nodes to  $v$  in  $G'$ .

5. For each  $v$ , compute the distances to the set  $N_{k'}(v)$  for  $k' = \tilde{O}(n^{1/4})$ , and update the estimates  $\{\delta(v, u)\}_{u \in N_{k'}(v)}$  accordingly.
6. Set  $\delta(u, v) = \min\{\delta(u, v), \min_{w \in N_{k'}(u) \cap N_{k'}(v)} \{\delta(u, w) + \delta(w, v)\}\}$ .
7. Construct a hitting set  $A'$  of size  $\tilde{O}(n^{3/4})$  that hits all the sets  $N_{k'}(v)$ .
8. Compute  $(1+\epsilon)$ -approximate shortest paths from  $A'$  to all the nodes in the graph  $G'$ . For each  $v \in V$  update the estimates  $\{\delta(v, u)\}_{u \in A'}$  accordingly.
9. Let  $p'(v) \in N_{k'}(v) \cap A'$  be a node from  $A'$  closest to  $v$  that is chosen by  $v$ . The node  $v$  sends  $p'(v)$  to all the nodes.
10. Set  $\delta(u, v) = \min\{\delta(u, v), \delta(u, p'(u)) + \delta(p'(u), v), \delta(v, p'(v)) + \delta(p'(v), u)\}$ .
11. Let  $\delta'(u, v) = \min\{\delta(u, u') + \delta(u', v') + \delta(v', v) \mid u' \in N_{k'}(u), v' \in N_{k'}(v), \{u', v'\} \in E'\}$ .
12. Set  $\delta(u, v) = \min\{\delta(u, v), \delta'(u, v)\}$ .

### 6.3.3 Complexity

**Lemma 19** *The complexity of the algorithm is  $O(\frac{\log^2 n}{\epsilon})$  rounds.*

**Proof** Implementing Line 1 is immediate. We choose  $k = \sqrt{n}$ . To implement Line 2, we use Lemma 1 to build  $A$  of size  $O(n \log n/k) = O(\sqrt{n} \log n)$  that hits all the sets  $N(v)$  for  $v$  with  $|N(v)| \geq k$ . Note that each node  $v$  knows the set  $N(v)$ . The complexity is  $O((\log \log n)^3)$  rounds. Each node updates all the nodes if it is in  $A$  which takes one round. We implement Line 3 using our MSSP algorithm from Theorem 2, which takes  $O(\frac{\log^2 n}{\epsilon})$  rounds since  $|A| = O(\sqrt{n} \log n)$ . Line 4 is equivalent to computing distances through the set  $A$  with respect to the distances  $\{\delta(u, v)\}_{u \in V, v \in A}$  computed in Line 3. This takes  $O\left(\frac{(\sqrt{n} \log n)^{2/3}}{n^{1/3}} + 1\right) = O(\log n)$  rounds by Theorem 7.

To implement Line 5, we choose  $k' = n^{1/4}$ , and use Theorem 5 to find the  $k'$ -nearest nodes in  $O\left(\left(\frac{k'}{n^{2/3}} + \log n\right) \log k'\right) = O(\log^2 n)$  rounds. To implement Line 6, we use our distance through sets algorithm, with respect to the sets  $N_{k'}(v)$ . Since  $k' = n^{1/4}$ , the complexity is  $O(1)$  by Theorem 7. To implement Line 7, we use Lemma 1 to build  $A'$  of size  $O(n \log n/k') = O(n^{3/4} \log n)$  that hits all the sets  $N_{k'}(v)$  for  $v \in G'$ . Note that all the nodes already computed the sets  $N_{k'}(v)$ . The complexity is  $O((\log \log n)^3)$  rounds. Each node updates all the nodes if it is in  $A'$  which takes one round. To implement Line 8, we use a sparse variant of our MSSP algorithm, as follows. First, we compute a  $(\beta, \epsilon)$ -hopset  $H$  of size  $O(n^{3/2} \log n)$  for  $G'$  using Theorem 8. The complexity is  $O(\frac{\log^2 n}{\epsilon})$  rounds and  $\beta = O(\frac{\log n}{\epsilon})$ . Now, the  $\beta$ -hop distances in  $G' \cup H$  give  $(1+\epsilon)$ -approximation for the distances in  $G'$ . Hence, to approximate the distances of all nodes from  $A'$ , we use our  $(S, d, k'')$ -algorithm in the graph  $G' \cup H$ , with  $S = A'$ ,  $k'' = |A'| = O(n^{3/4} \log n)$ ,  $d = \beta$ . In  $G'$  the maximum degree is  $k = \sqrt{n}$ , and  $H$  is of size  $O(n^{3/2} \log n)$ , hence the graph  $G' \cup H$  is of size  $O(n^{3/2} \log n)$ . Thus, by Theorem 6, the complexity is  $O\left(\left(\frac{(n^{3/2} \log n)^{1/3} \cdot (n^{3/4} \log n)^{2/3}}{n} + 1\right) \frac{\log n}{\epsilon}\right) = O(\frac{\log^2 n}{\epsilon})$ .

Implementing Line 9 takes one round since each node  $v$  knows  $N_{k'}(v)$  and  $A'$ . To implement Line 10 it is enough that each node  $v$  sends to each node  $u$  the values  $\{\delta(v, p'(v)), \delta(v, p'(u))\}$ . Finally, we explain how we implement Line 11. For this we multiply the following 3 matrices. The matrix  $M_1(u, w) = \delta(u, w)$  if  $w \in N_{k'}(u)$  and  $M_1(u, w) = \infty$  otherwise. The matrix  $M_3 = M_1^T$ . I.e.,  $M_3(w, u) = M_1(u, w)$ , and the matrix  $M_2(u, v) = \delta(u, v)$  if  $\{u, v\} \in E'$  and  $M_2(u, v) = \infty$  otherwise. Note that  $M_1, M_3$  have  $O(n^{5/4})$  entries which are not  $\infty$ , since the sizes of the sets  $N_{k'}(v)$  are  $n^{1/4}$ . Also, from the definition of  $G'$ , the matrix  $M_2$  has  $O(n^{3/2})$  entries which are not  $\infty$ . We multiply

$M_1 \cdot M_2 \cdot M_3$  in the min-plus semiring. For this, we first multiply  $M_1 \cdot M_2$  which takes  $O(1)$  rounds based on their sparsity, by Theorem 3. Then, we multiply the resulting matrix and  $M_3$ . Since the maximum degree in  $M_1$  is  $k' = n^{1/4}$  and the maximum degree in  $M_2$  is  $k = \sqrt{n}$ , we get that in  $M_1 \cdot M_2$  the maximum degree is  $n^{3/4}$ . Hence, the matrix  $M_1 \cdot M_2$  has at most  $n^{7/4}$  non- $\infty$  elements, and  $M_3$  has at most  $n^{5/4}$  non- $\infty$  elements. Hence multiplying  $(M_1 \cdot M_2) \cdot M_3$  takes  $O(1)$  rounds By Theorem 3 (choosing  $\widehat{\rho}_{ST} = n$ ). Since the multiplication is in the min-plus semiring and by the definition of the matrices, it holds that after the multiplication any pair of nodes  $u, v$  computed the value

$$\min_{\substack{u' \in N_{k'}(u), v' \in N_{k'}(v), \\ \{u', v'\} \in E'}} \{\delta(u, u') + \delta(u', v') + \delta(v', v)\}.$$

This completes the proof. □

### 6.3.4 Correctness

**Lemma 20** *By the end of the algorithm, for all  $u, v \in V$ , it holds that  $\delta(u, v) \leq (2 + \epsilon)d_G(u, v)$ .*

**Proof** If  $d_G(u, v) = 1$ , then after Line 1,  $\delta(u, v) = 1$  and we are done. Hence, we assume that  $d_G(u, v) \geq 2$ . Let  $P$  be a shortest path between  $u$  and  $v$  in  $G$ . We first consider the case that  $P$  has a high-degree node  $u'$  of degree at least  $k = \widetilde{O}(\sqrt{n})$ . Then,  $u'$  has a neighbour  $w \in A$ . For now, we use the notation  $d(u, v)$  for  $d_G(u, v)$ . From the triangle inequality,  $d(u, w) \leq d(u, u') + d(u', w) = d(u, u') + 1$  and  $d(v, w) \leq d(v, u') + 1$ . Also, since  $u' \in P$ , it holds that  $d(u, v) = d(u, u') + d(u', v)$ . Hence,  $d(u, w) + d(w, v) \leq d(u, u') + 1 + d(u', v) + 1 \leq d(u, v) + 2 \leq 2d(u, v)$ . In Line 3, we compute  $(1 + \epsilon)$ -approximations for the values  $d(u, w), d(w, v)$ , hence after Line 4, we have  $\delta(u, v) \leq (2 + 2\epsilon)d(u, v)$ .

We next consider the case that  $P$  has only low-degree nodes, hence it is contained in the graph  $G'$ . From now on we work with the graph  $G'$ , and use  $d(u, v)$  for  $d_{G'}(u, v)$ . Since  $P$  is contained in  $G'$ , it holds that  $d_{G'}(u, v) = d_G(u, v)$ . The proof is divided into 3 cases, as in the proof of Lemma 18.

**Case 1:** there is a node  $w \in P \cap (N_{k'}(u) \cap N_{k'}(v))$ . Then,  $d(u, v) = d(u, w) + d(w, v)$ . Also,  $u$  and  $v$  computed the values  $d(u, w), d(w, v)$  in Line 5, and  $\delta(u, v) = d(u, w) + d(w, v) = d_{G'}(u, v) = d_G(u, v)$  after Line 6.

**Case 2:** there is a node  $w \in P \setminus (N_{k'}(u) \cup N_{k'}(v))$ . Since  $w \in P$ ,  $p'(u) \in N_{k'}(u)$  and  $p'(v) \in N_{k'}(v)$ , we have  $d(u, p'(u)) + d(v, p'(v)) \leq d(u, w) + d(v, w) \leq d(u, v)$ . Exactly as in the proof of Case 2 in Lemma 18, we get that  $\min\{d(u, p'(u)) + d(p'(u), v), d(v, p'(v)) + d(p'(v), u)\} \leq 2d(u, v)$ . Since our algorithm computes  $d(u, p'(u)), d(v, p'(v))$  exactly in Line 5 and approximates

$d(u, p'(v)), d(v, p'(u))$  within  $(1 + \epsilon)$  factor in Line 8, after Line 10 we have  $\delta(u, v) \leq (2 + 2\epsilon)d(u, v)$ .

**Case 3:**  $P$  is composed of a path between  $u$  to a node  $u' \in N_{k'}(u)$ , an edge  $\{u', v'\} \in E'$  where  $v' \in N_{k'}(v)$  and a path between  $v'$  and  $v$ . In this case,  $\delta(u, u') = d(u, u')$  and  $\delta(v', v) = d(v, v')$  after Line 5, and  $\delta(u', v') = d(u, v)$  after Line 1. Hence, in Line 11 we set  $\delta(u, v) = \delta(u, u') + \delta(u', v') + \delta(v', v) = d(u, v)$ .

To conclude, in all cases we have  $\delta(u, v) \leq (2 + 2\epsilon)d_G(u, v)$  by the end of the algorithm. Choosing  $\epsilon' = 2\epsilon$  completes the proof. □

### 6.3.5 Conclusion

From Lemmas 19 and 20, we get the following.

**Theorem 10** *There is a deterministic  $(2 + \epsilon)$ -approximation algorithm for unweighted APSP in the CONGESTED CLIQUE model that takes  $O(\frac{\log^2 n}{\epsilon})$  rounds.*

## 7 Other applications

### 7.1 Exact single-source shortest paths

We show that our distance tools allow to get a faster algorithm for exact SSSP in undirected weighted graphs, improving the previous  $\widetilde{O}(n^{1/3})$ -round algorithm [13]. The algorithm is very simple, and is based on ideas used in [22,48,55]. We start by computing the distances to the nodes in  $N_k(v)$  for all  $v \in V$  using our  $k$ -nearest algorithm. Then, we add to the graph all the edges  $\{\{u, v\} | u \in V, v \in N_k(u)\}$  with the weights  $d(u, v)$  computed by the algorithm, which results in a new graph  $G'$ .  $G'$  is called the  $k$ -shortcut graph or  $k$ -shortcut hopset in [22,48], and it is proved in [48] that the shortest path diameter of  $G'$  is  $O(n/k)$ . I.e., for any pair of nodes  $u, v$  there is a shortest path between  $u$  and  $v$  of hop-distance at most  $O(n/k)$  in  $G'$ .

**Lemma 21** (Theorem 3.10 in [48]) *The shortest path diameter of  $G'$  is smaller than  $4n/k$ .*

Now, to compute SSSP from a certain node we just run the classic Bellman-Ford algorithm [9,25] in the graph  $G'$  which takes  $O(n/k)$  rounds in a graph with shortest path diameter  $O(n/k)$ , even in the more restricted CONGEST model (see e.g., [48]). To optimize the complexity, we choose  $k = n^{5/6}$ . Now, finding the  $k$ -nearest takes  $\widetilde{O}(k/n^{2/3}) = \widetilde{O}(n^{1/6})$  rounds by Theorem 5, and the Bellman-Ford exploration takes  $O(n^{1/6})$  rounds, proving the following.

**Theorem 11** *There is a deterministic algorithm for exact SSSP in undirected weighted graphs that takes  $\widetilde{O}(n^{1/6})$  rounds in the CONGESTED CLIQUE model.*

## 7.2 Diameter

We show that using our distance tools we can implement efficiently the algorithm of Roditty and Vassilevska Williams for approximating the diameter [54], which is an extension of the algorithm of Aingworth et al. [4]. We focus on an algorithm for unweighted graphs, and later discuss an extension to the weighted case.

The algorithm for approximating the diameter is based on computing BFS trees from  $\tilde{O}(\sqrt{n})$  nodes. A difference in our case is that we compute  $(1 + \epsilon)$ -approximate distances from the same nodes, and not the exact distances, which affects slightly the approximation obtained. Also, our algorithm works only for *undirected* graphs. Another difference from [54] is that we compute distances to the nodes in  $N_k(v)$ . This is useful in our case, since we can compute these distances *exactly* which is useful for the analysis, it also allows us to give a deterministic algorithm. The algorithm works as follows.

1. Each node  $v$  computes the distances to  $N_k(v)$  for  $k = \sqrt{n}$ .
2. We compute a hitting set  $S$  of size  $\tilde{O}(\sqrt{n})$  that hits all the sets  $\{N_k(v)\}_{v \in V}$ .
3. We compute  $(1 + \epsilon)$ -approximate distances from  $S$  to all the nodes.
4. Let  $p(v) \in S \cap N_k(v)$  be a closest node to  $v$  from  $S$ .  $v$  knows the distance to  $p(v)$  and sends to all the nodes the value  $d(v, p(v))$ .
5. Let  $w \in V$  be a node such that  $d(w, p(w)) \geq d(v, p(v))$  for all  $v \in V$ , we compute  $(1 + \epsilon)$ -approximate distances from the set  $N_k(w)$  (including  $w$ ) to all the nodes.
6. The estimate for the diameter is the maximum distance found in Line 3 or 5.

**Claim 22** *The algorithm takes  $O(\frac{\log^2 n}{\epsilon})$  rounds.*

**Proof** Let  $k = \sqrt{n}$ . Computing the distances to the  $k$ -nearest neighbours takes  $O(\log^2 n)$  rounds by Theorem 5. We compute a hitting set  $A$  of size  $O(\sqrt{n} \log n)$  that hits the sets  $N_k(v)$  using Lemma 1, and then each node tells all nodes if it is in  $S$  which takes 1 round. Line 3 is implemented using our MSSP algorithm that takes  $O(\frac{\log^2 n}{\epsilon})$  rounds since  $A$  is of size  $O(\sqrt{n} \log n)$ . The distance  $d(v, p(v))$  is already known to  $v$  since  $p(v) \in N_k(v)$ , and sending the values  $d(v, p(v))$  takes one round. After Line 4, all the nodes can deduce the node  $w$  (breaking ties by ids), since  $w$  knows the set  $N_k(w)$ , it can update all the nodes in  $N_k(w)$  that they are in this set, and in one additional round all the nodes learn  $N_k(w)$ . Then, implementing Line 5 is done using our MSSP algorithm which takes  $O(\frac{\log^2 n}{\epsilon})$  rounds since  $k = \sqrt{n}$ . Computing the maximum estimate takes one additional round. To conclude, the complexity is  $O(\frac{\log^2 n}{\epsilon})$  rounds.  $\square$

We next prove that the approximation returned by the algorithm is nearly  $3/2$ , following the proof in [54]. Since we compute  $(1 + \epsilon)$ -approximate distances, we may also get an estimate that is greater than  $D$  by a  $(1 + \epsilon)$  factor. If we want the estimate to be always at most  $D$  we can divide by  $1 + \epsilon$ .

**Claim 23** *Let  $G = (V, E)$  be an undirected graph with diameter  $D = 3h + z$ , where  $h \geq 0$  and  $z \in \{0, 1, 2\}$ , and let  $D'$  be the estimate returned by the algorithm. Then,  $2h + z \leq D' \leq (1 + \epsilon)D$  if  $z \in \{0, 1\}$ , and  $2h + 1 \leq D' \leq (1 + \epsilon)D$  if  $z = 2$ .*

**Proof** The analysis is based on the analysis in [54] (see Lemma 4). The difference in our algorithm compared to [54] is that in some cases we approximate distances instead of computing exact ones. However, all the distances  $d(v, p(v))$  are computed exactly, which shows that the definition of  $w$  in our algorithm is exactly as in [54]. Let  $a, b$  be such that  $d(a, b) = D$ , where  $D$  is the diameter of the graph. The analysis in [54] is divided to cases. In the first case,  $d(w, p(w)) \leq h$ , and they show that in this case there is a node at distance at least  $2h + z$  from  $p(a)$ . Our algorithm computes  $(1 + \epsilon)$ -approximate distances to all the nodes  $p(v)$ , hence at least one of our estimates is at least  $2h + z$  in this case.

A second case is that there is a node of distance at least  $2h + z$  from  $w$ . Since we compute approximate distances from  $w$ , one of our estimates is at least  $2h + z$  in this case. The last case to consider is that all the nodes are at distance smaller than  $2h + z$  from  $w$ , and  $d(w, p(w)) > h$ . In this case, they show that there exists a node  $w' \in N_k(w)$  such that  $d(a, w') \geq 2h + 1$ . Since we compute  $(1 + \epsilon)$ -approximate distances from all the nodes in  $N_k(w)$ , our estimate is at least  $2h + 1$ . In addition, since all the distances computed are  $(1 + \epsilon)$ -approximate distances, our estimate is at most  $(1 + \epsilon)D$ . This completes the proof.  $\square$

**Remark** As stated in [54], the same analysis works also for weighted graphs with non-negative weights with a loss of an additive  $W$  term where  $W$  is the maximum edge-weight. I.e., the estimate obtained satisfies  $\lfloor 2D/3 - W \rfloor \leq D' \leq (1 + \epsilon)D$ . The implementation is exactly the same in our case.

## 8 Discussion

In this paper, we provide new tools for distance computations in the CONGESTED CLIQUE model based on sparse MM. We demonstrate the power of these tools by providing efficient algorithms for many central problems, such as APSP, SSSP, and diameter approximation. We believe that these techniques may be useful for many additional tasks in distance computation in the CONGESTED CLIQUE, and possibly

also in additional related models. Many intriguing questions are still open.

First, while we provide constant approximations for many tasks in polylogarithmic time, the complexity of our exact SSSP algorithm is still polynomial. It would be interesting to study whether it is possible to obtain a polylogarithmic algorithm for the exact case. Surprisingly, even for the simple task of computing a BFS tree there is currently no sub-polynomial algorithm in the CONGESTED CLIQUE.

Second, while our MM algorithms work also for the directed case, our hopset construction is only for the undirected case and hence our results are for undirected graphs. Also, as we explain in Sect. 1.2, obtaining *any* approximation for directed APSP in sub-polynomial time would give a sub-polynomial algorithm for matrix multiplication. It would be interesting to study whether it is possible to provide sub-polynomial algorithms for directed SSSP.

Our unweighted APSP algorithm provides a  $(2 + \epsilon)$ -approximation, whereas in the weighted case there is currently an additional additive term, and a natural question is whether we can get a  $(2 + \epsilon)$ -approximation for weighted APSP in polylogarithmic time as well. Another interesting goal is to obtain *additive* approximations in sub-polynomial time.

**Acknowledgements** Open access funding provided by Institute of Science and Technology (IST Austria). We thank Mohsen Ghaffari, Michael Elkin and Merav Parter for fruitful discussions. This project has received funding from the European Union's Horizon 2020 Research And Innovation Program under Grant Agreement No. 755839.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Abboud, A., Bodwin, G., Pettie, S.: A hierarchy of lower bounds for sublinear additive spanners. *SIAM J. Comput.* **47**(6), 2203–2236 (2018)
2. Abboud, A., Censor-Hillel, K., Houry, S.: Near-linear lower bounds for distributed distance computations, even in sparse networks. *Proc. DISC* **2016**, 29–42 (2016)
3. Agarwal, U., Ramachandran, V., King, V., Pontecorvi, M.: A deterministic distributed algorithm for exact weighted all-pairs shortest paths in  $\tilde{O}(n^{3/2})$  rounds. *Proc. PODC* **2018**, 199–205 (2018). <https://doi.org/10.1145/3212734.3212773>
4. Aingworth, D., Chekuri, C., Indyk, P., Motwani, R.: Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.* **28**(4), 1167–1181 (1999)
5. Backurs, A., Roditty, L., Segal, G., Williams, V.V., Wein, N.: Towards tight approximation bounds for graph diameter and eccentricities. *Proc. STOC* **2018**, 267–280 (2018). <https://doi.org/10.1145/3188745.3188950>
6. Baswana, S., Goyal, V., Sen, S.: All-pairs nearly 2-approximate shortest-paths in  $O(n^2 \text{polylog } n)$  time. *Proc. STACS* **2005**, 666–679 (2005)
7. Baswana, S., Kavitha, T.: Faster algorithms for all-pairs approximate shortest paths in undirected graphs. *SIAM J. Comput.* **39**(7), 2865–2896 (2010)
8. Becker, R., Karrenbauer, A., Krinninger, S., Lenzen, C.: Near-optimal approximate shortest paths and transshipment in distributed and streaming models. In: *Proc. DISC* 2017 (2017). <https://doi.org/10.4230/LIPICs.DISC.2017.7>
9. Bellman, R.: On a routing problem. *Q. Appl. Math.* **16**(1), 87–90 (1958)
10. Berman, P., Kasiviswanathan, S.P.: Faster approximation of distances in graphs. *Proc. WADS* **2007**, 541–552 (2007)
11. Bernstein, A.: Fully dynamic  $(2 + \epsilon)$ -approximate all-pairs shortest paths with fast query and close to linear update time. *Proc. FOCS* **2009**, 693–702 (2009). <https://doi.org/10.1109/FOCS.2009.16>
12. Bernstein, A., Nanongkai, D.: Distributed exact weighted all-pairs shortest paths in near-linear time (2018). [ArXiv:1811.03337](https://arxiv.org/abs/1811.03337) [cs.DC]
13. Censor-Hillel, K., Kaski, P., Korhonen, J.H., Lenzen, C., Paz, A., Suomela, J.: Algebraic methods in the congested clique. In: *Proc. PODC* 2015, pp. 143–152. ACM (2015)
14. Censor-Hillel, K., Leitersdorf, D., Turner, E.: Sparse matrix multiplication and triangle listing in the congested clique model. In: *Proc. OPODIS* 2018 (2018). <https://doi.org/10.4230/LIPICs.OPODIS.2018.4>
15. Chang, Y.J., Pettie, S., Zhang, H.: Distributed triangle detection via expander decomposition. *Proc. SODA* **2018**, 821–840 (2018). <https://doi.org/10.1137/1.9781611975482.51>
16. Chechik, S., Larkin, D.H., Roditty, L., Schoenebeck, G., Tarjan, R.E., Williams, V.V.: Better approximation algorithms for the graph diameter. *Proc. SODA* **2014**, 1041–1052 (2014)
17. Cohen, E.: Polylog-time and near-linear work approximation scheme for undirected shortest paths. *J. ACM (JACM)* **47**(1), 132–166 (2000)
18. Cohen, E., Zwick, U.: All-pairs small-stretch paths. *J. Algorithm* **38**(2), 335–353 (2001)
19. Dolev, D., Lenzen, C., Peled, S.: “tri, tri again”: finding triangles and small subgraphs in a distributed setting. In: *Proc. DISC*, pp. 195–209 (2012)
20. Dor, D., Halperin, S., Zwick, U.: All-pairs almost shortest paths. *SIAM J. Comput.* **29**(5), 1740–1759 (2000)
21. Drucker, A., Kuhn, F., Oshman, R.: On the power of the congested clique model. *Proc. PODC* **2014**, 367–376 (2014)
22. Elkin, M.: Distributed exact shortest paths in sublinear time. In: *Proc. STOC* 2017, pp. 757–770. ACM (2017)
23. Elkin, M., Neiman, O.: Hopsets with constant hopbound, and applications to approximate shortest paths. *Proc. FOCS* **2016**, 128–137 (2016)
24. Elkin, M., Neiman, O.: Linear-size hopsets with small hopbound, and distributed routing with low memory (2017). [ArXiv:1704.08468](https://arxiv.org/abs/1704.08468) [cs.DS]
25. Ford Jr., L.R.: *Network flow theory*. Tech. Rep., Rand Corp Santa Monica, CA (1956)
26. Forster, S., Nanongkai, D.: A faster distributed single-source shortest paths algorithm. *Proc. FOCS* **2018**, 686–697 (2018). <https://doi.org/10.1109/FOCS.2018.00071>



27. Gall, F.L.: Powers of tensors and fast matrix multiplication. Proc. ISSAC **2014**, 296–303 (2014). <https://doi.org/10.1145/2608628.2608664>
28. Ghaffari, M.: Distributed MIS via all-to-all communication. Proc. PODC **2017**, 141–149 (2017). <https://doi.org/10.1145/3087801.3087830>
29. Ghaffari, M., Gouleakis, T., Konrad, C., Mitrovic, S., Rubinfeld, R.: Improved massively parallel computation algorithms for MIS, matching, and vertex cover. Proc. PODC **2018**, 129–138 (2018). <https://doi.org/10.1145/3212734.3212743>
30. Ghaffari, M., Li, J.: Improved distributed algorithms for exact shortest paths. Proc. STOC **2018**, 431–444 (2018). <https://doi.org/10.1145/3188745.3188948>
31. Ghaffari, M., Nowicki, K.: Congested clique algorithms for the minimum cut problem. In: Proc. PODC 2018, pp. 357–366. ACM (2018)
32. Ghaffari, M., Parter, M.: MST in log-star rounds of congested clique. Proc. PODC **2016**, 19–28 (2016). <https://doi.org/10.1145/2933057.2933103>
33. Hegeman, J.W., Pemmaraju, S.V.: Lessons from the congested clique applied to MapReduce. Theoret. Comput. Sci. **608**, 268–281 (2015)
34. Henzinger, M., Krinninger, S., Nanongkai, D.: A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. Proc. STOC **2016**, 489–498 (2016). <https://doi.org/10.1145/2897518.2897638>
35. Henzinger, M., Krinninger, S., Nanongkai, D.: Decremental single-source shortest paths on undirected graphs in near-linear total update time. J. ACM **65**(6), 36 (2018)
36. Holzer, S., Pinski, N.: Approximation of distances and shortest paths in the broadcast congest clique. In: Proc. OPODIS 2015 (2015). <https://doi.org/10.4230/LIPIcs.OPODIS.2015.6>
37. Holzer, S., Wattenhofer, R.: Optimal distributed all pairs shortest paths and applications. Proc. PODC **2012**, 355–364 (2012)
38. Huang, C., Nanongkai, D., Saranurak, T.: Distributed exact weighted all-pairs shortest paths in  $\tilde{O}(n^{5/4})$  rounds. Proc. FOCS **2017**, 168–179 (2017). <https://doi.org/10.1109/FOCS.2017.24>
39. Huang, S.E., Pettie, S.: Thorup-zwick emulators are universally optimal hopsets. Inf. Process. Lett. **142**, 9–13 (2019)
40. Jurdziński, T., Nowicki, K.: MST in  $O(1)$  rounds of congested clique. Proc. SODA **2018**, 2620–2632 (2018)
41. Korhonen, J.H., Suomela, J.: Towards a complexity theory for the congested clique. Proc. SPAA **2018**, 163–172 (2018). <https://doi.org/10.1145/3210377.3210391>
42. Le Gall, F.: Further algebraic algorithms in the congested clique model and applications to graph-theoretic problems. Proc. DISC **2016**, 57–70 (2016)
43. Lenzen, C.: Optimal deterministic routing and sorting on the congested clique. Proc. PODC **2013**, 42–50 (2013)
44. Lenzen, C., Patt-Shamir, B.: Fast routing table construction using small messages. In: Proc. STOC 2013, pp. 381–390. ACM (2013)
45. Lenzen, C., Patt-Shamir, B.: Fast partial distance estimation and applications. In: Proc. PODC 2015, pp. 153–162. ACM (2015)
46. Lenzen, C., Peleg, D.: Efficient distributed source detection with limited bandwidth. Proc. PODC **2013**, 375–382 (2013). <https://doi.org/10.1145/2484239.2484262>
47. Lotker, Z., Patt-Shamir, B., Pavlov, E., Peleg, D.: Minimum-weight spanning tree construction in  $o(\log \log n)$  communication rounds. SIAM J. Comput. **35**(1), 120–131 (2005)
48. Nanongkai, D.: Distributed approximation algorithms for weighted shortest paths. In: Proc. STOC 2014, pp. 565–573. ACM (2014)
49. Pandurangan, G., Robinson, P., Scquizzato, M.: On the distributed complexity of large-scale graph computations. Proc. SPAA **2018**, 405–414 (2018). <https://doi.org/10.1145/3210377.3210409>
50. Parter, M.:  $(\Delta + 1)$  coloring in the congested clique model. In: Proc. ICALP 2018, pp. 160:1–160:14 (2018). <https://doi.org/10.4230/LIPIcs.ICALP.2018.160>
51. Parter, M., Su, H.: Randomized  $(\Delta + 1)$ -coloring in  $O(\log^* \Delta)$  congested clique rounds. In: Proc. DISC 2018, pp. 39:1–39:18 (2018). <https://doi.org/10.4230/LIPIcs.DISC.2018.39>
52. Parter, M., Yogev, E.: Congested clique algorithms for graph spanners. In: Proc. DISC 2018, pp. 40:1–40:18 (2018). <https://doi.org/10.4230/LIPIcs.DISC.2018.40>
53. Peleg, D., Roditty, L., Tal, E.: Distributed algorithms for network diameter and girth. Proc. ICALP **2012**, 660–672 (2012)
54. Roditty, L., Vassilevska Williams, V.: Fast approximation algorithms for the diameter and radius of sparse graphs. In: Proc. STOC 2013, pp. 515–524. ACM (2013)
55. Shi, H., Spencer, T.H.: Time-work tradeoffs of the single-source shortest paths problem. J. Algorithms **30**(1), 19–32 (1999)
56. Thorup, M., Zwick, U.: Spanners and emulators with sublinear distance errors. Proc. SODA **2006**, 802–809 (2006)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.