


An Improved Random Shift Algorithm for Spanners and Low Diameter Decompositions

Sebastian Forster  

University of Salzburg, Austria

Martin Grösbacher 

University of Salzburg, Austria

Tijn de Vos  

University of Salzburg, Austria

Abstract

Spanners have been shown to be a powerful tool in graph algorithms. Many spanner constructions use a certain type of clustering at their core, where each cluster has small diameter and there are relatively few spanner edges between clusters. In this paper, we provide a clustering algorithm that, given $k \geq 2$, can be used to compute a spanner of stretch $2k - 1$ and expected size $O(n^{1+1/k})$ in k rounds in the CONGEST model. This improves upon the state of the art (by Elkin, and Neiman [TALG'19]) by making the bounds on both running time and stretch independent of the random choices of the algorithm, whereas they only hold with high probability in previous results. Spanners are used in certain synchronizers, thus our improvement directly carries over to such synchronizers. Furthermore, for keeping the *total* number of inter-cluster edges small in low diameter decompositions, our clustering algorithm provides the following guarantees. Given $\beta \in (0, 1]$, we compute a low diameter decomposition with diameter bound $O\left(\frac{\log n}{\beta}\right)$ such that each edge $e \in E$ is an inter-cluster edge with probability at most $\beta \cdot w(e)$ in $O\left(\frac{\log n}{\beta}\right)$ rounds in the CONGEST model. Again, this improves upon the state of the art (by Miller, Peng, and Xu [SPAA'13]) by making the bounds on both running time and diameter independent of the random choices of the algorithm, whereas they only hold with high probability in previous results.

2012 ACM Subject Classification Theory of computation \rightarrow Sparsification and spanners; Theory of computation \rightarrow Distributed algorithms

Keywords and phrases Spanner, low diameter decomposition, synchronizer, distributed graph algorithms

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2021.16

Related Version *Full Version:* <https://arxiv.org/abs/2111.08975>

Funding Supported by the Austrian Science Fund (FWF): P 32863-N.

Acknowledgements The first author would like to thank Merav Parter for mentioning that there was still some room for improvement in the spanner construction.

1 Introduction

Clustering has become an essential tool in dealing with large data sets. The goal of clustering data is to identify disjoint, dense regions such that the space between them is sparse. When working with graphs, this translates to partitioning the vertex set into clusters with relatively few edges between clusters such that the clusters satisfy a particular property. One can for example demand that the clusters have low diameter [4, 2, 9, 31], high conductance [22, 24, 36, 12, 13, 35, 14], or low effective resistance diameter [1]. In this paper, we focus on the low diameter decomposition and its connection to spanners. Low diameter decompositions are formally defined as follows.



© Sebastian Forster, Martin Grösbacher, and Tijn de Vos;
licensed under Creative Commons License CC-BY 4.0

25th International Conference on Principles of Distributed Systems (OPODIS 2021).

Editors: Quentin Bramas, Vincent Gramoli, and Alessia Milani; Article No. 16; pp. 16:1–16:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

16:2 Spanners and Low Diameter Decompositions

► **Definition 1.** Let $G = (V, E)$ be a weighted graph. A probabilistic (β, δ) -low diameter decomposition of G is a partition of the vertex set V into subsets V_1, \dots, V_l , called clusters, such that

- each cluster V_i has strong diameter at most δ , i.e., $d_{G[V_i]}(u, v) \leq \delta$ for all $u, v \in V_i$ ¹;
- the probability that an edge $e \in E$ is an inter-cluster edge is at most $\beta \cdot w(e)$, i.e., for $e = (u, v)$, the probability that $u \in V_i$ and $v \in V_j$ for $i \neq j$ is at most $\beta \cdot w(u, v)$.

In an unweighted graph, another typical definition of the low diameter decomposition replaces the second condition with an upper bound on the number of inter-cluster edges [31]. In this fashion, a probabilistic low diameter decomposition has $O(\beta m)$ inter-cluster edges in expectation.

Originally, low diameter decompositions were developed for distributed models, where they have been proven useful by reducing communication significantly in certain situations [4, 6]. Later, they also have shown to be fruitful in other models; they have been applied in shortest path approximations [15], cut sparsifiers [27], and tree embeddings with low stretch [2, 9, 10].

The clustering technique used for computing low diameter decompositions has implicitly been used to develop sparse spanners [11, 30, 17] and synchronizers [4, 7]. The main idea is to create the clusters, and add some, but not all, of the inter-cluster edges. In a sense, the inter-cluster edges are sparsified. We formalize this concept as follows.

► **Definition 2.** Let $G = (V, E)$ be an unweighted graph. A sparsified (ζ, δ) -low diameter decomposition of G is a partition of the vertex set V into subsets V_1, \dots, V_l , called clusters, together with a set of edges $F \subseteq E$ such that

- each cluster V_i has strong diameter at most δ , i.e., $d_{G[V_i]}(u, v) \leq \delta$ for all $u, v \in V_i$;
- for every edge, one of its endpoints has an edge from F into the cluster of the other endpoint, i.e., $\forall e = (u, v) \in E$, we have either $(u', v) \in F$ for some $u' \in C_u$ or $(u, v') \in F$ for some $v' \in C_v$ ².
- $|F| \leq \zeta$.

Moreover, we say that a sparsified (ζ, δ) -low diameter decomposition is tree-supported if for each cluster V_i we have a cluster center $c_i \in V_i$ and a tree of height at most $\delta/2$ spanning the cluster. All these trees together are called the support-forest.

Our main result is a clustering algorithm that produces a sparsified low diameter decomposition.

► **Theorem 3.** There exists an algorithm, such that for each unweighted graph $G = (V, E)$ and parameter $k \geq 2$ it outputs a tree-supported sparsified $(\zeta, 2k - 2)$ -low diameter decomposition, with $\zeta = O(n^{1+1/k})$ in expectation. The algorithm runs in k rounds in the CONGEST model, and in $O(k \log^* n)$ depth and $O(m)$ work in the PRAM model.

An important feature of this result is that the bounds on the strong diameter and number of rounds are not probabilistic; they are independent of the random choices in the algorithm. We show two applications of this theorem: constructing spanners and constructing synchronizers.

¹ For $U \subseteq V$, we write $G[U]$ for the graph induced by U , i.e., $G[U] := (U, U \times U \cap E)$.

² For $v \in V$, we write C_v for the cluster containing v .

Spanners. Given a graph $G = (V, E)$, we say that $H \subseteq G$ is a *spanner* of stretch α , if $d_H(u, v) \leq \alpha \cdot d_G(u, v)$, for every $u, v \in V$. It is straightforward that a tree-supported sparsified $(\zeta, 2k - 2)$ -low diameter decomposition gives a spanner of size $\zeta + n$ and stretch $2k - 1$, for details we refer to Section 3.1. This gives us the following corollary.

► **Corollary 4.** *There exists an algorithm, such that for each unweighted graph $G = (V, E)$ and parameter $k \geq 2$ it outputs a spanner H of stretch $2k - 1$. The expected size of H is at most $O(n^{1+1/k})$. The algorithm runs in k rounds in the CONGEST model, and in $O(k \log^* n)$ depth and $O(m)$ work in the PRAM model.*

Spanners themselves have been useful in computing approximate shortest paths [5, 16], distance oracles and labeling schemes [38, 32], and routing [33]. A simple greedy algorithm [3] gives a spanner of stretch $2k - 1$ and of size $O(n^{1+1/k})$, which is an optimal trade-off under the girth conjecture [18]. However, its fastest known implementation in the RAM model takes $O(kn^{2+1/(2k+1)})$ time [34]. Halperin and Zwick [23] gave a linear-time algorithm to construct spanners with an optimal trade-off for unweighted graphs in the RAM model. However, this algorithm does not adapt well to distributed and parallel models of computation. This problem can be overcome by exploiting the aforementioned relation with sparsified low diameter decompositions. This was (implicitly) done by Baswana and Sen [11], who provide an algorithm that computes a spanner of stretch $2k - 1$ and of size $O(kn^{1+1/k})$ in $O(k)$ rounds. The state of the art is by Elkin and Neiman [17], which builds off [30], and is also based on low diameter decompositions. They provide an algorithm that with probability $1 - 1/c$ computes a $(2k - 1)$ -spanner of expected size $O(c^{1/k}n^{1+1/k})$ in k rounds. Standard techniques of boosting the failure probability to something inverse polynomial (or “with high probability”) will require a logarithmic overhead. Alternatively, one can view the algorithm of Elkin and Neiman as an algorithm that outputs an α -spanner of expected size $O(c^{1/k}n^{1+1/k})$ in $O(\alpha)$ rounds, such that with probability $1 - 1/c$ we have that $\alpha = 2k - 1$.

Corollary 4 improves on the result of Elkin and Neiman by making the bounds on the stretch and the running time independent of the random choices in the algorithm. In particular, the algorithm of Elkin-Neiman involves sampling vertex values from an exponential distribution. The exponential distribution introduces an (as we show) unnecessary amount of randomness; we demonstrate that the geometric distribution suffices. We replace the extra random bits the exponential distribution provides by a tie-breaking rule on the vertex IDs, which we believe contributes to a more intuitive construction.

Synchronizers. The second application of Theorem 3 is in constructing synchronizers in the CONGEST model. A synchronizer gives a procedure to run a synchronous algorithm on an asynchronous network. More precisely, the goal is to run any synchronous $R(n)$ -round $M(n)$ -message complexity CONGEST model algorithm on an asynchronous network with minimal time and message overhead. The first results on synchronizers are by Awerbuch [4], called synchronizers α , β , and γ . Subsequently, these results were improved by Awerbuch and Peleg [8], and Awerbuch et al. [7], both having $O(R(n) \log^3 n)$ time and $O(M(n) \log^3 n)$ message complexity.

The synchronizer γ from [4] essentially consists of running a combination of the simple synchronizers α and β on a sparsified (ζ, δ) -low diameter decompositions. In that case, the bound ζ on the sparsified inter-cluster edges goes into the bound for the communication overhead of the synchronizer and the strong diameter bound δ goes into the bound for the time overhead of the synchronizer. Applying synchronizer γ on our clustering, we obtain the following result.

► **Theorem 5.** *There exists an algorithm that, given $k \geq 2$, can run any synchronous $R(n)$ -round $M(n)$ -message complexity CONGEST model algorithm on an asynchronous CONGEST network. In expectation, the algorithm uses a total of $O(M(n) + R(n)n^{1+1/k})$ messages. Provided that each message incurs a delay of at most one time unit, it takes $O(R(n)k)$ rounds. The initialization phase takes $O(k)$ time, using $O(km)$ messages.*

The running time claimed in this theorem is independent of the random choices in our algorithm, which is a direct result of Theorem 3. The previous sparsified low diameter decompositions (implicit in [17]) would provide similar bounds on the running time, but only with constant probability.

Low Diameter Decompositions. Perhaps unsurprisingly, we show that, with the right choice of parameters, our clustering algorithm can also compute *unsparsified* low diameter decompositions.

► **Theorem 6.** *There exists an algorithm, such that for each graph $G = (V, E)$ with integer weights $w: E \rightarrow \{1, \dots, W\}$ and parameter $\beta \in (0, 1]$ it outputs a low diameter decomposition, whose components are clusters of strong diameter of at most $O\left(\frac{\log n}{\beta}\right)$. Moreover, each edge is an inter-cluster edge with probability at most $\beta \cdot w(u, v)$. The algorithm runs in $O\left(\frac{\log n}{\beta}\right)$ rounds in the CONGEST model, and in $O\left(\frac{\log n \log^* n}{\beta}\right)$ depth and $O(m)$ work in the PRAM model.*

Similar to our spanner algorithm, the bounds on the running time and strong diameter hold independent of the random choices within the algorithm, as opposed to the previous state of the art [31], where they only hold with high probability. In the low diameter decomposition as discussed above, the trade-off between β and diameter bound $O\left(\frac{\log n}{\beta}\right)$ is essentially optimal [9].

Technical Overview

Our clustering algorithm follows an approach known as ball growing, related to the probabilistic partitions of [9, 10]. In a sequential setting, this consists of picking a vertex, and repeatedly adding the neighbors of the current vertices to the ball. This stops when a certain bound is reached, such as a bound on the diameter of the ball or on the number of edges between the current ball and the remainder of the graph. The algorithm repeats this procedure with the remainder of the graph until this is empty. Miller, Peng, and Xu [31] showed that this can be parallelized by letting each vertex create its own ball, but after a certain start time delay. In [31], this has been done by sampling the delays from the exponential distribution, which leads to the aforementioned probabilistic diameter guarantee, as the exponential distribution can take infinitely high values – albeit with small probability. Furthermore, multiple authors (see e.g. [19, 31]) argue that one can round the sampled values from the exponential distribution for most of the algorithm and solely use that the fractional values of the sampled value induce a random permutation of the nodes. In this paper, we show that even fewer random bits are needed: we do not require a random permutation of the nodes. We demonstrate that a tie-breaking rule based on the IDs is enough.

We sample with a capped-off geometric distribution, also used in [28, 7]. As opposed to the standard geometric distribution, the capped-off version can only take a finite number of values. We believe this leads to a more direct proof of the spanner algorithm of [17] and of the decomposition algorithm of [31]. Moreover, by making the sparsifier low diameter

decomposition explicit, the application to synchronizers is almost immediate. In the remainder of this paper, we will not think of the sampled values as start time delays, but as the distance to some conceptual source s , similar to the view in [31]. The rest of the clustering algorithm then consists of computing a shortest path tree rooted at s , which is easily calculated, both in the CONGEST and PRAM model. The clusters consists of the trees that remain when we disconnect the shortest path tree by removing the root s .

As an anonymous reviewer pointed out, in the case of low diameter decompositions, the algorithm of Miller et al. [31] admits an alternative approach. We can exploit the fact that the exponential delays are bounded with high probability. In case the delays exceed the bound, we could return a sub-optimal clustering, without any central communication. As this only happens with low probability, it does not impact the expected number of inter-cluster edges. Note however, that the spanner construction of Elkin and Neiman [17] is not in this high-probability regime, therefore this straightforward approach would not work. We additionally believe that, beyond the result itself, our algorithm provides a more streamlined view.

2 The Clustering Algorithm

Let $G = (V, E)$ be a graph with integer weights $w: E \rightarrow \{1, \dots, W\}$. Let $p \in (0, 1)$ and $r \in \mathbb{N}$ be parameters, to be chosen according to the application of our algorithm. In the following, we provide an algorithm for computing a clustering, where the strong diameter of these clusters will be $2r$. In particular, we will show that each cluster is tree-supported by a tree of height r . The number of inter-cluster edges depends on both p and r , and can be bounded in two ways. The first approach, detailed in Section 3, shows we have a sparsified low diameter decomposition. Here, for each vertex we compute the expected number of edges in the sparsified set of inter-cluster edges, which gives a bound that does not depend on m , but only on n, p and r . The second approach, detailed in Section 4, shows we have a probabilistic low diameter decomposition, by computing the probability that any edge is an inter-cluster edge.

2.1 Construction

First we conceptually add a node s to the graph G to form the graph G' . The node s will function as an artificial source for a shortest path tree. Each vertex will have a distance to s in G' depending on some random offset. Hereto, each vertex samples a value δ_v from the *capped exponential distribution* $\text{GeomCap}(p, r)$, defined by

$$\mathbb{P}[\text{GeomCap}(p, r) = i] = \begin{cases} p(1-p)^i & \text{if } 0 \leq i \leq r-1; \\ (1-p)^r & \text{if } i = r; \\ 0 & \text{else.} \end{cases}$$

This distribution corresponds to the model where we repeat at most r Bernoulli trials, and measure how many trials occur (strictly) before the first success, or whether there is no success in the first r trials. We check that GeomCap is indeed a probability distribution on $\{0, 1, \dots, r\}$:

$$\sum_{i=0}^r \mathbb{P}[\text{GeomCap}(p, r) = i] = \sum_{i=0}^{r-1} p(1-p)^i + (1-p)^r = p \frac{1 - (1-p)^r}{1 - (1-p)} + (1-p)^r = 1.$$

As the intuition suggests, GeomCap has a memoryless property as long as the cap is not reached, i.e., $\mathbb{P}[\text{GeomCap}(p, r) = i \mid \text{GeomCap}(p, r) \geq i] = p$ for $i \leq r - 1$. The proof is completely analogous to the proof of the memoryless property of the geometric distribution.

For each vertex $v \in V$, we conceptually add an edge (s, v) to G' , with weight $w(s, v) := r - \delta_v$. We define $d^{(u)}(s, x) := w(s, u) + d_G(u, x)$, which is the minimal path length, for a path from s to x over u . Now we have that the distance from s to x equals $d_{G'}(s, x) = \min_{u \in V} \{d^{(u)}(s, x)\}$. We call this the *level* of x , ranging from 0 (closest to s) to r (furthest from s). Moreover, we define $p_u(x)$ to be the predecessor of x on an arbitrary but fixed shortest path from u to x . Next, we construct a shortest path tree T rooted at s . When necessary, we do tie-breaking according to IDs: let v be such that $d_{G'}(s, x) = d^{(v)}(s, x)$ and $\text{ID}(v) < \text{ID}(u)$ for all $u \in V$ satisfying $d_{G'}(s, x) = d^{(u)}(s, x)$. Then we connect x to the shortest path tree using the edge $(p_v(x), x)$. Moreover, we add x to the cluster of v and write $c_x := v$ for the corresponding cluster center. Intuitively, the clusters correspond to the connected components that remain when we remove the source s from the created shortest path tree. The formal argument for this can be found in the proof of Lemma 7. The computation of this shortest path tree is model-specific, we provide details in Section 2.3.

The algorithm outputs the shortest path tree T , and for each $x \in V$, the center of its cluster center c_x and its level. The knowledge of cluster centers immediately gives a clustering, where – by the remark above – each cluster has radius at most r . In Section 3, we show how to construct a set of edges $F \subseteq E$ from the cluster centers and levels, such that $H := T \cup F$ is a spanner.

In the above, we only need an arbitrary ordering of the vertices. If we assume that each vertex $v \in V$ has a unique identifier, $\text{ID}(v) \in \{1, \dots, N\}$, we can provide an alternative way of constructing the same shortest path tree. We construct a graph where $w(s, v) = r - \delta_v + \text{ID}(v)/(N + 1)$, and compute a shortest path tree rooted at s . This embeds the tie-breaking rule in the weight of the added edges, and thus in the distances. For generality – and suitable implementation in distributed models with limited bandwidth – the remainder of this paper relies on the former characterization using the tie-breaking rule.

2.2 Tree-Support

Next, we will show that the created clusters are tree-supported by a tree of height r . We have already chosen cluster centers, and we will show that we can identify trees rooted at these centers that satisfy the tree-support condition.

► **Lemma 7.** *Each cluster is tree-supported by a tree of height at most r .*

Proof. Let $v \in V$ be a vertex, which is part of the cluster centered at c_v . We show that there is a path from c_v to v contained in this cluster, which has length at most $d_G(c_v, v) \leq r$. We proceed to show by induction on $d_G(c_v, v)$ that there is a path from c_v to v contained in their cluster, which has length at most $d_G(c_v, v)$. The base case, $v = c_v$, is trivial. Let u be the predecessor of v on some path from c_v to v of length $d_G(c_v, v)$. It suffices to show that u is in the same cluster, then the result follows from the induction hypothesis. By definition of c_v , we have that

$$d_{G'}(s, v) = d_{G'}(s, c_v) + d_G(c_v, v) = d_{G'}(s, c_v) + d_G(c_v, u) + w(u, v) = d^{(c_v)}(s, u) + w(u, v).$$

By the triangle inequality we have $d_{G'}(s, v) \leq d_{G'}(s, u) + w(u, v)$. Combining this, we see $d^{(c_v)}(s, u) \leq d_{G'}(s, u)$. As the distance $d_{G'}(s, u)$ is minimal, by definition we have $d^{(c_v)}(s, u) = d_{G'}(s, u)$. Now suppose that u is part of some cluster c_u . Then we have $d^{(c_u)}(s, u) = d_{G'}(s, u)$

and $\text{ID}(c_u) \leq \text{ID}(c_v)$. However, this implies that $d^{(c_u)}(s, v) \leq d^{(c_u)}(s, u) + w(u, v) = d^{(c_v)}(s, u) + w(u, v) = d^{(c_v)}(s, v)$. Hence by the tie-breaking for v we have $\text{ID}(c_v) \leq \text{ID}(c_u)$ and thus $c_u = c_v$. \blacktriangleleft

As an immediate corollary, we obtain a bound on the strong diameter of the clusters.

► **Corollary 8.** *Each cluster has a strong diameter of $2r$.*

2.3 Implementation and Running Time

For the RAM model, the implementation is straightforward and can be done in linear time [37]. The implementation in distributed and parallel models requires a little more attention. For both models, the computational aspect is very similar to prior work [31, 17].

2.3.1 Distributed Model

The algorithm as presented, can be implemented efficiently both in the LOCAL and in the CONGEST model. It runs in $r + 1$ rounds as follows. In the initialization phase, each vertex v samples its value δ_v and sets its initial distance to the conceptual vertex s as $r - \delta_v$. In the first round of communication, v sends the tuple $(r - \delta_v, \text{ID}(v))$ to its neighbors. In each round, v updates its distance to s according to received messages. It then broadcasts the tuple of its updated distance and the ID corresponding to the first vertex on the path from s to v . Note that at the end of the algorithm, each node knows its own level and cluster center, and the level and cluster center of each of its neighbors.

When the algorithm is applied with $r = O(n)$ (if $r \geq n$, we can simply return the connected components of the graph as clusters), we maintain a bound on the message size of $O(\log n)$, so there are no digit precision consideration for the CONGEST model. Moreover, each vertex v has distance at most $r - \delta_v \leq r$ to s , the algorithm terminates within $r + 1$ rounds.³

2.3.2 PRAM Model

The implementation in the PRAM model is slightly different to the CONGEST model. Instead of broadcasts by each vertex in each round, a vertex v updates its distance only once: either after one of its neighbors updated its distance, or after time $r + 1 - \delta_v$ it sets its distance to $r - \delta_v$. The total required depth differs on the exact model of parallelism, it is $O(r \log^* n)$ in the CRCW model of parallel computation. To show this, we follow the general lines of [25], but we have to be careful: during the shortest path computation, we might need to apply our tie-breaking rule, i.e., finding the minimum ID among all options. Note that in the PRAM model, we can assume without loss of generality that the IDs are labeled 1 to n in the adjacency list representation. Finding the minimum can be done with high probability in $O(\log^* n)$ depth and $O(n)$ work, as we can sort a list of integers between 1 and n in $O(\log^* n)$ depth and $O(n)$ work [21]. If we exceed the $O(\log^* n)$ depth bound, we stop and output the trivial clustering consisting of singletons. This clustering clearly satisfies the diameter bound, and as we only output it with low probability, it has no effect on the expected number of inter-cluster edges. So we can conclude that the additional sorting overhead for the tie-breaking is a factor $O(\log^* n)$. The algorithm has total work $O\left(m + \frac{n}{1-p}\right)$, where the contribution of $O\left(\frac{n}{1-p}\right)$ comes from sampling from the geometric distribution. In this paper, this factor vanishes as we always have p such that $O\left(\frac{n}{1-p}\right) = O(m)$.

³ The “+1” appears, as nodes in the lowest level have distance 0 to the source s .

3 Constructing a Sparsified Low Diameter Decomposition

In this section, we show how the clustering algorithm leads to a sparsified low diameter decomposition. The procedure is as follows: given $k \geq 2$, we set $r = k - 1$, $p = 1 - n^{-1/k}$, and compute a clustering using the algorithm of Section 2. We denote $F \subseteq E$ for the sparsified set of inter-cluster edges. Intuitively, for each vertex $v \in V$, we add an edge to F for each cluster in which we have a neighbor on one level below, or a neighbor on the same level as v with the ID of the cluster center smaller than the ID of center of the cluster of v .

► **Lemma 9.** *There exists a set of edges $F \subseteq E$ of expected size $O(n^{1+1/k})$, such that for every edge, one of its endpoints has an edge from F into the cluster of the other endpoint.*

Proof. We define, $F := \bigcup_{x \in V} C(x)$, where $C(x)$ consists of the following edges

$$C(x) := \{(x, p_u(x)) : d^{(u)}(s, x) = d_{G'}(s, x)\} \\ \cup \{(x, p_u(x)) : d^{(u)}(s, x) = d_{G'}(s, x) + 1 \text{ and } \text{ID}(u) < \text{ID}(c_x)\}.$$

First, we show that F satisfies the property stated in the lemma, then we consider its size. Let $(x, y) \in E$, without loss of generality, we assume $d_{G'}(s, x) \geq d_{G'}(s, y)$, and in case of equality we assume $\text{ID}(c_x) > \text{ID}(c_y)$. We will show that there is an edge $(x, p_{c_y}(x)) \in C(x)$ to the cluster of y . First of all, notice that $d_{G'}(s, y) \geq d_{G'}(s, x) - 1$ by the triangle inequality. If $d_{G'}(s, y) = d_{G'}(s, x) - 1$, then $d^{(c_y)}(s, x) \leq d_{G'}(s, x)$. Because of minimality of $d_{G'}(s, x)$, we have $d^{(c_y)}(s, x) = d_{G'}(s, x)$, and thus $(x, p_{c_y}(x)) \in C(x)$ by definition of $C(x)$. If $d_{G'}(s, y) = d_{G'}(s, x)$, we have $\text{ID}(c_x) > \text{ID}(c_y)$. Moreover, we have $d^{(c_y)}(s, x) \leq d_{G'}(s, x) + 1$. So again it follows that $(x, p_{c_y}(x)) \in C(x)$ by definition of $C(x)$.

Now, we turn to the expected size of F . By linearity of expectation, we have $\mathbb{E}[F] = \sum_{x \in V} \mathbb{E}[C(x)]$. We will show that for each $x \in V$ the expected size of $C(x)$ is at most $2n^{1/k}$. For each $u \in V$, we potentially add an edge $(x, p_u(x))$ to $C(x)$. First, we calculate the probability that at least t such vertices u contribute an edge. Hereto, we look at the random variables $X_u = d^{(u)}(s, x) = k - \delta_u + d_G(u, x)$. According to these random variables, we order all vertices: $V = \{u_1, u_2, \dots, u_n\}$, such that for $i < j$ we satisfy one of the following properties

- $X_{u_i} < X_{u_j}$;
- $X_{u_i} = X_{u_j}$ and $\text{ID}(u_i) < \text{ID}(u_j)$.

We calculate $\mathbb{P}[|C(X)| \geq t]$, i.e., the probability that $\{(x, p_{u_1}(x)), \dots, (x, p_{u_t}(x))\} \subseteq C(x)$. We do this by conditioning on $u_t = v$. We observe

$$\mathbb{P}[\{(x, p_{u_1}(x)), \dots, (x, p_{u_t}(x))\} \subseteq C(x) \mid u_t = v] = \\ \prod_{i=1}^{t-1} \mathbb{P}[(x, p_{u_i}(x)) \in C(x) \text{ and } (x, p_v(x)) \in C(x) \mid u_t = v].$$

So we calculate $\mathbb{P}[(x, p_{u_i}(x)) \in C(x) \text{ and } (x, p_v(x)) \in C(x) \mid u_t = v]$ for $i < t$. By definition of $C(x)$, this can only hold if either x 's closest neighbors in the clusters centered at u_i and v are on the same level (in which case we have $\text{ID}(u_i) < \text{ID}(v)$, as $v = u_t$ and $i < t$) or the neighbor from the cluster centered at u_i is at a level lower and $\text{ID}(v) < \text{ID}(u_i)$. Note that the level of the closest neighbor in the cluster of u_i or v corresponds to the distance $d^{(u_i)}(s, x) = X_{u_i}$ or $d^{(v)}(s, x)$ respectively. As the allowed distances depend on the ID of u_i , we split the vertices according to ID:

$$V_{<} := \{u \in V : \text{ID}(u) < \text{ID}(v)\}; \\ V_{>} := \{u \in V : \text{ID}(u) > \text{ID}(v)\}.$$

If $u_i \in V_{<}$, then we know $X_{u_i} \leq d^{(v)}(s, x)$. If both $(x, p_{u_i}(x))$ and $(x, p_v(x))$ are in $C(x)$, we must have $X_{u_i} = d^{(v)}(s, x)$. So for every $i < t$, we are looking at

$$\begin{aligned} & \mathbb{P}[(x, p_{u_i}(x)) \in C(x) \text{ and } (x, p_v(x)) \in C(x) \mid u_i \in V_{<} \text{ and } u_t = v] \\ & \leq \mathbb{P}\left[X_{u_i} = d^{(v)}(s, x) \mid u_i \in V_{<} \text{ and } u_t = v\right], \\ & = \mathbb{P}\left[X_{u_i} = d^{(v)}(s, x) \mid u_i \in V_{<} \text{ and } X_{u_i} \leq d^{(v)}(s, x)\right], \end{aligned}$$

where the last equality holds as we only rewrote the condition using the order of the u_j 's. We fill in the definition of X_{u_i} and use that the probability of the event we are looking at is independent of $\text{ID}(u_i)$:

$$\begin{aligned} & \mathbb{P}\left[X_{u_i} = d^{(v)}(s, x) \mid u_i \in V_{<} \text{ and } X_{u_i} \leq d^{(v)}(s, x)\right] \\ & = \mathbb{P}\left[k - \delta_{u_i} + d_G(u_i, x) = d^{(v)}(s, x) \mid u_i \in V_{<} \text{ and } k - \delta_{u_i} + d_G(u_i, x) \leq d^{(v)}(s, x)\right] \\ & = \mathbb{P}\left[k - \delta_{u_i} + d_G(u_i, x) = d^{(v)}(s, x) \mid k - \delta_{u_i} + d_G(u_i, x) \leq d^{(v)}(s, x)\right] \\ & = \mathbb{P}\left[\delta_{u_i} = k + d_G(u_i, x) - d^{(v)}(s, x) \mid \delta_{u_i} \geq k + d_G(u_i, x) - d^{(v)}(s, x)\right]. \end{aligned}$$

When $k + d_G(u_i, x) - d^{(v)}(s, x) \leq k - 2$, this equals p , by the memoryless property of the geometric distribution. To distinguish this, we partition the vertices u with $\text{ID}(u) < \text{ID}(v)$ into two set

$$\begin{aligned} V_{<,1} & := \{u \in V : \text{ID}(u) < \text{ID}(v) \text{ and } k + d_G(u_i, x) - d^{(v)}(s, x) \leq k - 2\}; \\ V_{<,2} & := \{u \in V : \text{ID}(u) < \text{ID}(v) \text{ and } k + d_G(u_i, x) - d^{(v)}(s, x) = k - 1\}. \end{aligned}$$

Now for $u_i \in V_{>}$, we obtain by the same reasoning

$$\begin{aligned} & \mathbb{P}[(x, p_{u_i}(x)) \in C(x) \text{ and } (x, p_v(x)) \in C(x) \mid u_i \in V_{>} \text{ and } u_t = v] \\ & = \mathbb{P}\left[\delta_{u_i} = k + d_G(u_i, x) - d^{(v)}(s, x) + 1 \mid \delta_{u_i} \geq k + d_G(u_i, x) - d^{(v)}(s, x) + 1\right]. \end{aligned}$$

As before, when $k + d_G(u_i, x) - d^{(v)}(s, x) + 1 \leq k - 2$, this equals p , by the memoryless property of the geometric distribution. And again, we partition $V_{>}$ into two sets

$$\begin{aligned} V_{>,1} & := \{u \in V : \text{ID}(u) > \text{ID}(v) \text{ and } k + d_G(u_i, x) - d^{(v)}(s, x) + 1 \leq k - 2\}; \\ V_{>,2} & := \{u \in V : \text{ID}(u) > \text{ID}(v) \text{ and } k + d_G(u_i, x) - d^{(v)}(s, x) + 1 = k - 1\}. \end{aligned}$$

If we define $V_1 = V_{<,1} \cup V_{>,1}$ and $V_2 = V_{<,2} \cup V_{>,2}$, we can summarize our results as

$$\mathbb{P}[(x, p_{u_i}(x)) \in C(x) \text{ and } (x, p_v(x)) \in C(x) \mid u_i \in V_1 \text{ and } u_t = v] \leq p.$$

Next, we split the expected value of $C(x)$ depending on V_1 and V_2 :

$$\mathbb{E}[|C(x)| \mid u_t = v] = \mathbb{E}[|C(x) \cap V_1| \mid u_t = v] + \mathbb{E}[|C(x) \cap V_2| \mid u_t = v].$$

We bound the first summand with $n^{1/k}$, independent of v . Hereto, we observe that for any non-negative discrete random variable X we have

$$\mathbb{E}[X] = \sum_{s=1}^{\infty} s \mathbb{P}[X = s] = \sum_{s=1}^{\infty} \sum_{t=1}^s \mathbb{P}[X = s] = \sum_{t=1}^{\infty} \sum_{s=t}^{\infty} \mathbb{P}[X = s] = \sum_{t=1}^{\infty} \mathbb{P}[X \geq t].$$

16:10 Spanners and Low Diameter Decompositions

Using this, we obtain

$$\begin{aligned}
& \mathbb{E}[|C(x) \cap V_1| \mid u_t = v] \\
&= \sum_{t=1}^n \mathbb{P}[|C(x) \cap V_1| \geq t \mid u_t = v] \\
&\leq \sum_{t=1}^n \prod_{i=1}^{t-1} \mathbb{P}[(x, p_{u_i}(x)) \in C(x) \text{ and } (x, p_v(x)) \in C(x) \mid u_i \in V_1 \text{ and } u_t = v] \\
&\leq \sum_{t=1}^n p^{t-1} \\
&\leq \sum_{t=0}^{\infty} p^t \\
&= \frac{1}{1-p} \\
&= n^{1/k},
\end{aligned}$$

where the last equality holds by definition of p . For the second summand, we look at all v simultaneously.

$$\begin{aligned}
\sum_{v \in V} \mathbb{E}[|C(x) \cap V_2| \mid u_t = v] \mathbb{P}[u_t = v] &\leq \sum_{v \in V} \mathbb{E}[|V_2| \mid u_t = v] \mathbb{P}[u_t = v] \\
&\leq \sum_{v \in V} \mathbb{E}[|\{u \in V : \delta_u = k-1\}| \mid u_t = v] \mathbb{P}[u_t = v] \\
&= \mathbb{E}[|\{u \in V : \delta_u = k-1\}|],
\end{aligned}$$

where the last equality holds by the law of total probability. We bound this as follows

$$\mathbb{E}[|\{u \in V : \delta_u = k-1\}|] \leq n \mathbb{P}[\delta_u = k-1] = n(1-p)^{k-1} = n^{1/k},$$

where the last equality holds by definition of p . In total, this gives us $\mathbb{E}[|C(x)|] \leq 2n^{1/k}$. ◀

Together Lemma 7 and Lemma 9 imply the following theorem.

► **Theorem 3 (Restated).** *There exists an algorithm, such that for each unweighted graph $G = (V, E)$ and parameter $k \geq 2$ it outputs a tree-supported sparsified $(\zeta, 2k-2)$ -low diameter decomposition, with $\zeta = O(n^{1+1/k})$ in expectation. The algorithm runs in k rounds in the CONGEST model, and in $O(k \log^* n)$ depth and $O(m)$ work in the PRAM model.*

3.1 Constructing a Spanner

Now, we can construct a spanner from the tree supported low diameter decomposition in the following manner. Let T denote the support forest, and let F denote the set as given in Lemma 9. We define the spanner $H := (V, T \cup F)$. As any forest has at most $n-1$ edges, the expected size of H is at most $O(n^{1+1/k})$. Actually, one could also show that $T \subseteq F$ in our construction of F , but this would not impact the asymptotic size bound. To show that H is a spanner, we show its of limited stretch.

► **Lemma 10.** *H is a spanner of stretch $2k-1$.*

Proof. We will show that for every edge $(x, y) \in E$, there exists a path from x to y in H of length at most $2k - 1$. Consequently we have that $d_H(x, y) \leq (2k - 1)d_G(x, y)$ for every $x, y \in V$, hence H is a spanner of stretch $2k - 1$.

Let $(x, y) \in E$. By definition of F , one of the endpoints has an edge in F into the cluster of the other endpoint. Without loss of generality, let there be an edge $(x, z) \in F$ with z in the cluster of y . By Corollary 8, there is path of length at most $2(k - 1)$ from z to y , so in total we have a path of length at most $2(k - 1) + 1 = 2k - 1$ from x to z to y . ◀

Now, the following corollary follows from Theorem 3 and Lemma 10.

► **Corollary 11 (Restated).** *There exists an algorithm, such that for each unweighted graph $G = (V, E)$ and parameter $k \geq 2$ it outputs a spanner H of stretch $2k - 1$. The expected size of H is at most $O(n^{1+1/k})$. The algorithm runs in k rounds in the CONGEST model, and in $O(k \log^* n)$ depth and $O(m)$ work in the PRAM model.*

3.2 Constructing a Synchronizer

Suppose we are given a synchronous CONGEST model algorithm, but we want to run it on an asynchronous CONGEST network. That is, the messages sent in the network can now have arbitrary delays and, in an event-driven manner, nodes become active each time they receive a message. For the purpose of analyzing the time complexity of the algorithm, it is often assumed that the delay is at most one unit of time, however, the algorithm should behave correctly under *any* finite delays. In this situation, a node should start simulating its next (synchronous) round when it has received all the messages from the previous round from its neighbors. The problem is that it cannot tell the difference between the situation if a message from a particular neighbor has not arrived yet or if this same neighbor is not sending any message in that round at all. We say that a node is *safe* if all the messages it has sent have arrived at their destination. In order to determine whether all neighboring nodes are safe, additional messages are sent. The procedure governing these additional messages is called the *synchronizer*. There are two things to take into account when analyzing synchronizers. First, the time overhead: how much time is needed to send the additional messages for each synchronous round. Second, the message-complexity (or communication) overhead: how many additional messages are sent. For more details on synchronizers see e.g. [29, 26].

Let us first consider two simple synchronizers: synchronizer α and synchronizer β , see [4]. In synchronizer α , when a node receives a message from a neighbor, it sends back an “acknowledge” message. When a node has received acknowledge messages for all its sent messages, it marks itself safe and reports this to all its neighbors. The synchronizer α uses, for each simulated synchronous round, additional $O(1)$ time, and $O(m)$ messages.

Synchronizer β will produce a different trade off between time and message overhead. It uses an initialization phase in which it creates a rooted spanning tree, where the root is declared the leader. Now after sending messages of a certain synchronous round, again nodes that receive messages reply with an acknowledge message to each. Nodes that are safe, and whose children in the constructed tree are also safe communicate this to their parent in the tree. Eventually the leader will learn that the whole graph is safe, and will broadcast this along the spanning tree. Synchronizer β uses $O(D)$ time and $O(n)$ messages per synchronous round.

Now we are ready to consider a little more involved example, called synchronizer γ , see [4]. This synchronizer makes use of clustering, where within each cluster synchronizer β is used and between clusters synchronizer α is used. In the LOCAL model, the cluster centers can simply select a communication link for each neighboring cluster to communicate

16:12 Spanners and Low Diameter Decompositions

individually with the neighboring cluster centers [4]. However, in the CONGEST model, communicating information about neighboring clusters to the cluster center might lead to congestion problems. Using a slightly more careful analysis, the procedure can be adapted to the CONGEST model.

► **Lemma 11** (Implicit in [4]). *Given a $T(n)$ -round synchronous CONGEST model algorithm for constructing a sparsified (ζ, δ) -low diameter decomposition, any synchronous $R(n)$ -round $M(n)$ -message complexity CONGEST model algorithm can be run on an asynchronous CONGEST network with a total of $O(M(n) + R(n)(\zeta + n))$ messages, and, provided that each message incurs a delay of at most one time unit, in time $O(R(n)\delta)$. The initialization phase takes $O(T(n))$ time, using $O((T(n) + \delta)m)$ messages.*

For a sketch of the algorithm, we refer to Appendix A. If we plug in our clustering, we obtain the following theorem.

► **Theorem 5** (Restated). *There exists an algorithm that, given $k \geq 2$, can run any synchronous $R(n)$ -round $M(n)$ -message complexity CONGEST model algorithm on an asynchronous CONGEST network. In expectation, the algorithm uses a total of $O(M(n) + R(n)n^{1+1/k})$ messages. Provided that each message incurs a delay of at most one time unit, it takes $O(R(n)k)$ rounds. The initialization phase takes $O(k)$ time, using $O(km)$ messages.*

4 Constructing a Low Diameter Decomposition

In this section, we show that for an integer weighted graph the computed clustering is a probabilistic low diameter decomposition. To be precise, if we set $p = \frac{\beta}{4}$, and $r = \left\lceil \frac{1}{p} \ln \left(\frac{n^2}{p} \right) + \frac{1}{4p} \right\rceil$ we obtain a $(\beta, O(\frac{\log n}{\beta}))$ -low diameter decomposition. By Corollary 8, we know that each of the clusters has a strong diameter of at most $2r = 2 \left\lceil \frac{1}{p} \ln \left(\frac{n^2}{p} \right) + \frac{1}{4p} \right\rceil = O\left(\frac{\log n}{\beta}\right)$. Now, we show that the probability that an edge $e \in E$ is an inter-cluster edge is at most $4p \cdot w(e) = \beta \cdot w(e)$. We use a general proof structure from [31], but make it more streamlined; we avoid an artificially constructed “midpoint” on the edge (u, v) . Further, our proof borrows the idea of conditioning on the event $E_{u', v', \alpha}$ from Xu [39], which we adapt to our situation.

► **Lemma 12.** *For $(u, v) \in E$, the probability that u and v belong to different clusters is at most $4p \cdot w(u, v)$.*

Proof. Suppose (u, v) is an inter-cluster edge. Without loss of generality, we assume $d_{G'}(s, v) \leq d_{G'}(s, u)$. By the triangle inequality, we have $d_{G'}(s, u) \leq d_{G'}(s, v) + w(u, v)$, hence we have $d_{G'}(s, u) - d_{G'}(s, v) \leq w(u, v)$. Using this, we can upper bound the probability that an edge (u, v) is an inter-cluster edge by the probability that this inequality holds. Note that we can assume $4p \cdot w(u, v) < 1$, otherwise the statement is trivially true.

We want to condition on the cluster center v' that satisfied $d^{(v')}(s, u) = d_{G'}(s, v)$, and on the cluster center $u' \neq v'$ that minimizes $d^{(u')}(s, u)$. Moreover, we ask that these cluster respect the tie-breaking rule, i.e., both have minimal ID among all centers with equal distance. Finally, we condition on the value of $d_{G'}(s, u)$, which we set to α . We call this event $E_{u', v', \alpha}$, which we formally define to hold when the following four conditions are satisfied

1. $d^{(v')}(s, u) \leq \alpha$;
2. for $w' \in V \setminus \{v'\}$ we either have $d^{(w')}(s, v) > d^{(v')}(s, v)$, or we have $d^{(w')}(s, v) = d^{(v')}(s, v)$ and $\text{ID}(v') < \text{ID}(w')$;

3. $d^{(u')}(s, v) = \alpha$;
4. for $w' \in V \setminus \{u', v'\}$ we either have $d^{(w')}(s, u) > d^{(u')}(s, u)$, or we have $d^{(w')}(s, u) = d^{(u')}(s, u)$ and $\text{ID}(u') < \text{ID}(w')$.

Now, we condition on $E_{u', v', \alpha}$ and use the law of total probability:

$$\begin{aligned}
& \mathbb{P}[(u, v) \text{ is an inter-cluster edge}] \\
&= \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha} \mathbb{P}[(u, v) \text{ is an inter-cluster edge} \mid E_{u', v', \alpha}] \mathbb{P}[E_{u', v', \alpha}] \\
&\leq \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha} 2 \mathbb{P} \left[d^{(u')}(s, u) - d^{(v')}(s, v) \leq w(u, v) \mid E_{u', v', \alpha} \right] \mathbb{P}[E_{u', v', \alpha}].
\end{aligned}$$

For simplicity, we omit the bounds for the sum over α , which is a finite sum as we always have $\alpha \leq r + m$. The factor two appears because the event assumes $d_{G'}(s, v) \leq d_{G'}(s, u)$, hence we gain a factor two by symmetry of u and v . We look at the first probability more closely. We can loosen some of the event's restrictions, and just maintain $d^{(v')}(s, v) \leq \alpha$, as the event we examine is independent of conditions 2, 3, and 4 of the event $E_{u', v', \alpha}$. We obtain

$$\begin{aligned}
& \mathbb{P} \left[d^{(u')}(s, u) - d^{(v')}(s, v) \leq w(u, v) \mid E_{u', v', \alpha} \right] \\
&= \mathbb{P} \left[\alpha - d^{(v')}(s, v) \leq w(u, v) \mid d^{(v')}(s, v) \leq \alpha \right] \\
&= \mathbb{P} [\delta_{v'} \leq r + d_G(v', v) - \alpha + w(u, v) \mid \delta_{v'} \geq r + d_G(v', v) - \alpha],
\end{aligned}$$

where the last equality holds by definition of $d^{(v')}(s, v)$. Now if $r + d_G(v', v) - \alpha + w(u, v) \leq r$ (or equivalently, $\alpha \geq d_G(v', v) + w(u, v)$), we stay away from our cap on the geometric distribution, and hence we can apply the memoryless property of the geometric distribution to obtain

$$\begin{aligned}
& \mathbb{P} [\delta_{v'} \leq r + d_G(v', v) - \alpha + w(u, v) \mid \delta_{v'} \geq r + d_G(v', v) - \alpha] \\
&= 1 - (1 - p)^{w(u, v)} \\
&\leq pw(u, v),
\end{aligned}$$

where the last step holds by Bernoulli's inequality. If we have $r + d_G(v', v) - \alpha + w(u, v) > r$ (or equivalently, $\alpha < d_G(v', v) + w(u, v)$), we show that the probability $\mathbb{P}[E_{u', v', \alpha}]$ of the event taking place is already small:

$$\begin{aligned}
\mathbb{P}[E_{u', v', \alpha}] &\leq \mathbb{P} \left[d^{(v')}(s, v) \leq \alpha \text{ and } d^{(u')}(s, u) = \alpha \right] \\
&= \mathbb{P} [\delta_{v'} \geq r + d_G(v', v) - \alpha] \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right] \quad (\text{since the events are independent}) \\
&\leq \mathbb{P} [\delta_{v'} \geq r - w(u, v)] \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right],
\end{aligned}$$

where the last equality holds as $r + d_G(v', v) - \alpha + w(u, v) > r$. We bound this probability as follows:

$$\mathbb{P} [\delta_{v'} \geq r - w(u, v)] \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right] \leq (1 - p)^{r - w(u, v)} \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right].$$

16:14 Spanners and Low Diameter Decompositions

Now we use that $r = \left\lceil \frac{1}{p} \ln \left(\frac{n^2}{p} \right) + \frac{1}{4p} \right\rceil$, to obtain

$$\begin{aligned} (1-p)^{r-w(u,v)} &\leq (1-p)^{\frac{1}{p} \ln \left(\frac{n^2}{p} \right) + \frac{1}{4p} - w(u,v)} \\ &\leq \left((1-p)^{1/p} \right)^{\ln \left(\frac{n^2}{p} \right)} \quad (\text{as } 4p \cdot w(u,v) \leq 1) \\ &\leq \frac{p}{n^2}. \end{aligned}$$

Combining all of this, we obtain

$$\begin{aligned} &\mathbb{P}[(u, v) \text{ is an inter-cluster edge}] \\ &\leq 2 \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha \geq d_G(v', v) + w(u, v)} \mathbb{P} \left[d^{(u')}(s, u) - d^{(v')}(s, v) \leq w(u, v) \mid E_{u', v', \alpha} \right] \mathbb{P} [E_{u', v', \alpha}] \\ &\quad + 2 \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha < d_G(v', v) + w(u, v)} \mathbb{P} \left[d^{(u')}(s, u) - d^{(v')}(s, v) \leq w(u, v) \mid E_{u', v', \alpha} \right] \mathbb{P} [E_{u', v', \alpha}] \\ &\leq 2 \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha \geq d_G(v', v) + w(u, v)} p \cdot w(u, v) \mathbb{P} [E_{u', v', \alpha}] \\ &\quad + 2 \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha < d_G(v', v) + w(u, v)} \frac{p}{n^2} \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right] \\ &\leq 2p \cdot w(u, v) \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha} \mathbb{P} [E_{u', v', \alpha}] + 2 \frac{p}{n^2} \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha} \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right]. \end{aligned}$$

Next, we notice that all events $E_{u', v', \alpha}$ are disjoint by design, so

$$\sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \sum_{\alpha} \mathbb{P} [E_{u', v', \alpha}] = 1.$$

Clearly we have $\sum_{\alpha} \mathbb{P} \left[d^{(u')}(s, u) = \alpha \right] = 1$, as this is just a sum over all possible values of $d^{(u')}(s, u)$. Filling both in, we conclude

$$\mathbb{P}[(u, v) \text{ is an inter-cluster edge}] \leq 2p \cdot w(u, v) + 2 \sum_{u' \in V} \sum_{v' \in V \setminus \{u'\}} \frac{p}{n^2} \leq 4p \cdot w(u, v). \quad \blacktriangleleft$$

Together with Corollary 8, this gives us the following theorem.

► **Theorem 6 (Restated).** *There exists an algorithm, such that for each graph $G = (V, E)$ with integer weights $w: E \rightarrow \{1, \dots, W\}$ and parameter $\beta \in (0, 1]$ it outputs a low diameter decomposition, whose components are clusters of strong diameter of at most $O\left(\frac{\log n}{\beta}\right)$. Moreover, each edge is an inter-cluster edge with probability at most $\beta \cdot w(u, v)$. The algorithm runs in $O\left(\frac{\log n}{\beta}\right)$ rounds in the CONGEST model, and in $O\left(\frac{\log n \log^* n}{\beta}\right)$ depth and $O(m)$ work in the PRAM model.*

5 Conclusion

We have presented an algorithm that computes a clustering, more precisely, a tree-supported sparsified low diameter decomposition. This directly leads to a sparse spanner and can be applied to compute a synchronizer for the CONGEST model. Moreover, we show that we also improve upon the state-of-the-art for low diameter decompositions. By showing

that clustering can be done using a capped geometric distribution, we improve on existing algorithms for spanners and low diameter decompositions in two ways. First, we obtain bounds on the diameter/stretch and running time that are independent of the random choices of the algorithm. Second, the discreteness of the geometric distribution fits the discrete nature of graph theoretical problems better than a continuous distribution. We believe this leads to a more intuitive algorithm.

A natural question that remains is whether it would be possible to give a with-high-probability bound on the total number of inter-cluster edges or the size of the spanner rather than an in-expectation bound. A more ambitious goal is to develop a completely deterministic algorithm with the same bounds, improving on the work of Ghaffari and Kuhn [20].

References

- 1 Vedat Levi Alev, Nima Anari, Lap Chi Lau, and Shayan Oveis Gharan. Graph clustering using effective resistance. In *Proc. of the Innovations in Theoretical Computer Science Conference (ITCS)*, pages 41:1–41:16, 2018.
- 2 Noga Alon, Richard M Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k -server problem. *SIAM Journal on Computing*, 24(1):78–100, 1995.
- 3 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.
- 4 Baruch Awerbuch. Complexity of network synchronization. *Journal of the ACM*, 32(4):804–823, 1985.
- 5 Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. Near-linear cost sequential and distributed constructions of sparse neighborhood covers. In *Proc. of the Conference on Foundations of Computer Science (FOCS)*, pages 638–647, 1993.
- 6 Baruch Awerbuch, Andrew V Goldberg, Michael Luby, and Serge A Plotkin. Network decomposition and locality in distributed computation. In *Proc. of the Conference on Foundations of Computer Science (FOCS)*, volume 30, pages 364–369, 1989.
- 7 Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Michael Saks. Adapting to asynchronous dynamic networks. In *Proc. of the Symposium on Theory of Computing (STOC)*, pages 557–570, 1992.
- 8 Baruch Awerbuch and David Peleg. Network synchronization with polylogarithmic overhead. In *Proc. of the Symposium on Foundations of Computer Science (FOCS)*, pages 514–522, 1990.
- 9 Yair Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Proc. of the Conference on Foundations of Computer Science (FOCS)*, pages 184–193, 1996.
- 10 Yair Bartal. On approximating arbitrary metrics by tree metrics. In *Proc. of the Symposium on Theory of Computing (STOC)*, pages 161–168, 1998.
- 11 Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Structures & Algorithms*, 30(4):532–563, 2007.
- 12 Keren Censor-Hillel, Bernhard Haeupler, Jonathan Kelner, and Petar Maymounkov. Global computation in a poorly connected world: fast rumor spreading with no dependence on conductance. In *Proc. of the Symposium on Theory of Computing (STOC)*, pages 961–970, 2012.
- 13 Yi-Jun Chang, Seth Pettie, and Hengjie Zhang. Distributed triangle detection via expander decomposition. In *Proc. of the Symposium on Discrete Algorithms (SODA)*, pages 821–840, 2019.
- 14 Yi-Jun Chang and Thatchaphol Saranurak. Deterministic distributed expander decomposition and routing with applications in distributed derandomization. In *Proc. of the Symposium on Foundations of Computer Science (FOCS)*, pages 377–388, 2020.
- 15 Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. In *Proc. of the Symposium on Theory of Computing (STOC)*, pages 16–26, 1994.

16:16 Spanners and Low Diameter Decompositions

- 16 Edith Cohen. Fast algorithms for constructing t -spanners and paths with stretch t . *SIAM Journal on Computing*, 28(1):210–236, 1998.
- 17 Michael Elkin and Ofer Neiman. Efficient algorithms for constructing very sparse spanners and emulators. *ACM Transactions on Algorithms (TALG)*, 15, 2019.
- 18 Arnold Filtser and Shay Solomon. The greedy spanner is existentially optimal. In *Proc. of the Symposium on Principles of Distributed Computing (PODC)*, pages 9–17, 2016.
- 19 Sebastian Forster and Gramoz Goranci. Dynamic low-stretch trees via dynamic low-diameter decompositions. In *Proc. of the Symposium on Theory of Computing (STOC)*, pages 377–388, New York, NY, USA, 2019.
- 20 Mohsen Ghaffari and Fabian Kuhn. Derandomizing distributed algorithms with small messages: Spanners and dominating set. In *Proc. of the Symposium on Distributed Computing (DISC)*, pages 29:1–29:17, 2018.
- 21 Joseph Gil, Yossi Matias, and Uzi Vishkin. Towards a theory of nearly constant time parallel algorithms. In *Proc. of the Symposium of Foundations of Computer Science (FOCS)*, pages 698–710, 1991.
- 22 Oded Goldreich and Dana Ron. A sublinear bipartiteness tester for bounded degree graphs. *Combinatorica*, 19(3):335–373, 1999.
- 23 Shay Halperin and Uri Zwick. Personal communication, 1993.
- 24 Ravi Kannan, Santosh Vempala, and Adrian Vetta. On clusterings: Good, bad and spectral. *Journal of the ACM*, 51(3):497–515, 2004.
- 25 Philip N Klein and Sairam Subramanian. A randomized parallel algorithm for single-source shortest paths. *Journal of Algorithms*, 25(2):205–220, 1997.
- 26 Ajay D Kshemkalyani and Mukesh Singhal. *Distributed computing: principles, algorithms, and systems*. Cambridge University Press, 2011.
- 27 Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999.
- 28 Nathan Linial and Michael Saks. Low diameter graph decompositions. *Combinatorica*, 13(4):441–454, 1993.
- 29 Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.
- 30 Gary L Miller, Richard Peng, Adrian Vladu, and Shen Chen Xu. Improved parallel algorithms for spanners and hopsets. In *Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 192–201, 2015.
- 31 Gary L Miller, Richard Peng, and Shen Xu. Parallel graph decompositions using random shifts. *Proc. of the Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 196–203, 2013.
- 32 David Peleg. Proximity-preserving labeling schemes. *Journal of Graph Theory*, 33(3):167–176, 2000.
- 33 David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, 36(3):510–530, 1989.
- 34 Liam Roditty and Uri Zwick. On dynamic shortest paths problems. In *Proc. of the European Symposium on Algorithms (ESA)*, pages 580–591, 2004.
- 35 Thatchaphol Saranurak and Di Wang. Expander decomposition and pruning: Faster, stronger, and simpler. In *Proc. of the Symposium on Discrete Algorithms (SODA)*, pages 2616–2635, 2019.
- 36 Daniel A Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- 37 Mikkel Thorup. Undirected single-source shortest paths with positive integer weights in linear time. *Journal of the ACM*, 46(3):362–394, 1999.
- 38 Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM*, 52(1):1–24, 2005.
- 39 Shen Chen Xu. *Exponential Start Time Clustering and its Applications in Spectral Graph Theory*. PhD thesis, Carnegie Mellon University, Pittsburgh, 2017.

A Using Sparsified Low Diameter Decompositions for Synchronization

In the following, we turn to the algorithm realizing Lemma 11, i.e., we show how we can run a synchronous CONGEST algorithm on an asynchronous CONGEST network, using a sparsified low diameter decomposition. Hereto, we present an implementation of the synchronizer γ in the CONGEST model, using sparsified low diameter decompositions for the communication. We refer to [4] for a proof of correctness of the synchronizer γ .

The initialization phase consists of three steps. First, we compute the sparsified (ζ, δ) -low diameter decomposition. To do this in the asynchronous CONGEST model, we use the synchronizer α (for details see [4], or textbooks, e.g., [29, 26]). Hence this takes $O(T(n))$ time, and $O(T(n)m)$ messages. Second, we pick a cluster center for each cluster and construct a tree rooted at the cluster center spanning the cluster. We can do this in $O(\delta)$ time, using $O(\delta m)$ messages, again using the synchronizer α . Note that if the computed decomposition was tree-supported, these trees are already given. As a third and final step of the initialization phase, each vertex needs to be aware of its incident sparsified inter-cluster edges, as it will use these to communicate to neighboring clusters. This might be already determined during the construction of the clustering. It could also be the case that for each sparsified inter-cluster edge, only one of the two incident vertices knows this. In $O(1)$ time, using $O(m)$ messages, this can be communicated using the synchronizer α . In total, we use $O(T(n))$ time for the initialization phase, and $O((T(n) + \delta)m)$ messages.

Now we are set up for the simulation of the $R(n)$ -round, $M(n)$ -message complexity synchronous CONGEST model algorithm. In each simulation of a synchronous round of this algorithm, vertices respond to each message with an acknowledge message, same as in the synchronizers α and β . When a vertex has received acknowledge messages for each sent message, it declares itself safe. If a vertex and all its children in the cluster tree are safe, it notifies its parent in the cluster tree. Once the cluster center has received confirmation that the whole cluster is safe, it down-casts this information to the whole cluster. Each vertex communicates that its cluster is safe over its sparsified inter-cluster edges. Once a vertex received a message of being safe over each sparsified inter-cluster edge, it declares itself *ready*. When a vertex and all its children in the cluster tree are ready, it sends a ready-message to its parent in the cluster tree. Once a cluster center received ready-messages from the whole cluster, it down-casts a message “cluster ready” to all cluster vertices.

Assuming that each message incurs a delay of at most one time unit, we need at most $O(\delta)$ time for this procedure, as we send information along the trees of height δ for a total of four times. See [4] for the argument explaining why confirmation that neighboring clusters are done suffices. Moreover, the only communication links participating in this procedure, are the edges from the sparsified low diameter decomposition (consisting of at most ζ inter-cluster edges and n tree edges). Each of these edges sends up to four messages in total, giving a total bound on the number of messages of $O(R(n)(\zeta + n))$ for the synchronization. This gives a total time bound of $O(R(n)\delta)$, and message complexity bound of $O(M(n) + R(n)(\zeta + n))$.