# Deterministic Distributed Sparse and Ultra-Sparse Spanners and Connectivity Certificates

Marcel Bezdrighin ETH Zurich mbezdrighin@student.ethz.ch Michael Elkin Ben-Gurion University elkinm@cs.bgu.ac.il

Christoph Grunau ETH Zurich

cgrunau@inf.ethz.ch

Bernhard Haeupler ETH Zurich and CMU

Václav Rozhoň ETH Zurich rozhony@inf.ethz.ch

bernhard.haeupler@inf.ethz.ch

Mohsen Ghaffari ETH Zurich ghaffari@inf.ethz.ch

Saeed Ilchi ETH Zurich saeed.ilchi@inf.ethz.ch

#### Abstract

This paper presents efficient distributed algorithms for a number of fundamental problems in the area of graph sparsification:

- We provide the first deterministic distributed algorithm that computes an ultra-sparse spanner in polylog(n) rounds in weighted graphs. Concretely, our algorithm outputs a spanning subgraph with only n + o(n) edges in which the pairwise distances are stretched by a factor of at most  $O(\log n \cdot 2^{O(\log^* n)})$ .
- We provide a  $\operatorname{polylog}(n)$ -round deterministic distributed algorithm that computes a spanner with stretch (2k-1) and  $O(nk+n^{1+\frac{1}{k}}\log k)$  edges in unweighted graphs and with  $O(n^{1+\frac{1}{k}}k)$  edges in weighted graphs.
- We present the first polylog(n) round randomized distributed algorithm that computes a sparse connectivity certificate. For an *n*-node graph *G*, a certificate for connectivity *k* is a spanning subgraph *H* that is *k*-edge-connected if and only if *G* is *k*-edge-connected, and this subgraph *H* is called sparse if it has O(nk) edges. Our algorithm achieves a sparsity of (1 + o(1))nk edges, which is within a 2(1 + o(1)) factor of the best possible.

# 1 Introduction

This paper studies graph sparsification problems, particularly connectivity-preserving and distancepreserving problems. Concretely, given a network G = (V, E), we are interested in computing a spanning subgraph H = (V, E'), where the subset  $E' \subset E$  has as few as possible edges, while H preserves some particular connectivity or distance properties of G, that we will soon elaborate on. Such sparsifications have various applications. The prototypical usage is that, instead of operating on the entire network G, we now operate on the sparser network H and this would save various costs. For instance, this cost might be the monetary price paid to the network provider, the total amount of communication and thus the energy consumed, or the overall computational complexity. All of these can be proportional to the number of edges in use, and thus operating on a sparser network saves costs.

The two families of properties that we focus on in this paper are distances and connectivity. The former concept is usually called spanner, while the latter is often studied under titles such as sparse k-connectivity certificates and approximation algorithms for k-edge-connected subgraphs. Below, we discuss

two categories separately, in each case first reviewing the definitions and the state of the art and then stating our contributions. Before that, let us briefly recall the standard message-passing model of distributed computation, which we use throughout.

**Distributed Model.** We work with the standard synchronous message-passing model of distributed computation, often called the CONGEST model [Pel00]. Here, the network is abstracted as an *n*-node graph G = (V, E) where each node represents one computer/processor in the network. Each processor/node has a unique  $O(\log n)$ -bit identifier. Communications take place in synchronous rounds, where per round each node can send one  $O(\log n)$ -bit message to each of its neighbors. We comment that if the model is relaxed to allow unbounded message sizes, then it is referred to as the LOCAL model. Throughout the paper, we will work with the CONGEST model but we will discuss in the related work some results in the relaxed LOCAL model. Generally, the input and output are represented in a distributed fashion. At the start of the algorithm, the processors/nodes do not know the topology of the network G; each of them knows just its own neighbors and is able to communicate with each of those neighbors once per round. At the end of the algorithm, each node should know its own part of the output, e.g., when discussing sparsification problems, each node should know which of its edges is the computed subgraph H of G.

#### 1.1 Spanners: Background

Graph spanners were introduced by Peleg and Schäffer [PS89]. Since then, spanners and efficient distributed, parallel, and sequential constructions for them have been studied extensively [ADD+93, ABCP93, Coh93, ACIM99, DHZ00, EP04, Elk05, TZ06, BS07, Pet09, Pet10, BKMP10, Che13, AB17, EN18, GP17, GK18, CHKPY18, PY18a, EM19].

**Definition 1.1** ( $\alpha$ -spanner). For a graph G = (V, E), a subgraph H = (V, E') is a spanner with stretch  $\alpha > 1$  —or simply an  $\alpha$ -spanner— if and only if  $d_H(u, v) \leq \alpha \cdot d_G(u, v)$  for all  $u, v \in V$ . We assume that G is undirected, and we are interested in minimizing |E'| in terms of integers n and  $\alpha$ .

Of particular interest in this paper are sparse spanners, which have O(n) edges, and ultra-sparse spanners, which have n + o(n) edges. More concretely, an ultra-sparse spanner of an *n*-vertex graph has at most n + n/tedges for some t > 1. That is, it is made of n - 1 edges—as necessary for a spanning tree—and only n/t + 1extra edges. As t gets larger, the structure gets closer and closer to a spanning tree, which is the minimal subgraph to keep the graph connected.

Motivation and Applications. Spanners have found numerous applications, including in packet routing, constructing synchronizers, and algorithmics of various graph problems. See e.g., [ABCP93, AP92, Coh93, TZ05, TZ01, GPPR04, LL18] and the recent survey of Ahmed et al. [ABS<sup>+</sup>20].

Again, of particular interest in this paper are sparse and ultra-sparse spanners. A direct application of sparse and ultra-sparse spanners is that we can consider them as a sparse skeleton, i.e., a subgraph that retains connectivity, with an asymptotically optimal number of edges. Furthermore, ultra-sparse spanners have found important applications in the context of solving symmetrically diagonally-dominant (SDD) linear systems, primarily as ultra-sparse spanners can be seen as one spanning tree and only few extra edges. The applications include estimating effective resistances, and computing spectral sparsifiers or ultra-sparse for cuts and flows [SS11, BGK<sup>+</sup>14]. Since cuts, flows, and distances are easy in trees and since ultra-sparse spanners are trees together with a very small number of extra edges, these ultra-sparse structures have emerged as a powerful tool for (recursively) reducing the complexity of numerous fundamental optimization and graph problems. These include maximum flow [Pen16], min-cost flow, lossy flow problems, several variants of min-cut problems [DS08], as well as approximate shortest paths and transshipment problems [Li20].

State of the art, small stretch spanners. We first review the state of the art for spanners in the general case of k and then focus on the particularly interesting regime of sparse and ultra-sparse spanners.

Althöfer et al. [ADD<sup>+</sup>93] provided a simple greedy algorithm for constructing (2k - 1)-spanner with  $n^{1+\frac{1}{k}}$  edges, in unweighted and weighted graphs. This size bound is tight conditioned on the Erdős girth conjecture: There is a family of graphs with girth 2k + 2 and  $\Omega(n^{1+1/k})$  edges [Erd63]. Note that by setting  $k = \omega(\log n)$ , one obtains ultra-sparse spanners with stretch  $\omega(\log n)$ .

In the distributed setting, we discuss prior results in two parts of randomized and deterministic algorithms. Baswana and Sen [BS07, Pet10] presented an O(k) rounds randomized algorithm for (2k-1)-stretch unweighted spanners with expected size  $O(nk+n^{1+\frac{1}{k}}\log k)$ .<sup>1</sup> For weighted graphs, their output has  $O(n^{1+\frac{1}{k}}k)$  edges. Inspired by the work of Miller et al. [MPVX15], Elkin and Neiman [EN18] gave a randomized (2k-1)-spanner for unweighted graphs that runs in O(k) rounds and matches the centralized bound of  $O(n^{1+\frac{1}{k}})$  with a constant probability; concretely, for any  $\varepsilon > 0$ , it has size  $O(n^{1+\frac{1}{k}}/\varepsilon)$  with probability at least  $1 - \varepsilon$ .

For deterministic algorithms, the work of Barenboim et al. [BEG18] gave an unweighted spanner with size  $O(n^{1+\frac{1}{k}})$  but with the weaker stretch  $O(\log^{k-1} n)$  in  $(\log^{k-1} n)$  rounds. Derbel et al. [DMZ10] compute (2k-1)-stretch unweighted spanners with  $O(n^{1+1/k})$  edges but with a rather high round complexity of  $O(n^{1-\frac{1}{k}})$ . Grossman and Parter [GP17] attain the same stretch with size  $O(kn^{1+1/k})$  and improved round complexity of  $O(2^k n^{\frac{1}{2}-\frac{1}{k}})$ . The first polylog(n) rounds algorithm for (2k-1)-spanners in unweighted graphs was devised by Ghaffari and Kuhn [GK18]. Their output has size  $O(nk+kn^{1+\frac{1}{k}}\log n)$ . If we allow unbounded message sizes, there is a deterministic distributed algorithm by Derbel et al. [DGPV08] with  $O(kn^{1+\frac{1}{k}})$  edges and in O(k) rounds in the LOCAL model.

State of the art, sparse and ultra-sparse spanners. The problem of devising efficient distributed and parallel algorithms for computing ultra-sparse spanners in unweighted graphs was extensively studied in [DMZ10, Pet10, RV11, BEG18, EN18]. Pettie [Pet10] presented a distributed randomized algorithm for computing a  $O(2^{\log^* n} \log n)$ -spanner with O(n) edges in  $O(2^{\log^* n} \log n)$  rounds. As discussed before, the randomized algorithm of Elkin and Neimann [EN18] provides a (2k-1)-spanner with  $O(n^{1+\frac{1}{k}})$  edges. With  $k = \Theta(\log n)$ , this automatically gives an O(n)-size sparse spanner. Indeed, their algorithm can compute, with a constant probability, a spanner with  $O(t \log n)$ -stretch and n + n/t edges, in  $O(t \log n)$  rounds. For a discussion on the known centralized and parallel approaches to ultra-sparse spanners, see Appendix A. The state of the art parallel algorithm is that of Li [Li20], and our results improve on it as we mention later.

Unfortunately, the above distributed algorithms do not provide ultra-sparse spanners in weighted graphs. We note that for many of the modern applications of ultra-sparse spanners in algorithmic problems mentioned above we need ultra-sparse spanners for weighted graphs.

We comment that there is a standard and simple reduction from weighted graphs to unweighted graphs, but this reduction does not provide ultra-sparse or even sparse spanners. The reduction works as follows: to compute a spanner for a weighted graph G, we round the weight of each edge to a multiple of  $(1 + \varepsilon)$ , and then compute an unweighted  $\alpha$ -spanner for the edges of each of the  $O(\frac{\log(U+1)}{\varepsilon})$  weight classes separately, where U is the aspect ratio of the edge weights. The union of these spanners forms a spanner with stretch  $(1 + \varepsilon)\alpha$  for G. Besides the  $(1 + \varepsilon)$ -factor loss in stretch, this reduction loses an  $O(\frac{\log(U+1)}{\varepsilon})$ -factor in the number of edges. For ultra-spanners, this reduction completely destroys the ultra-sparsity as the union of even just two ultra-spanners is no longer ultra-sparse. Also, in the standard cases of weighted graphs, we usually assume  $\log U = O(\log n)$ , and thus this reduction does not give even a sparse spanner with O(n)edges.

Lower bounds for ultra-sparse spanners It is known that for every n and t, there exists some graph for which any spanner with at most n+n/t edges has stretch at least  $\Omega(t \log n)$ . Furthermore, computing any spanner with n+n/t edges requires  $\Omega(t \cdot \log n)$  rounds of distributed computation [Elk07, DGPV09]. Indeed, these lower bounds hold even for unweighted graphs, randomized computations, and the LOCAL model. The LOCAL model is much more permissive than the CONGEST model considered in this paper in that it allows nodes to exchange messages of unbounded size and perform arbitrarily complex local computations.

#### **1.2** Spanners: Our Contribution

As a first-order summary of prior work, to the best of our knowledge, there are no known deterministic distributed algorithms for ultra-sparse or even sparse spanners with O(n) edges, even if we restrict ourselves to unweighted graphs. Furthermore, even for spanners of higher density and small stretch 2k-1 for  $k \leq \log n$ ,

<sup>&</sup>lt;sup>1</sup>In the original paper [BS07], it is claimed that the spanner has  $O(nk + n^{1+\frac{1}{k}})$  edges. Pettie [Pet10] noticed a gap in their argument and provides a bound but with an extra log k factor. The original claim remains unproven.

the number of edges achieved by the state of the art deterministic algorithm [GK18] is higher than the corresponding randomized algorithms [BS07].

Our contributions resolve this situation:

- 1. We present deterministic algorithms that compute ultra-sparse spanners with n + o(n) edges and close to  $O(\log n)$  stretch, even in weighted graphs. We do this by showing a reduction from the ultrasparse case to sparse spanners, and by derandomizing a randomized sparse spanner construction of Pettie [Pet10], and extending it to weighted graphs (see Table 1).
- 2. For general stretch parameter k, we provide a derandomization of Baswana-Sen [BS07] that matches its stretch-size tradeoff and improves on [GK18] (see Table 2).
- 3. These derandomization-based algorithms use local computations that exceed  $O(mpoly \log n)$  and thus are less suitable for adaptation to the PRAM model of parallel computation. To avoid that, we show in addition a work-efficient reduction from spanners to so-called weak-diameter clusterings. Using this connection and our reduction from ultra-sparse spanners to spanners of higher density, we achieve a work-efficient distributed/parallel algorithm for ultra-sparse weighted spanners with poly(log n) stretch.

| Paper      | Stretch                       | Size              | Weighted?    | Deterministic? | # rounds                       |
|------------|-------------------------------|-------------------|--------------|----------------|--------------------------------|
| [Pet10]    | $O(\log n + 2^{\log^* n})$    | O(n)              | ×            | ×              | $O(\log n) \cdot 2^{\log^* n}$ |
| [EN18]     | $O(\log n)$                   | n + o(n) expected | ×            | ×              | $O(\log n)$                    |
| This paper | $O(\log n + 2^{O(\log^* n)})$ | n + o(n)          | $\checkmark$ | $\checkmark$   | $\operatorname{poly}(\log n)$  |

In the next three subsections, we elaborate on these contributions.

**Table 1:** This table summarizes the relevant results for construction of spanners that are very sparse. Ideally, we aim for ultra-sparse spanners with size n + o(n). This property is crucial in some applications, as discussed in the introduction. We highlight that our result is the first ultra-sparse spanner construction that is deterministic. Our result is also the first that handles weighted graphs.

| Paper      | Stretch | Size                      | Weighted?    | Deterministic? | # rounds                       |
|------------|---------|---------------------------|--------------|----------------|--------------------------------|
| [BS07]     | 2k - 1  | $O(n^{1+1/k}\log k + nk)$ | ×            | ×              | O(k)                           |
| [BS07]     | 2k - 1  | $O(n^{1+1/k}k)$           | $\checkmark$ | ×              | O(k)                           |
| [GK18]     | 2k - 1  | $O(n^{1+1/k}k\log n)$     | ×            | $\checkmark$   | $\operatorname{poly}(\log(n))$ |
| This paper | 2k - 1  | $O(n^{1+1/k}\log k + nk)$ | ×            | $\checkmark$   | poly(log(n))                   |
| This paper | 2k - 1  | $O(n^{1+1/k}k)$           | $\checkmark$ | $\checkmark$   | $\operatorname{poly}(\log(n))$ |

**Table 2:** This table summarizes the relevant results for construction of spanners with small stretch (think of k as a constant). We highlight that our results match the bounds of [BS07] while our results are deterministic.

#### **1.2.1** Reduction from Ultra-Sparse Spanners to (Sparse) Spanners

Our first contribution is a general reduction from the problem of computing ultra-sparse spanners in weighted graphs to the problem of computing their spanners. This allows us to efficiently move the extra factor in the number of edges of the final spanner to its stretch:

Before we state the theorem, we briefly introduce the notion of a *cluster graph*: Given a graph G, a cluster graph H is a graph that we get by contracting some of its disjoint subgraphs that we call clusters. If all of those clusters are connected and have radius at most r, we say that H is an r-cluster-graph of G. We say that a distributed algorithm works on an N-node r-cluster-graph in T(N, r) rounds if the underlying communication network is G, but the output of the algorithm should be for a given r-cluster-graph H of G with |V(H)| = N. For a more precise definition of a cluster graph and an algorithm that operates on a cluster graph see Section 2.

**Theorem 1.2.** Suppose that there exists a deterministic distributed algorithm A which computes an  $\alpha$ spanner with  $N \cdot s(N)$  edges for any N-node weighted r-cluster-graph in T(N,r) rounds. Then, for any  $t \geq 1$ , there is a deterministic distributed algorithm A' that computes an  $O(t \cdot s(n) \cdot \alpha)$ -spanner with n + n/tedges for any n-node weighted graph, in

$$O\left(t \cdot s(n) \cdot \log^* n + T\left(\frac{n}{t \cdot s(n)}, O\left(t \cdot s\left(n\right)\right)\right)\right)$$

rounds.

We present the proof of Theorem 1.2 in Section 4. We note that the stretch-size tradeoff in our reduction is asymptotically optimal. Indeed, plugging an  $O(\log n)$ -spanner construction with O(n) edges [DGPV09, EN18] in our reduction would result in an ultra-spanner with n + n/t edges and stretch  $O(t \log n)$ , which is asymptotically optimal, even in unweighted graphs.

**Example Implication.** As one concrete corollary of Theorem 1.2, let us discuss how this reduction gives an improvement on the results of Li [Li20]. By applying Theorem 1.2 to the Pettie's [Pet10] randomized sparse spanner construction, we obtain a randomized algorithm for weighted ultra-sparse spanner with n+n/t edges and  $O(t \log n \cdot 4^{\log^* n})$  stretch. More concretely, we have:

**Theorem 1.3.** There is a randomized distributed algorithm that computes an  $O(t \log n \cdot 4^{\log^* n})$ -spanner of any n-node weighted graph and any  $t \ge 1$  with expected n + n/t edges. The algorithm works with high probability and it runs in  $t \cdot \operatorname{polylog}(n)$  rounds in the CONGEST model. A parallel variant of this algorithm works with  $\operatorname{polylog}(n)$  depth and  $m \cdot \operatorname{polylog}(n)$  work in the PRAM model.

This stretch is optimal up to the  $4^{\log^* n}$  term. We comment that Pettie presents the algorithm only for unweighted graphs, but we show in Theorem 1.5 that by simple modifications, we can extend the algorithm to weighted graphs, with only a minimal loss of increasing the stretch factor from  $O(\log n \cdot 2^{\log^* n})$  to  $O(\log n \cdot 4^{\log^* n})$ . Also, as the algorithm readily runs in the PRAM model with polylog(n) depth and  $m \cdot$ polylog(n) work. This improves on the PRAM ultra-sparse spanner construction of Li [Li20], which had stretch  $O(t^2 \log^3 n \log^2 \log n)$  for n + n/t edges.

In the next two subsubsections we discuss how invoking the reduction of Theorem 1.2 atop our (deterministic) spanner constructions leads to our final results on ultra-sparse spanners.

#### 1.2.2 Spanners via Derandomization

Our second contribution is derandomization-based deterministic spanner constructions (see Section 3). We present these in two groups: (A) focused on stretch and especially for sub-logarithmic values of stretch, and (B) focused on low sparsity for roughly logarithmic values of stretch. We note that the algorithms below achieve a good stretch for the given number of edges, but they involve large computations in each node.

(A) Low stretch spanners, by derandomizing Baswana-Sen [BS07]: In the first category, our focus is on small values of stretch. We show polylog(n)-round deterministic algorithms for (2k-1)-spanners, with  $O(nk + n^{1+1/k} \log k)$  edges in unweighted graphs and  $O(n^{1+1/k}k)$  edges in weighted graphs:

**Theorem 1.4.** There are polylog(n)-rounds deterministic distributed algorithms that compute a (2k - 1)spanner with  $O(nk + n^{1+1/k} \log k)$  and  $O(nk + n^{1+1/k} \log k)$  edges for unweighted and weighted graphs, respectively.

Our approach is inspired by the work of Ghaffari and Kuhn [GK18] on derandomizing the Baswana-Sen randomized algorithm. Ghaffari and Kuhn [GK18] get the same stretch but with a worse sparsity bound of  $O(n^{1+1/k}k \log n)$  edges, and only for unweighted graphs. In contrast, our bounds match the best known analysis of the randomized Baswana-Sen algorithm, in both unweighted and weighted graphs.

(B) Low sparsity spanners, by Derandomizing Pettie [Pet10]: In the second category, our focus is on the sparsity of the spanner, and we show polylog(n)-round deterministic algorithms that compute sparse spanners with O(n) edges, in both unweighted and weighted cases, which achieve a stretch of almost  $O(\log n)$ :

**Theorem 1.5.** There are polylog(n) rounds deterministic distributed algorithms that compute a spanner with O(n) edges. For unweighted graphs, the stretch of the spanner is  $O(\log n \cdot 2^{\log^* n})$  and for weighted graphs, it is  $O(\log n \cdot 4^{\log^* n})$ .

These deterministic algorithms are also obtained via derandomization, but when applied to Pettie's randomized algorithm [Pet10]. We also note that the weighted sparse spanner stated in Theorem 1.5 was not stated in prior work even for randomized algorithms; we obtain this result via a minor modification of Pettie's unweighted approach.

**Implications.** Via Theorem 1.2 with Theorem 1.5, we get the first polylog(n)-round deterministic distributed algorithm for ultra-sparse spanners:

**Theorem 1.6.** There are  $t \cdot \operatorname{polylog}(n)$  distributed deterministic algorithms that compute a spanner with n + n/t edges with stretch  $O\left(t \log n \cdot 2^{\log^* n}\right)$  for unweighted graphs and with stretch  $O\left(t \log n \cdot 4^{\log^* n}\right)$  for weighted graphs.

#### 1.2.3 Spanners via Low-Diameter Clusterings

As our third contribution, we show a connection with the so-called weak-diameter clusterings. As a result of this connection, we will achieve deterministic spanner constructions that are work efficient—i.e., doing only m poly log n computations—but have weaker bounds on the stretch; the latter is merely due the sub-optimal bounds in the state of the art radius for deterministic low-diameter clusterings.

Our most general result in this direction is the following near-optimal end-to-end reduction of spanner construction to clustering construction. Before we state the somewhat technical theorem, we need to briefly introduce the notion of weak-diameter and separated clusterings. A clustering is a collection of disjoint vertex sets (clusters). Next, think actually of every cluster as a pair of the vertex set  $C \subseteq V(G)$  and a tree  $T_C$  with  $V(T_C) \supseteq C$ . That is, each cluster also carries a tree with it collecting its nodes. A clustering has weak-diameter D if every cluster C of it has the property that the diameter of  $T_C$  is at most D. Namely, we require any two nodes of C to be close in G, but the cluster C itself might be even disconnected. A clustering is k-separated if the distance of any two different clusters of it is at least k. Finally, the average overlap  $\xi_{AVG}$  of a clustering  $\{C_1, \ldots, C_t\}$  is defined as  $\xi_{AVG} = \left(\sum_{i=1}^t |V(T_{C_i})|\right)/n$ . That is, it is the average overlap of the trees of the clustering.

**Theorem 1.7.** Suppose there exists an algorithm A which for any unweighted n-vertex graph finds a 3separated clustering with weak-diameter D(n) with average overlap at most  $\xi_{AVG}(n)$ . Then, there is an algorithm A' that builds a  $\beta$ -spanner H on an unweighted graph such that (1)  $|E(H)| = O(\xi_{AVG}(n)n)$ , (2)  $\beta = O(D(n))$ .

Furthermore, if A is deterministic then so is A' and

- If A requires at most T(n) CONGEST rounds then A' requires  $O(\log(n)T(n))$  CONGEST rounds.
- If A requires at most T(n,r) CONGEST rounds on a r-cluster-graph, then A' requires  $O(\log(n)(T(n,r)+r))$ CONGEST rounds on a r-cluster-graph.
- If A is a PRAM algorithm with T(n) depth and W(m, n) work then A' is a PRAM algorithm with depth  $O(\log(n)T(n))$  and work  $O(\operatorname{poly}(\log(n))(W(m, n) + m + n\xi_{AVG}(n)).$

Theorem 1.7 is optimal, up to constants, in terms of its density-stretch tradeoff. This is because there are 3-separated clusterings with diameter  $\log n$  and constant overlap. These would give the desired stretch factor of  $O(\log n)$ . Also, Theorem 1.7 can be extended to produce spanners in weighted graphs with an additional  $O(\log(U+1))$  overhead in the number of edges in the spanner, by applying the standard reduction mentioned earlier in Section 1.1.

**Implications.** Plugging the state-of-the-art deterministic distributed clusterings [RG20] into Theorem 1.7, and then applying Theorem 1.2 on top, gives the following result:

**Theorem 1.8.** There is a deterministic work-efficient CONGEST algorithm that, given any n-vertex weighted graph G and  $t \ge 1$ , computes in  $O(t \log^{10} n)$  rounds an ultra-sparse spanner with n + n/t edges and stretch  $O(t \log^4 n \log(U + 1))$ , where  $U \ge 1$  is the aspect ratio of the weights. There is also a deterministic PRAM algorithm that computes such a spanner in polylog(n)-time and with  $m \cdot polylog(n)$  work.

#### 1.3 Connectivity Certificates: Background and Our Contribution

**Definition and Motivation.** For a graph G = (V, E), a k-connectivity certificate is a spanning subgraph H = (V, E') such that if G is k-edge-connected, so is H. This high connectivity property can be relevant for resilience against link failures or for higher communication capacity. Concretely, if G can withstand the failure/removal of any k of its edges and it would still remain connected, the same should be true also for H. Similarly, if the minimum cut in G—which is in some sense the communication bottleneck in the network as it is the smallest number of edges between one set of nodes and the rest—has size k, the same should be true for H.

Any k-edge-connected graph must have at least nk/2 edges, as each node must have degree at least k. Hence, the sparsest that the connectivity certificate H can be is to have nk/2 edges. Considering this, any connectivity certificate that has O(nk) edges is called a *sparse connectivity certificate*. We note that an approximation version of this problem has also been widely studied under the notion of k-edge-connected spanning subgraph (k-ECSS), where for each given k-edge-connected graph G, the objective is to compute the sparsest possible k-edge-connected subgraph H, and the performance is measured in terms of the ratio of the number of edges in the computed subgraph H to the smallest possible.

State of the Art. Thurimella [Thu97] gave the first distributed algorithm for computing a sparse connectivity certificate, with k(n-1) edges, and the round complexity of  $O(k(D + \sqrt{n}))$  in the CONGEST model. Here, D denotes the network diameter. Primarily coming from the side of the k-ECSS problem, Censor-Hillel and Dory [CHD20] investigated the special case of k = 2 and they gave an O(D) round randomized distributed algorithm that computes a 2-connectivity certificate with O(n) edges with high probability, i.e., an O(1) approximation for 2-ECSS. Then, Dory [Dor18] provided an  $O(D \log^3 n)$  round randomized distributed algorithm that computes a 3-connectivity certificate with  $O(n \log n)$  edges with high probability, thus an  $O(\log n)$  approximation for 3-ECSS. Daga et al. [DHNS19] improved the algorithm of Thurimella [Thu97] to achieve a round complexity of  $\tilde{O}(D + \sqrt{nk})$ ).

Finally, Parter [Par19] improved the complexity significantly by providing an  $O(kpoly(\log n))$ -round randomized algorithm that computes a k-connectivity certificate with O(kn) edges, with high probability. Thus, this also gives an O(1) approximation for the k-ECSS problem in  $O(kpoly(\log n))$  rounds. Notice that this complexity can still grow larger even up to  $npoly(\log n)$  as k grows. Parter [Par19] also gave a faster  $poly(\log n)$ -round algorithm but for a slightly weaker notion of approximate-certificate. Concretely, given the parameter k, the algorithm runs in  $poly(\log n)/\varepsilon^2$  rounds and computes a spanning subgraph H with O(kn) edges such that, if G is k-edge-connected, then H is  $k(1 - \varepsilon)$ -edge-connected, with high probability. It remained open whether the same  $poly(\log n)$  round complexity can also suffice for sparse connectivity certificates, without reducing the connectivity value to the approximate version.

**Our contribution.** We present a simple algorithm that resolves the above question. Concretely, we show a poly $(\log n)$ -round randomized distributed algorithm that computes a k-connectivity certificate with (1 + o(1))kn edges, with high probability.

**Theorem 1.9.** For any  $\varepsilon < 1/2$ , there is a randomized distributed algorithm that computes a k-connectivity certificate in  $\frac{\text{polylog}(n)}{\varepsilon^3}$  rounds with  $kn(1+\varepsilon)$  edges.

This sparsity is within a 2 + o(1) factor the best possible as any k-connectivity certificate needs at least kn/2 edges. Hence, our algorithm gives a 2 + o(1) approximation for the (unweighted) k-ECSS problem in poly(log n) rounds. Due to the space limitations, the entire proof of Theorem 1.9 is deferred to Appendix G.

We note that, similar to Daga et al. [DHNS19] and Parter [Par19], we also use Karger's edgesampling [Kar99] to split the graph into many edge disjoint parts and paralellize the work. Daga et al. [DHNS19] perform a tree packing in each of these parts. Parter packs sparse spanners, in each part, but loses a  $1 + \varepsilon$  in the connectivity. We also used packings of ultra-sparse spanners, and we show by a simple case analysis cut sizes that we reach the exact k-connectivity, without the  $1 + \varepsilon$  factor loss. We note that this  $1 + \varepsilon$  factor loss can be important in applications, e.g., in the framework of Daga et al. [DHNS19] for computing min-cut, this loss would downgrade exact min-cut algorithms to  $1 + \varepsilon$  approximation, which is a much easier problem.

#### 2 Basic Notations

The input is a graph G, with n nodes and m edges. If G is weighted, we assume all edge weights are non-negative and are bounded by poly(n). We denote via  $d_G(u, v)$  the weight (i.e., length) of the shortest path between two nodes  $u, v \in V(G)$ . We may drop the subscript when the graph is clear from context. The distance function extends to sets by  $d(A, B) = \min_{a \in A, b \in B} d(a, b)$  and we write d(u, S) instead of  $d(\{u\}, S)$ . We denote by diam(G) the diameter of a graph G which is the maximum pairwise distances between nodes of G. We may use diam(C) to denote the diameter of the induced subgraph G[C].

For a weighted graph G, a cluster  $C \subseteq V(G)$  is simply a subset of its nodes. A clustering C is a set of disjoint clusters. We say that a clustering C is a partition if  $V(G) = \bigcup_{C \in C} C$ . The clustering (partition) C is an *r*-clustering (*r*-parition) if there is a rooted tree with radius at most r in the subgraph induced by each of its cluster. A rooted tree has radius r if the maximum hop-distance (i.e. the number of edges in the shortest path) between a leaf and the root is r. We say an edge is a boundary-edge of a cluster C if exactly one of its endpoints is in C. An edge is an *inside-edge* of C if both of its endpoints are in C. A graph H is called an r-cluster-graph of G if it is obtained from G by contracting each cluster of an r-clustering to a single node. If v is a node of H, then  $inv^{H \to G}(v)$  represents the set of nodes that are contracted to v. Let us emphasize that in the definition of r-clustering and r-cluster-graphs, the parameter r does not depend on edge weights and only depends on hop-distances.

### 3 Derandomization

**Baswana-Sen Algorithm** [BS07]. The spanner is constructed in k iterations. At first, all nodes and edges of G are alive. The input of iteration i is a (i - 1)-partition of the nodes that remain alive after the first i - 1 iterations. During one iteration, some nodes die. When a node dies, all of its incident edges die as well. Moreover, some edges incident to an alive node may die during one iteration. By adding a relatively small set of edges to the spanner in iteration i, we ensure that all edges that die in this iteration have a stretch at most 2i - 1. We also ensure that all nodes die after the last iteration. So all edges die and as a result, we have a spanner with stretch 2k - 1 at the end. The output of iteration i is the input of iteration i + 1. The input for the first iteration is the trivial partition of all nodes (one cluster for each node). The details of iteration i are given in the following:

- (1) We sample each cluster with probability  $p = n^{-1/k}$  for  $i \le k 1$ . In the last iteration (when i = k), we sample each cluster with probability zero, i.e. no cluster is sampled.
- (2) Each node adds some (possibly zero) edges to the spanner. For a node v with d adjacent clusters, let  $e_i$  be the edge with the smallest weight  $w_i$  among all edges between v and its *i*-th adjacent cluster  $C_i$  and assume  $w_1 \leq w_2 \leq \cdots \leq w_d$ . A cluster C is adjacent to v if v has a neighbor in C. If v is in a sampled cluster, it does nothing. If v is an unsampled cluster, let i be the smallest integer such that  $C_i$  is sampled. If there is such an i, node v joins  $C_i$ . The edge  $e_i$  along with all  $e_j$  for which  $w_j$  is strictly less than  $w_i$  are added to the spanner. If there is no such an i, node v dies and all d edges  $e_1, \ldots, e_d$  are added to the spanner. In all of the above cases for v, whenever we add an edge  $e_i$  to the spanner, all edges between v and  $C_i$  die.
- (3) The output of iteration i is an *i*-partition on the set of nodes that are alive after iteration i. The partition has one cluster C' for each sampled cluster C along with all nodes that are joined to it.

When a node v joins C, its parent in C is the node to which it has an edge with smallest weight (note that this edge is in the spanner from step (2)). Observe that the radius of C' is at most the radius of C plus one.

In the following lemma we collect deterministic properties of the above construction proven in [BS07].

**Lemma 3.1** ([BS07]). For any cluster C in the output of iteration i, we have: Radius of C is at most i. For any alive boundary-edge  $\{u \notin C, v \in C\}$  of C with weight w, all edges in the unique path from v to the root of C have weight at most w. For any alive inside-edge  $\{u \in C, v \in C\}$  of C with weight w, all edges in the unique path between u and v in C have weight at most w. It holds that all dead edges in iteration i have stretch 2i - 1. So the final spanner has stretch 2k - 1 as all nodes die in the last iteration. All these properties are deterministic in the sense that they hold regardless of the way we sample clusters.

The lemma above provides a stretch guarantee we need. It remains to show the expected size of final spanner. For that, let us first define a hitting-event. We write all the stretch analysis in terms of these kind of events as we need this for derandomization. We discuss the reason in a moment.

**Definition 3.2.** A binary random variable E is a *hitting-event* over the set of binary random variables  $\{X_1, \ldots, X_c\}$ , if there is a subset  $S \subseteq [c]$  such that  $E = \bigvee_{j \in [c]} X_j$ .

Number of clusters. For i > 1, we have

$$\mathbb{E}[c_i] = \sum_{j \in [c_i]} \Pr[X_j^{(i)} = 1] = p \cdot \mathbb{E}[c_{i-1}].$$

Since  $c_1 = n$ , we have  $\mathbb{E}[c_i] = np^{i-1}$ .

**Last iteration.** For i = k, there are  $np^{k-1} = n^{1/k}$  clusters in expectation. Since there are n nodes, at most  $n^{1+1/k}$  edges are added in expectation for i = k.

**First** k-1 **iterations.** Consider an alive node v with d adjacent clusters in iteration i. Without loss of generality, assume that its j-th adjacent cluster is  $C_j^{(i)}$  and the smallest edge weight between v and  $C_j^{(i)}$  is larger or equal than the smallest edge weight between v and  $C_{j-1}^{(i)}$ . Let  $X_j^{(i)}$  be the indicator random variable that the cluster  $C_j^{(i)}$  is selected in the *i*-th iteration. The node v adds at most

$$1 + \sum_{j \in [d]} \Pr[\bigvee_{\ell \in [j-1]} X_{\ell}^{(i)} = 0] < 1 + \sum_{j=0}^{\infty} (1-p)^j = O(1/p)$$

edges in expectation. There are n nodes, so each iteration adds O(n/p) edges. In total and including the last iteration, we get the claimed size bound  $O(nk/p) = O(n^{1+1/k}k)$ .

**Unweighted graphs.** In this case, when a node remains alive during an iteration, it adds at most one edge to the spanner (since we only add edges with **strictly** smaller weights). So, in total, at most n(k-1) edges are added from nodes that remain alive during an iteration. For the contribution of dead nodes, consider an alive node v with d adjacent clusters in iteration i. Let  $S \subseteq [c_i]$  with |S| = d be the set of adjacent clusters of S. Node v dies in iteration i with probability at most

$$\Pr[\bigvee_{j \in S} X_j^{(i)} = 0] = (1 - p)^d \le e^{-pd}.$$

So a node can add up to  $de^{-pd}$  edges, in expectation. Function  $x \to xe^{-px}$  is maximized at x = 1/p where its value is  $\Theta(1/p)$ . From this and since there are *n* nodes, each iteration adds O(n/p) edges, in expectation. This results in the final expected size of O(nk/p). This is the same as in the weighted case.

To improve this bound, suppose d is larger than a threshold  $\tau = \ln(k)/p$ . Function  $x \to xe^{-px}$  is decreasing in the range  $[1/p, +\infty)$ . Thus the total expected contribution of such v in iteration i is  $n \cdot (\tau e^{-p\tau}) =$ 

 $O(n \log k/(pk))$ . There are k-1 iterations, so  $O(n \log k/p)$  edges added from these nodes in expectation. The bound for nodes with  $d \leq \tau$  is deterministic. The overall contribution of these nodes in all the first k-1 iterations is at most  $n\tau$  (since each node dies only once and there are at most n nodes). Considering the contribution of  $n^{1+1/k}$  edges in the last iteration, the expected size of the final spanner is the claimed bound  $O(nk + n^{1+1/k} \log k)$ .

**High-degree nodes.** For implementation, we need one more ingredient. An alive node v with d adjacent cluster in iteration i is called *high-degree* if  $d \ge \xi = 10 \ln n/p$ . The probability that such a node dies in iteration i is a hitting-event over  $\{X_j^{(i)}\}_{j \in [c_i]}$  and is at most  $(1-p)^{\xi} \le e^{-p\xi} = n^{-10}$ . So by union bound, no high-degree node dies during iteration i with probability at least  $1 - n^{-9}$ .

**Reducing randomness.** In the above, we assume full independence for sampling in each iteration and between the iterations. The analysis still goes through with less randomness inside each iteration but keeping the independence across the iterations. More concretely, for each *i*, there is a distribution  $\mathcal{P}$  from which we can sample with  $O(\log n \log \log n)$  random bits to generate  $\{X_j^{(i)}\}_{j \in [c_i]}$  such that it approximates each hitting-event with additive factor  $1/\operatorname{poly}(n)$ . To be more precise, let *E* be a hitting-event over  $\{X_j^{(i)}\}_{j \in [c_i]}$  and let  $p_E$  be the probability that *E* is 1 assuming that each  $X_j^{(i)}$ s is sampled independently of the other. Let  $\tilde{p}_E$  be  $\Pr_{\mathcal{P}}[E=1]$ . Then  $|p_E - \tilde{p}_E| \leq n^{-10}$ . Let us emphasize that the distribution  $\mathcal{P}$  is independent of *E*. For details, please see Appendix B. In derandomization, we use the method of conditional expectation where we fix bits of the random seed one by one. So for a polylogarithmic rounds algorithm, such a short seed is needed.

The only randomized part of each iteration is cluster sampling. If we derandomize this part, the whole algorithm becomes deterministic. The next lemma gives a formal statement of our derandomization for sampling. It is written in a parametric form as we need it later for computing linear size spanners. There is no randomness in the last iteration, so it is already deterministic and we ignore that. A direct implication of the following lemma is Theorem 1.4.

**Lemma 3.3.** For any positive integer g and real number  $p \in (1/n, 1)$ , there is a deterministic distributed algorithm that "simulates" g iterations of Baswana-Sen with sampling probability p in  $g^2 \cdot \operatorname{polylog}(n)$  rounds. That is, (1) it adds O(ng/p) and  $O(ng + \frac{n \log g}{p})$  edges to the spanner for weighted and unweighted graphs, respectively, (2) the number of clusters in the output of last iteration is at most  $np^g$ , (3) no high-degree node dies in any of g iterations.

*Proof.* We derandomize each iteration separately. For iteration  $i \in [g]$ , we have three objectives:

- (a) The number of clusters in the output of iteration *i*, variable  $c_{i+1}$ , is at most  $np^i$ .
- (b) Let  $\iota$  be a large enough constant. For weighted graphs, we should add at most  $\iota n/p$  edges. For unweighted graphs, nodes that die in the iteration and has at least  $\tau = \ln g/p$  adjacent clusters should add at most  $(\iota n \log g)/(pg)$  edges. From earlier discussion, the contribution of nodes that remain alive or the one that dies but has  $d \leq \tau$  is within the final size budget, deterministically. So we ignore them.
- (c) Ensuring that no high-degree node (having more than  $10 \ln n/p$  adjacent clusters) dies.

We combine all our objectives into one random variable (also known as utility function). For weighted graphs, we define  $U_i^w$  as

$$U_i^{w} = \iota/p^{i+1} \sum_{j \in [c_i]} X_j^{(i)} + \sum_{v \in V_i} (b_v + n^5 h_v).$$
(3.1)

For unweighted graphs, we define  $U_i^{\text{uw}}$  as

$$U_i^{\text{uw}} = (\iota \ln g) / (gp^{i+1}) \sum_{j \in [c_i]} X_j^{(i)} + \sum_{v \in V_i} (b_v + n^5 h_v).$$
(3.2)

In the definitions of  $U_i^{w}$  and  $U_i^{uw}$ , set of alive nodes in iteration *i* is denoted by  $V_i$ . Random variable  $b_v$  is the number of edges added by node v (in the unweighted case, we set  $b_v$  to zero if it is an ignored node). Random variable  $h_v$  is one if v is high-degree and dies in iteration *i*. It is zero otherwise. Suppose there is an assignment of  $\{X_j^{(i)}\}_{j\in[c_i]}$  that makes  $U_i^{w}$  at most in/p. So in this assignment  $c_{i+1} = \sum_{j\in[c_i]} X_j^{(i)} \le np^i$  and  $\sum_{v\in V_i} b_v \le in/p$ . Moreover,  $\sum_{v\in V_i} h_v = 0$  as  $in/p = O(n^2)$  and the summation should be an integer. So all three required conditions (a), (b), and (c) hold. Similarly, for the unweighted case, an assignment with  $U^{uw} \le (in \log g)(pg)$  suffices. But why does such an assignment exist? For that, suppose we sample each  $X_j^{(i)}$  independently with probability p/4 (and not p). From induction, we know that  $c_i \le np^{i-1}$  so  $\mathbb{E}[\sum_{j\in[c_i]} X_j^{(i)}] \le np^i/4$ , and using earlier discussion, we know that  $\sum_{v\in V_i} \mathbb{E}[b_v] \le in/(4p)$  and  $\sum_{v\in V_i} \mathbb{E}[h_v] \le n^{-5}$ . So, we have:

$$\mathbb{E}[U_i^{\mathsf{w}}] \le \iota n/(4p) + \iota n/(4p) + n^{-5} < \iota n/p.$$

Similarly, we can show that  $\mathbb{E}[U_i^{uw}] < (\iota n \log g)/(pg)$ .

So there is such a good assignment for  $\{X_j^{(i)}\}_{j \in [c_i]}$ . Observe that each utility function can be written as  $O(n^2)$  hitting-events, so if we approximate full independence with the distribution in Appendix B, the above analysis still goes through as it only incurs  $O(n^2 \cdot n^{-10}) = O(n^{-8})$  error. To find such an assignment deterministically in polylog(n) rounds, we use the method of conditional expectation running over a network decomposition. The details are in Appendix C.

**Computational Aspects.** We do not get work-efficient algorithms in the sense of having only  $m \, \cdot \, \text{polylog}(n)$  computation. In fact, the amount of computation of each node is slightly super-polynomial  $2^{O(\log n \log \log n)} = n^{O(\log \log n)}$ . However, this is somewhat similar to (and only better than) recent works that use conditional expectation for derandomization. In [GK18, DKM19], they use  $O(\log n)$ -wise independence, and in [PY18b], they use a distribution with  $\log n(\log \log n)^3$  random bits. Hence, they need  $n^{O(\log n)}$  and  $n^{O((\log \log n)^3)}$  local computations, respectively.

#### 3.1 Linear Size Spanners

We also provide a polylog(n) rounds derandomization of Pettie's algorithm [Pet10] for weighted and unweighted graphs. Our result is stated in Theorem 1.5 and its proofs is deferred to Appendix D. In the following, we define *stretch-friendly* clustering which plays the key role in the stretch analysis of this algorithm and Section 4.

**Definition 3.4.** An *r*-cluster *C* is stretch-friendly if for any boundary-edge  $\{u \notin C, v \in C\}$  of *C* with weight w, all edges in the unique path from v to the root of *C* have weight at most w. Moreover, for any inside-edge  $\{u \in C, v \in C\}$  of *C* with weight w, all edges in the unique path between u and v in *C* have weight at most w. An *r*-clustering (*r*-partition) is stretch-friendly if all of its clusters are stretch-friendly.

**Observation 3.5.** Let C be a stretch-friendly r-partition of G and H be the r-cluster-graph induced by C. Denote by  $T_C$  the set of edges in the tree of cluster  $C \in C$ . The union of an  $\alpha$ -spanner of H with  $\bigcup_{C \in C} T_C$  is a  $((2r+1)(\alpha+1)-1)$ -spanner of G.

# 4 Deterministic Ultra-Sparse to Sparse Reduction

The main tool of this section is the following lemma.

**Lemma 4.1.** There is a deterministic distributed algorithm that computes a stretch-friendly O(t)-partition with at most n/t clusters in  $O(t \log^* n)$ .

To our knowledge, such a result is only known for unweighted graphs, (see Kutten and Peleg [KP98]). Note that in the unweighted case, any clustering is stretch-friendly. Round complexity of [KP98] is also  $O(t \log^* n)$  and our algorithm is arguably simpler.

*Remark.* The algorithm of Lemma 4.1 can be run in polylog(n) depth and  $m \cdot polylog(n)$  work in the PRAM model.

Next, we prove Theorem 1.2 and Theorem 1.6.

Proof of Theorem 1.2. First, we find a stretch-friendly  $O(t \cdot s(n))$ -clustering  $\mathcal{C}$  using Lemma 4.1. Then, using A, we compute an  $\alpha$ -spanner of a  $O(t \cdot s(n))$ -cluster-graph that is induced by  $\mathcal{C}$ . This spanner along with trees corresponding to the clusters in  $\mathcal{C}$  is a  $O(t \cdot s(n)\alpha)$ -spanner of G according to Observation 3.5. The spanner has n + O(n/t) edges since at most n - 1 edges are in the union of trees of  $\mathcal{C}$  and algorithm A adds  $s(n/(t \cdot s(n)))\frac{n}{t \cdot s(n)} = O(n/t)$  edges. Multiplying t by a large enough constant gives a spanner with  $O(t \cdot s(n)\alpha)$  stretch and n + n/t edges. The round complexity follows from the round complexity of A and Lemma 4.1.

Proof of Theorem 1.6. We apply Theorem 1.2 with algorithm A being the linear size algorithm of Theorem 1.5. For this A, the output has size O(N), so s(N) = O(1). For function T(N, r), recall that A is split into  $O(\log^* n)$  phases. The input of each phase is a cluster-graph. It is discussed in the proof of Theorem 1.5 that the radius of input for each phase is linearly multiplied in its round complexity as it stretches the dilation by r. So  $T(N, r) = r \cdot \text{polylog}(N)$  which completes the proof.

The algorithm of Lemma 4.1 is described in the following. Its proof is deferred to Appendix E. We gradually construct the partition in  $\lceil \log t \rceil$  iterations. The output of iteration *i* (input of iteration *i*+1) is a stretch-friendly  $(3 \cdot 2^i - 1)$ -partition with each cluster has size at least  $2^i$ . The input of first iteration is the trivial partition (one cluster for each node). Details of iteration *i* with input partition C are as follows:

- (1) Each cluster  $C \in \mathcal{C}$  computes it size.
- (2) Each cluster C finds its minimum weight boundary-edge (breaking ties arbitrarily) and orients it out from C.
- (3) A 3-coloring of the cluster-graph induced by  $\mathcal{C}$  considering only oriented edges is computed.
- (4) A maximal matching between *small* clusters is computed. A cluster is *small* if its size is less than  $2^i$  and is *large* otherwise.
- (5) A partition  $\mathcal{C}'$  (input of iteration i+1) is created: First, we merge matched clusters and put one cluster for each merged cluster in  $\mathcal{C}'$ . Then, each large cluster is added to  $\mathcal{C}'$ . In the end, each unmatched cluster is merged to the cluster of its outgoing neighbor in  $\mathcal{C}'$  (this neighbor is in  $\mathcal{C}'$  as it is either a matched small cluster or a large cluster). When we merge two clusters  $C_1$  and  $C_2$  with an edge oriented from  $C_1$  to  $C_2$ , the root of the new cluster is the root of  $C_2$ .

# 5 Deterministic Unweighted Spanner via Low-Diameter Clusterings

In this section, we sketch how one can efficiently compute ultra-sparse spanners with n + n/t edges by computing  $O(\log n)$  t'-separated strong diameter clusterings for  $t' = O(t \log n)$ .

**Definition 5.1** (*t*-separated Strong Diameter Clustering). Assume we are given an unweighted graph G and a parameter t > 0. A *t*-separated low diameter clustering with strong diameter D is a clustering C such that:

- 1. For each  $C \in \mathcal{C}$  we have diam $(C) \leq D$ .
- 2. For each  $C_1 \neq C_2 \in \mathcal{C}$  we have  $d(C_1, C_2) \geq t$ .
- 3. We have  $|\bigcup_{C \in \mathcal{C}} C| \ge n/2$ .

The high-level idea for the spanner construction is to compute a low-diameter clustering covering all the nodes such that the total number of neighboring clusters is at most n/t. Then, adding for each cluster a low-diameter spanning tree to the spanner and for each of the at most n/t neighboring clusters an arbitrary

edge between the two clusters results in a spanner with n + n/t edges and the stretch being on the order of the diameter of each cluster.

Note that obtaining a clustering that covers all the nodes is simple: we iteratively compute a low-diameter clustering of the yet unclustered nodes. In each iteration, the number of unclustered nodes decreases by a factor of two, and therefore each node is clustered after  $O(\log n)$  iterations.

To achieve a small number of neighboring clusters, we can do the following in each of the  $O(\log n)$  iterations. Start with a clustering with separation  $100t \log n$ . Now, repeatedly grow a given cluster by adding its neighbors to it until it would grow less than by a multiplicative factor of 1 + 1/t. The growth stops after at most  $10t \log n$  steps as  $(1 + 1/t)^{10t \log n} > n$ . The clusters still remain well-separated and the property that each final cluster C neighbors with at most |C|/t nodes implies that the cluster is "responsible" for at most |C|/t neighboring clusters in the final clustering. This means that in the end there are at most n/t neighboring pairs of clusters, as desired.

A formal statement of this reduction together with its proof can be found in Appendix F. In fact, the proof strengthens the idea presented above and shows that it suffices that the clusters have a separation of 100t instead of  $100t \log n$ .

In Appendix F, we extend the result by showing how to compute spanners with O(n) edges from so-called weak-diameter clusterings, a more relaxed notion compared to strong-diameter clusterings, for which more efficient CONGEST algorithms exist [RG20, GGR21] (compared to strong-diameter clusterings [CG21]). We conclude Appendix F by showing that, using the reduction from ultra-sparse spanners to sparse spanners and the folklore reduction from weighted to unweighted spanners, we can efficiently compute weighted ultrasparse spanners.

### Acknowledgment

M.E. was supported by the ISF grant No. (2344/19). M.G., C.G., S. I., and V.R. were supported in part by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement No. 853109) and the Swiss National Foundation (project grant 200021-184735). B.H. was supported in part by NSF grants CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808, a Sloan Research Fellowship, funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement 949272), and the Swiss National Foundation (project grant 200021-184735).

# References

- [AB17] Amir Abboud and Greg Bodwin. The 4/3 additive spanner exponent is tight. Journal of the ACM (JACM), 64(4):1–20, 2017.
- [ABCP93] Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. Near-linear cost sequential and distribured constructions of sparse neighborhood covers. In 34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993, pages 638– 647, 1993.
- [ABS<sup>+</sup>20] Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. Computer Science Review, 37:100253, 2020.
- [ACIM99] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). SIAM J. Comput., 28(4):1167– 1181, 1999.
- [ADD<sup>+</sup>93] 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.
- [AP92] Baruch Awerbuch and David Peleg. Routing with polynomial communication-space trade-off. SIAM J. Discrete Math., 5(2):151–162, 1992.

- [BEG18] Leonid Barenboim, Michael Elkin, and Cyril Gavoille. A fast network-decomposition algorithm and its applications to constant-time distributed computation. *Theoretical Computer Science*, 751:2–23, 2018.
- [BGK<sup>+</sup>14] Guy E. Blelloch, Anupam Gupta, Ioannis Koutis, Gary L. Miller, Richard Peng, and Kanat Tangwongsan. Nearly-linear work parallel SDD solvers, low-diameter decomposition, and lowstretch subgraphs. *Theory Comput. Syst.*, 55(3):521–554, 2014.
- [BKMP10] Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. Additive spanners and  $(\alpha, \beta)$ -spanners. ACM Trans. Algorithms, 7(1):5:1–5:26, 2010.
- [BS07] Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007.
- [CG21] Yi-Jun Chang and Mohsen Ghaffari. Strong-diameter network decomposition. In *Proceedings* of the 2021 ACM Symposium on Principles of Distributed Computing, pages 273–281, 2021.
- [CHD20] Keren Censor-Hillel and Michal Dory. Fast distributed approximation for tap and 2-edgeconnectivity. *Distributed Computing*, 33(2):145–168, 2020.
- [Che13] Shiri Chechik. New additive spanners. In Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013, pages 498–512, 2013.
- [CHKPY18] Keren Censor-Hillel, Telikepalli Kavitha, Ami Paz, and Amir Yehudayoff. Distributed construction of purely additive spanners. *Distributed Computing*, 31(3):223–240, 2018.
- [Coh93] Edith Cohen. Fast algorithms for constructing t-spanners and paths with stretch t. In 34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993, pages 648–658, 1993.
- [DGPV08] Bilel Derbel, Cyril Gavoille, David Peleg, and Laurent Viennot. On the locality of distributed sparse spanner construction. In Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing, pages 273–282, 2008.
- [DGPV09] Bilel Derbel, Cyril Gavoille, David Peleg, and Laurent Viennot. Local computation of nearly additive spanners. In *International Symposium on Distributed Computing*, pages 176–190. Springer, 2009.
- [DHNS19] Mohit Daga, Monika Henzinger, Danupon Nanongkai, and Thatchaphol Saranurak. Distributed edge connectivity in sublinear time. In *Proceedings of the 51st Annual ACM SIGACT Sympo*sium on Theory of Computing, pages 343–354, 2019.
- [DHZ00] Dorit Dor, Shay Halperin, and Uri Zwick. All-pairs almost shortest paths. SIAM J. Comput., 29(5):1740–1759, 2000.
- [DKM19] Janosch Deurer, Fabian Kuhn, and Yannic Maus. Deterministic distributed dominating set approximation in the congest model. In *Proceedings of the 2019 ACM Symposium on Principles* of Distributed Computing, pages 94–103, 2019.
- [DMP<sup>+</sup>05] Devdatt P. Dubhashi, Alessandro Mei, Alessandro Panconesi, Jaikumar Radhakrishnan, and Aravind Srinivasan. Fast distributed algorithms for (weakly) connected dominating sets and linear-size skeletons. J. Comput. Syst. Sci., 71(4):467–479, 2005.
- [DMZ10] Bilel Derbel, Mohamed Mosbah, and Akka Zemmari. Sublinear fully distributed partition with applications. *Theory of Computing Systems*, 47(2):368–404, 2010.

[Dor18] Michal Dory. Distributed approximation of minimum k-edge-connected spanning subgraphs. In Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, pages 149-158, 2018. [DS08] Samuel I. Daitch and Daniel A. Spielman. Faster approximate lossy generalized flow via interior point algorithms. In Cynthia Dwork, editor, Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008, pages 451-460. ACM, 2008. [Elk05]Michael Elkin. Computing almost shortest paths. ACM Transactions on Algorithms (TALG), 1(2):283-323, 2005.[Elk07] Michael Elkin. A near-optimal distributed fully dynamic algorithm for maintaining sparse spanners. In Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing, pages 185–194, 2007. [EM19] Michael Elkin and Shaked Matar. Near-additive spanners in low polynomial deterministic congest time. In Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, pages 531–540, 2019. [EN18] Michael Elkin and Ofer Neiman. Efficient algorithms for constructing very sparse spanners and emulators. ACM Transactions on Algorithms (TALG), 15(1):1-29, 2018. [EP04] Michael Elkin and David Peleg.  $(1+\varepsilon, \beta)$ -spanner constructions for general graphs. SIAM Journal on Computing, 33(3):608-631, 2004. [Erd63] Paul Erdős. Extremal problems in graph theory. In Proceedings of the Symposium on Theory of Graphs and its Applications, page 2936, 1963. [GGR21] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhoň. Improved deterministic network decomposition. In Proc. of the 32nd ACM-SIAM Symp. on Discrete Algorithms (SODA), page 2904–2923, USA, 2021. Society for Industrial and Applied Mathematics. [GK18] Mohsen Ghaffari and Fabian Kuhn. Derandomizing distributed algorithms with small messages: Spanners and dominating set. In 32nd International Symposium on Distributed Computing (DISC 2018), 2018. [GP17] Ofer Grossman and Merav Parter. Improved deterministic distributed construction of spanners. In 31 International Symposium on Distributed Computing, 2017. [GPPR04] Cyril Gavoille, David Peleg, Stéphane Pérennes, and Ran Raz. Distance labeling in graphs. J. Algorithms, 53(1):85–112, 2004. [GY20]Parikshit Gopalan and Amir Yehudayoff. Concentration for limited independence via inequalities for the elementary symmetric polynomials. Theory of Computing, 16(1):1–29, 2020. [Kar99] David R Karger. Random sampling in cut, flow, and network design problems. Mathematics of Operations Research, 24(2):383–413, 1999. [KP98] Shay Kutten and David Peleg. Fast distributed construction of smallk-dominating sets and applications. Journal of Algorithms, 28(1):40-66, 1998. [Li20] Jason Li. Faster parallel algorithm for approximate shortest path. In Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, page 308–321, New York, NY, USA, 2020. Association for Computing Machinery. [Lin87] Nathan Linial. Distributive graph algorithms-global solutions from local data. In 28th Annual Symposium on Foundations of Computer Science, Los Angeles, California, USA, 27-29 October 1987, pages 331-335, 1987.

- [LL18] Christoph Lenzen and Reut Levi. A centralized local algorithm for the sparse spanning graph problem. In 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, pages 87:1–87:14, 2018.
- [MPVX15] Gary L. Miller, Richard Peng, Adrian Vladu, and Shen Chen Xu. Improved parallel algorithms for spanners and hopsets, 2015.
- [Par19] Merav Parter. Small cuts and connectivity certificates: A fault tolerant approach. In 33rd International Symposium on Distributed Computing, 2019.
- [Pel00] David Peleg. Distributed computing: a locality-sensitive approach. SIAM, 2000.
- [Pen16] Richard Peng. Approximate undirected maximum flows in <math>o(mpolylog(n)) time. In *Proceedings* of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms, pages 1862–1867, 2016.
- [Pet09] Seth Pettie. Low distortion spanners. ACM Trans. Algorithms, 6(1):7:1–7:22, 2009.
- [Pet10] Seth Pettie. Distributed algorithms for ultrasparse spanners and linear size skeletons. *Distributed Computing*, 22(3):147–166, 2010.
- [PS89] David Peleg and Alejandro A Schäffer. Graph spanners. *Journal of graph theory*, 13(1):99–116, 1989.
- [PY18a] Merav Parter and Eylon Yogev. Congested clique algorithms for graph spanners. In 32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018, pages 40:1–40:18, 2018.
- [PY18b] Merav Parter and Eylon Yogev. Congested clique algorithms for graph spanners. In 32nd International Symposium on Distributed Computing, page 3, 2018.
- [RG20] Václav Rozhoň and Mohsen Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *Proceedings of the 52nd Annual ACM SIGACT* Symposium on Theory of Computing, pages 350–363, 2020.
- [RV11] L Shankar Ram and Elias Vicari. Distributed small connected spanning subgraph: Breaking the diameter bound. *Technical report/Swiss Federal Institute of Technology Zurich, Department of Computer Science*, 530, 2011.
- [RZ04] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. In Algorithms ESA 2004, 12th Annual European Symposium, Bergen, Norway, September 14-17, 2004, Proceedings, pages 580–591, 2004.
- [SS11] Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. SIAM J. Comput., 40(6):1913–1926, 2011.
- [Thu97] Ramakrishna Thurimella. Sub-linear distributed algorithms for sparse certificates and biconnected components. *Journal of Algorithms*, 23(1):160–179, 1997.
- [TZ01] Mikkel Thorup and Uri Zwick. Compact routing schemes. In Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA 2001, Heraklion, Crete Island, Greece, July 4-6, 2001, pages 1–10, 2001.
- [TZ05] Mikkel Thorup and Uri Zwick. Approximate distance oracles. Journal of the ACM (JACM), 52(1):1–24, 2005.
- [TZ06] Mikkel Thorup and Uri Zwick. Spanners and emulators with sublinear distance errors. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 802–809, 2006.

# A Some Other Related Work

**Centralized and parallel algorithms for ultra-sparse weighted spanners** We next discuss algorithms from the centralized and parallel computation settings that compute ultra-sparse weighted spanners, and elaborate why they are not applicable in the distributed setting.

One (centralized) approach to construct ultra-sparse spanners with good stretch for weighted graphs is to give an efficient implementation for the greedy algorithm of  $[\text{ADD}^+93]$ . This was done by Roditty and Zwick [RZ04] who showed that the greedy algorithm can be implemented in  $\tilde{O}(n^2)$ time in the centralized model of computation. Unfortunately, this implementation is inherently sequential and not amenable to extend to work-efficient distributed or parallel implementations. Dubhashi et al. [DMP<sup>+</sup>05] devised a distributed implementation of the greedy algorithm in the LOCAL model of distributed computation, but this algorithm only works for unweighted graphs and is not work-efficient <sup>2</sup>.

In terms of work-efficient algorithms, to the best of our knowledge, there are currently only two parallel algorithms for building ultra-sparse sub-graphs with distance-approximation guarantees for weighted graphs, namely Blelloch et al. [BGK<sup>+</sup>14] and Li [Li20]. Both algorithms seem inherently randomized and not suitable for distributed computational models:

The algorithm of Blelloch et al. [BGK<sup>+</sup>14] provides ultra-sparse subgraphs with a guarantee on the *average* stretch of edges of the original graph, as opposed to the worst-case stretch guarantee of spanners. While average-stretch is good enough for some spectral-sparsification algorithms, it is insufficient for many other applications, e.g., for Li's recursive scheme [Li20] for computing  $(1 + \varepsilon)$ -transshipment and  $(1 + \varepsilon)$ -approximate single-source shortest paths.

The second randomized parallel algorithm, due to Li [Li20], provides for any  $t \ge 1$  a spanners with n + n/t edges and  $O(t^2 \cdot \log^2 n \cdot \log^2 \log n)$  stretch using  $O(t \cdot \log^3 n \cdot \log \log n \log^* n)$  depth and  $O(m \log n)$  work, while succeeding with constant probability.<sup>3</sup> Because the approach of [MPVX15] for weighted graphs, which is also utilized by Li [Li20], requires  $\Omega(\frac{\log n}{\log \log n})$  rounds of contractions, these two algorithms cannot be implemented distributedly.

### **B** Short Seed Length

We want to approximate a hitting-event E over n binary random variables  $X_1, \ldots, X_n$  where each  $X_i$  is sampled independently with probability p. For this, we can use the family of hash functions provided in the work of Gopalan and Yehudayoff [GY20].

**Theorem B.1** ([GY20], Theorem 1.9). For positive integers N and M, and for error parameter  $\delta > 0$ , there is a family of hash functions  $\mathcal{H}$  from [N] to [M] with size  $2^r$  where  $r = O((\log \log N + \log(M/\delta)) \log \log(M/\delta))$ , such that for every subsets  $T_1, \ldots, T_N \subseteq [M]$ , we have:

$$|\operatorname{Pr}_{h\sim\mathcal{H}}[\forall i\in[N],h(i)\in T_i]-\operatorname{Pr}_{u\sim\mathcal{U}}[\forall i\in[N],u(i)\in T_i]|\leq\delta$$

where  $\mathcal{U}$  is the family of all functions from [N] to [M]. We can sample from  $\mathcal{H}$  in polynomial time.

Set N to n, M to 1/p, and  $\delta$  to  $n^{-10}$ . So the seed length is  $O(\log n \log \log n)$  as in our applications, we always have  $p \ge 1/\operatorname{poly}(n)$ . To sample  $\{X_i\}_{i\in[n]}$ , we sample a random hash function from  $\mathcal{H}$ ,

<sup>&</sup>lt;sup>2</sup>The LOCAL model allows sending of arbitrarily large messages and performing arbitrarily heavy local computations. The algorithm of  $[DMP^+05]$  indeed employs such heavy local computations.

<sup>&</sup>lt;sup>3</sup>When run on an *n*-vertex *m*-edge graph with a parameter *k*, the algorithm of [Li20], with constant probability, produces an  $O(k^2)$ -spanner with  $n + O(\frac{m \cdot \log n}{k})$  edges in  $O(k \log^2 n \log^* n)$  time and  $O(m \log^2 n)$  work. By first building an  $O(\log n)$ -spanner with  $m = O(n \log \log n)$  edges via the algorithm of [MPVX15], and then running the algorithm of [Li20] on top of this spanner with  $k = t \log n \log \log n$ , we obtain the bounds cited here.

and then set  $X_i$  to 1 if h(i) is 1 and set it to zero if  $h(i) \neq 1$ . This approximates any hitting event E with error  $\delta$ . To show this, suppose  $S \subseteq [n]$  is the corresponding subset of E, that is

$$\Pr[E=0] = \Pr[\bigwedge_{i \in S} X_i = 0] = (1-p)^{|S|}$$

For  $i \in [N]$ , we set  $T_i$  to  $[M] \setminus \{1\}$  if  $i \in S$  and we set  $T_i$  to [M] if  $i \notin S$ . Note that:

$$\Pr_{u \sim \mathcal{U}}[\forall i \in [N], u(i) \in T_i] = \Pr_{u \sim \mathcal{U}}[\forall i \in S, u(i) \in T_i]$$
$$= \Pr_{u \sim \mathcal{U}}[\forall i \in S, u(i) \neq 1]$$
$$= (1 - 1/M)^{|S|} = (1 - p)^{|S|}$$

On the other hand, we have:

$$\begin{aligned} \Pr_{h \sim \mathcal{H}}[\forall i \in [N], h(i) \in T_i] &= \Pr_{h \sim \mathcal{H}}[\forall i \in S, h(i) \in T_i] \\ &= \Pr_{h \sim \mathcal{H}}[\forall i \in S, X_i \neq 1] \\ &= \Pr_{h \sim \mathcal{H}}[\bigwedge_{i \in S} X_i = 0] \end{aligned}$$

So with this setting, we can capture hitting-event E. And so the distribution that is provided by  $\mathcal{H}$  gives a  $\delta = n^{-10}$  approximation of this event with only  $O(\log n \log \log n)$  random bits.

*Remark.* Note that p might not be in the form of 1/M. In that case, assuming  $p = \Omega(1/\text{poly}(n))$ , we can write p = M'/M - 1/poly(n) with M and M' being poly(n) bounded integers. If in the above, we set  $T_i$  to  $[M] \setminus [M']$  for  $i \in S$  and  $T_i$  to [M] for  $i \notin S$ , we again get approximation with polynomially small error with  $O(\log n \log \log n)$  seed length.

#### C Baswana-Sen Derandomization

We deterministically assign sample/unsample to clusters using the method of conditional expectation running on network decomposition. But first, let us recall the definition of network decomposition.

A (Q, D) weak-diameter network decomposition of an unweighted graph G is a partitioning of nodes of G into clusters, colored with Q colors, such that nodes of each cluster are at most at distance D in G. Moreover, it guarantees that there is no edge between clusters with the same color. According to the work of Rozhoň and Ghaffari [RG20], one can find a weak-diameter network decomposition in polylog(n) rounds deterministically. Indeed, they show that we can find a network decomposition of any constant power of G in polylog(n) rounds. The *t*-th power of G = (V, E), denoted by  $G^t$ , is a graph on V with an edge between any two nodes with distance at most t in G. For our derandomization, we need network decomposition of  $G^2$ . Concretely, it means that any two clusters C and C' with the same color have a distance of at least three in G.

**Theorem C.1** ([RG20], Theorem 2.12). There is a deterministic distributed algorithm that computes a  $(\log n, \operatorname{polylog}(n))$  weak-diameter network decomposition of  $G^2$ .

Moreover, for each color class and each cluster C on this class, a Steiner tree  $T_{\mathcal{C}}$  with radius polylog(n) is associated. Terminal nodes of  $T_C$  are the set of nodes of C and each edge of G appears in O(polylog(n)) of these Steiner trees.

The second part of the theorem above guarantees that for any color q, we can simulate convergecast and broadcast on the collections of clusters with color q simultaneously as each cluster can communicate through its Steiner tree. Since each edge is in only few Steiner trees, the congestion overhead of simultaneous communication cannot be more than polylog(n) rounds.

Proof of Lemma 3.3. Let C be the input clustering of iteration i (an (i-1)-partition of nodes that are alive at the beginning of iteration i) and H be the cluster-graph induced by C. First, we find a weak-diameter network decomposition of  $H^2$  into  $O(\log n)$  color classes such that each cluster has weak diameter polylog(n). For that, We can run the polylog(n) rounds algorithm of [RG20] on H. The algorithm can be easily adapted to run on cluster-graphs as its operation only needs broadcast and convergecast, and moreover, it can run on a network where each node has polylog(n) bits of memory. Since each node of H corresponds to a tree with radius i - 1 of G, the total running time is  $O(i \cdot \text{polylog}(n))$  rounds. To avoid confusion, we reserve the term "cluster" for referring to the clusters in Baswana-Sen algorithm and we refer to each cluster of network decomposition by "ND-cluster".

The contribution of each node to both  $U_i^w$  and  $U_i^{uw}$  (see (3.1) and (3.2)) can be written as the contribution of each individual alive node of G. The contribution of  $X_j^{(i)}$  (indicator random variable that determines whether *j*-th cluster in iteration *i* is sampled or not) is considered as part of a contribution of the root of the *j*-cluster. Moreover, observe that each node of G can compute its contribution by looking at its 1-hop neighborhood. In the following, we show how we can derandomize a utility function with these properties.

We go over the color classes of network decomposition one by one. Suppose we assign sample/unsample labels for each Baswana-Sen cluster that is in a ND-cluster with color 1 to q - 1and now we are working on color q. Suppose all ND-clusters with color q, independent of each other, sample the Baswana-Sen clusters that are inside them with  $O(\log n \log \log n)$  bits from the distribution in Appendix B<sup>4</sup>. In order to derandomize sampling of each ND-cluster, we go over  $O(\log n \log \log n)$  bits of its random seed one by one and fix them in such a way that the expected value of the utility function condition on those fixed bits is at most be the expected value of utility function without fixing those bits.

Consider one ND-cluster  $\mathcal{D}$  with color q. Let  $V_{\mathcal{D}}$  be the set of alive nodes of G in  $\mathcal{D}$  and let  $N(\mathcal{D})$  be the union of  $V_{\mathcal{D}}$  and all neighbors of  $V_{\mathcal{D}}$  in G. Observe that the randomness of  $\mathcal{D}$  only affects the contribution of nodes in  $N(\mathcal{D})$ . Since we do a network decomposition for the second power of H, we have two properties: (a) all nodes of one Baswana-Sen cluster is in one ND-cluster, and (b) for any two ND-clusters  $\mathcal{D}$  and  $\mathcal{D}'$ , the set of affected nodes by the randomness of  $\mathcal{D}$  and  $\mathcal{D}'$  are disjoint, i.e.  $N(\mathcal{D}) \cap N(\mathcal{D}') = \emptyset$ . Here is how we fix the first random bit for  $\mathcal{D}$ . Other bits can be fixed similarly. The first bit is fixed to  $b \in \{0, 1\}$  if

 $\mathbb{E}[\text{contribution of } N(\mathcal{D}) \mid \text{first bit is } b] \leq \mathbb{E}[\text{contribution of } N(\mathcal{D}) \mid \text{first bit is } 1-b]$ 

To decide for which b the conditional expectation is smaller, each node in  $N(\mathcal{D})$  computes two values: its contribution to the utility function if the first random bit equals zero/one. Each of those two values can be computed by taking a sum over all assignments of bits that are not yet fixed. So it may take  $2^{O(\log n \log \log n)} = n^{O(\log \log n)}$  time. Then for each Baswana-Sen cluster C in  $\mathcal{D}$ , all nodes in C aggregates the sum of their contributions in the root of the cluster. This takes O(i)rounds<sup>5</sup>. Each cluster C in  $\mathcal{D}$  is indeed represents a node in H. Now, these nodes in H aggregates their computed values into the leader of the ND-cluster  $\mathcal{D}$ . This can be done in  $i \cdot \text{polylog}(n)$  rounds via the Steiner trees of  $\mathcal{D}$  (as it is mentioned in Theorem C.1). The leader then decides which bit should be fixed and broadcast it to all nodes in  $N(\mathcal{D})$ . Since fixing one bit takes  $i \cdot \text{polylog}(n)$  rounds

 $<sup>^{4}</sup>$ Note that if two distributions approximate the independent distribution, their product is also an approximation.

<sup>&</sup>lt;sup>5</sup>We trim the computed values into  $O(\log n)$  bits. This is safe as it only incurs 1/poly(n) error.

and as there are  $O(\log n \log \log n)$  bits, fixing all of them also takes  $i \cdot \operatorname{polylog}(n)$  rounds. There are  $O(\log n)$  color classes in the network decomposition, so simulating the sampling for iteration i of Baswana-Sen takes  $i \cdot \operatorname{polylog}(n)$  rounds. There are g iterations, so the final round complexity is  $\sum_{i=1}^{g} i \cdot \operatorname{polylog}(n) = g^2 \cdot \operatorname{polylog}(n)$ .

## D Linear Size Spanners

Proof of Theorem 1.5. The algorithm consists of  $P = O(\log^* n)$  phases. Phase *i* is given an input graph  $G_i$  with  $n_i$  nodes. Using the algorithm of Lemma 3.3, we run the Baswana-Sen algorithm for  $g_i = (1 + I_w)x_i(1 + 2\frac{\log \log x_i}{\log x_i})$  iterations and with sampling probability  $\frac{1}{x_i}$ . Here,  $I_w$  is 1 if we are in the weighted case and 0 if we are in the unweighted case. This is why there is an extra  $2^{\log^* n}$  factor in the stretch bound of the spanner for the weighted case. Parameter  $x_i$  is determined later. At the end of phase *i*, we define graph  $G_{i+1}$  (input of the next phase) as the  $g_i$ -cluster-graph that is induced by  $g_i$ -clustering of running  $g_i$  iterations of Baswana-Sen. The input to the first phase is the network graph G. We show that the there is a sequence  $x_1, \ldots, x_P$  such that the edges that are added to the spanner in this P phases gives us a spanner with the claimed size and stretch.

The goal of phase *i* is to construct a spanner for  $G_i$  with stretch  $s_i = \prod_{j=i}^{P} (2g_j + 1)$ . Note that since  $G_1$  is G, proving this implies that our output is a  $s_1$ -spanner of G. In phase *i*, after running the sampling for  $g_i$  iterations, all dead edges of  $G_i$  has stretch  $2g_i - 1 \leq s_i$  according to Lemma 3.1. So we do not need to worry about them in terms of stretch and we can safely remove them from further consideration. Again, according to Lemma 3.1, observe that the  $g_i$ -clustering that we have after Baswana-Sen iterations is stretch-friendly. From Observation 3.5, if we can find an  $s_{i+1}$ -spanner for  $G_{i+1}$ , we get a spanner of stretch  $s_i$  for  $G_i$ . So if we show that the last phase gives a spanner with stretch  $s_P$ , the stretch of other phases follows automatically by induction. For that, we design the sequence  $x_1, \ldots, x_P$  such that all nodes of  $G_P$  are dead at the end of the last phase and as a result, all edges of  $G_p$  are dead which means that their stretch in the spanner is bounded by  $2g_P - 1 < s_P$ .

Let P be the largest integer such that  $\log^{(P)}(n) \ge \alpha_0$ . Here,  $\log^{(P)} n$  represents the P-th iteration of log and  $\alpha_0$  is the constant in Lemma D.1. Observe that  $P \le \log^* n$  and that  $\log^{(P)} n$  is a constant. We define the sequence of  $x_i$  as follows:

$$x_1 = \frac{\log^{(P)} n}{\log^{(P+1)} n}, x_2 = \frac{\log^{(P-1)} n}{\log^{(P)} n}, \dots, x_P = \frac{\log n}{\log \log n}$$

Note that:

$$s_{1} = \prod_{i=1}^{P} (2g_{i} + 1) \leq 2^{P} \left( \prod_{i=1}^{P} g_{i} \right) e^{\sum_{i=1}^{P} \frac{1}{2g_{i}}}$$
$$\leq (2(1 + I_{w}))^{P} \left( \prod_{i=1}^{P} x_{i} \right) e^{\sum_{i=1}^{P} \frac{2\log\log x_{i}}{\log x_{i}} + \frac{1}{2g_{i}}}$$
$$= (2(1 + I_{w}))^{P} \cdot (\log n) \cdot O(1).$$

We use the inequality  $1 + x \leq e^x$  and the fact that the summations  $\sum_{i=1}^{P} \frac{1}{g_i}$  and  $\sum_{i=1}^{P} \frac{2\log\log x_i}{\log x_i}$  are constants. This is because the sequence of  $x_i$  (and so  $g_i$ ) are exponentially growing and so the summations are dominated by their first term. So the output spanner has stretch  $O(\log n \cdot 2^{\log^* n})$  for unweighted graphs and  $O(\log n \cdot 4^{\log^* n})$  for weighted graphs. There is one remaining ingredient

though. We have to show that in last phase, all nodes of  $G_P$  die. Note that at the end of this phase, we have at most  $n_P/x_P^{g_P}$  clusters according to Lemma 3.3. On the other hand, from Lemma D.1, we have  $x_P^{g_P} \ge n \ge n_P$  (set  $\alpha$  to n in the lemma). So there is no cluster at the end of last phase and so all nodes are dead. This concludes the proof for the claimed stretch.

To bound the spanner size, let us first consider the unweighted case. Note that according to Lemma 3.3, we add  $O(n_ig_i + n_ix_i \log g_i) = O(n_ix_i \log x_i)$  edges to the spanner in phase *i*. From Lemma D.1,  $x_i \log x_i \leq x_{i-1}^{g_{i-1}}$  (set  $\alpha$  to  $\log^{(P-i+1)} n$  in the lemma). On the other hand, recall that  $n_i \leq n_{i-1}/x_{i-1}^{g_{i-1}}$ . So for any i > 1, the contribution of phase *i* is  $O(n_{i-1})$ . In the first phase,  $x_1$  is a constant so  $n_1x_1 \log g_1 = O(n_1)$ . So size of the spanner is bounded by  $O(n_1 + \sum_{i=2}^{P} n_{i-1})$ . Again, using  $n_i \leq n_{i-1}/x_{i-1}^{g_{i-1}}$ , we can write  $\sum_{i=2}^{t} n_{i-1} = O(n_1)$  as the summation is dominated by its first term. So in total, the size of the spanner is  $O(n_1) = O(n)$ . For the weighted case, note that in phase *i* we add  $O(n_ig_ix_i) = O(n_ix_i^2)$  edges. So  $x_i \log x_i$  in the unweighted case, is replaced by  $x_i^2$ . On the other hand, we double the number of iterations of each phase. So similar to the unweighted case, we can show that phase *i* in the weighted case is also adds  $O(n_{i-1})$  edges and so the final spanner size is O(n).

To implement the algorithm, the main challenge is that how we can run Lemma 3.3 on a clustergraph. For each node v in  $G_i$ , there is a tree between nodes of  $\operatorname{inv}^{G_i \to G_{i-1}}(v)$  with radius at most  $g_{i-1}$ . Subsequently, there is a tree among  $\operatorname{inv}^{G_i \to G}(v)$  with radius at most  $\frac{s_1}{s_i} = \prod_{j=1}^{i-1} (2g_j + 1)$ . This rooted tree can be constructed in  $O(\frac{s_1}{s_i}) = O(\log n)$  rounds. The root of this tree is responsible for the node v when we run phase i. Recall that in phase i, we run Baswana-Sen for  $g_i$  iterations with sampling probability  $\frac{1}{x_i}$ . To run that, each node of  $G_i$  should know its adjacent clusters. If a node has d adjacent clusters, then the root can discover this d neighbors in  $O(d \cdot \frac{s_1}{s_i})$  rounds. This can be problematic if d is large. However, note that according to Lemma 3.3 all nodes with  $\Omega(x_i \log n)$  adjacent clusters remain alive in each iteration. So a node only adds at most  $\min(\Theta(x_i \log n), d) = O(\log^2 n)$  edges in each iteration. So it should only know the first  $O(\log^2 n)$ clusters (when we order the d adjacent clusters of node v according to the smallest weight of an edge between them and v). So v can find these clusters in polylog(n) rounds. The other ingredient of derandomization in Lemma 3.3 is running the algorithm of Theorem C.1 for network decomposition on cluster-graphs. But recall that we already discussed do this in the proof of Lemma 3.3. This concludes that the round complexity of each phase and as a result the whole algorithm (since there are only  $O(\log^* n)$  phases) is polylog(n). 

**Lemma D.1.** There is a sufficiently large constant  $\alpha_0$ , such that for any  $\alpha > \alpha_0$ , we have

$$x \log x \le \alpha \le y^z \tag{D.1}$$

where  $x = \frac{\alpha}{\log \alpha}$ ,  $y = \frac{\log \alpha}{\log \log \alpha}$ , and  $z = y(1 + \frac{2 \log \log y}{\log y})$ .

*Proof.* For the ease of notation, we first take a log from (D.1):

$$\log x + \log \log x \le \log \alpha \le z \log y \tag{D.2}$$

Since  $\log x = \log \alpha - \log \log \alpha$  and  $\log \log x \le \log \log \alpha$ , we have  $x \log x \le \log \alpha$ . Next, we show that  $\log \alpha \le z \log y$ . Note that:

$$z\log y = y\log y + 2y\log\log y \tag{D.3}$$

Replacing  $\log y$  with  $\frac{\log \alpha}{\log \log \alpha}$  in the first term, we get

$$y \log y = \log \alpha - \log \alpha \cdot \frac{\log \log \log \alpha}{\log \log \alpha} = \log \alpha - y \log \log \log \alpha$$
(D.4)

If we show that  $2 \log \log y \ge \log \log \log \alpha$ , then the second summand in the RHS of (D.3) is larger than the negative term of (D.4) and we are done. For this, note that

$$\log \log y = \log(\log \log \alpha - \log \log \log \alpha)$$

Since  $\alpha$  is sufficiently large, we have  $\log \log \log \alpha \leq \frac{\log \log \alpha}{2}$ . So  $\log \log y \geq \log \log \log \alpha - 1$ . This implies:

$$2\log\log y \ge 2(\log\log\log\alpha - 1) \ge \log\log\log\alpha$$

and concludes the proof.

### E Stretch-Friendly Clustering

Proof of Lemma 4.1. Consider a cluster  $C' \in C'$ . We have to prove three properties for C': (a) Size of C' is at least  $2^i$ , (b) Radius of C' is less than  $3 \cdot 2^i$ , and (c) C' is stretch-friendly. Let us first show that C' satisfies (a) and (b). For this, we consider two cases based on how C' is created:

- C' is consisting of a big cluster C with a set of small clusters that has an outgoing edge to C. The size of C' is at least  $2^i$  as it contains C. Each small cluster has size at most  $2^i - 1$  (so their hop-diameter is bounded by  $2^i - 2$ ). According to the definition of merging two clusters, the root of C' is the same as the root of C. Since the radius of C is at most  $3 \cdot 2^{i-1} - 1$ , the final radius of C' is at most  $(2^i - 2 + 1) + (3 \cdot 2^{i-1} - 1) < 3 \cdot 2^i$ .
- C' is consisting of two matched small clusters  $C_1$  and  $C_2$  with a set of small clusters that has an outgoing edge to  $C_1$  and a set of small clusters that has an outgoing edge to  $C_2$ . Without loss of generality, suppose the matched edge between  $C_1$  and  $C_2$  is oriented from  $C_1$  to  $C_2$ . So the root of C' is the root of  $C_2$ . There are at least two small clusters in C', so its size is at least  $2 \cdot 2^{i-1} = 2^i$ . As we have discussed before, each small cluster has hop-diameter  $2^i - 2$ . So the radius of C' is at most  $2 \cdot (2^i - 2 + 1) + (2^i - 2) < 3 \cdot 2^i$ .

It only remains to show the property (c), i.e. C' is stretch-friendly. Note that each cluster in C' is constructed from merging several clusters in C. And we know by induction that C is a stretch-friendly partition. So if we show that two stretch-friendly clusters are merged to a stretch-friendly cluster, we are done.

Consider two stretch-friendly clusters  $C_1$  and  $C_2$  are merged to a cluster  $C_{12}$ . Suppose the merge edge  $\{u_1, u_2\}$  where  $u_1 \in C_1$  and  $u_2 \in C_2$  is oriented from  $C_1$  to  $C_2$  and has weight w. Moreover, let root<sub>1</sub> and root<sub>2</sub> be the roots of  $C_1$  and  $C_2$ , respectively. From the definition of merge, the root of  $C_{12}$  is root<sub>2</sub>. Consider a boundary edge with weight w' of  $C_{12}$  that is incident to a node  $v \in C_{12}$ . If  $v \in C_1$ , then  $w' \geq w$ . Each edge in the path from v to root<sub>1</sub> has weight at most w' as  $C_1$  is stretch-friendly. Each edge in the path from root<sub>1</sub> to  $u_1$  has weight at most w (again, because  $C_1$ is stretch-friendly) and each edge in the path from  $u_2$  to root<sub>2</sub> has weight at most w (since  $C_2$  is stretch-friendly). So since  $w' \geq w$ , all the edges in the path from v to the root of  $C_{12}$  which is root<sub>2</sub> have weight at most w'. The other is that  $v \in C_2$ . This case is trivial as root<sub>2</sub> is root of both  $C_2$ and  $C_{12}$ . So for all boundary edges, C' is stretch-friendly. It remains to show the stretch-friendly property for all inside-edges of  $C_{12}$ . For inside edges of  $C_1$  and  $C_2$ , we already know that the property holds. So suppose an inside-edge with weight w' between  $v_1 \in C_1$  and  $v_2 \in C_2$ . Exact similar argument, shows that all edges in the path from  $v_1$  to root<sub>2</sub> has weight at most w' and all edges in the path from  $v_2$  to root<sub>2</sub> has weight at most w' as well. So all the edges in the path from

 $v_1$  to  $v_2$  has weight at most w'. This completes the proof that  $C_{12}$  is stretch-friendly and as a result shows that  $\mathcal{C}'$  (the output of iteration i) is stretch-friendly.

Since there are  $\lceil \log t \rceil$  iterations, the output of the last iteration is a stretch-friendly O(t)-partition where each cluster has size at least t. So there are at most n/t clusters.

The only remaining piece of proving Lemma 4.1 is proving the claimed round complexity. For that, notice that steps (1), (2), (4), and (5) of the algorithm can be easily run in  $O(2^i)$  rounds since the hop-diameter of each cluster is  $O(2^i)$ . Step (3) needs  $O(2^i \log^* n)$  rounds as it is well-known that a graph with out-degree one can be 3-colored in  $O(\log^* n)$  rounds [Lin87]. And in our case, each node of a graph is a cluster with hop-diameter  $O(2^i)$ . So the total number of rounds of all the  $\lceil \log t \rceil$  iterations is  $O(\sum_{i=1}^{\log t} 2^i \log^* n) = O(t \log^* n)$ .

# F Spanners via Low-Diameter Clusterings

In Appendix F.1, we show that an efficient algorithm for computing a well-separated strongdiameter clustering implies an efficient algorithm for unweighted ultrasparse spanners. In Appendix F.2, we show that an efficient algorithm for computing a 3-separated weak-diameter clustering implies an efficient algorithm for unweighted sparse spanners. Finally, in Appendix F.3, we combine the result from the previous section together with (1) the weak-diameter clustering algorithm of [RG20], (2) the folklore reduction from weighted to unweighted spanners, and (3) the reduction from ultra-sparse spanners to sparse spanners to show that we can efficiently compute a weighted ultra-sparse spanner.

#### F.1 Unweighted Ultra-Sparse Spanners via Strong-Diameter Clusterings

In this section, we show that unweighted ultra-sparse spanners can be efficiently computed given an efficient algorithm for computing well-separated strong-diameter clusterings.

**Theorem F.1.** Let t > 0. Assume that for any t' one can construct a t'-separated strongdiameter clustering with diameter D(n,t') in T(n,t') rounds of the deterministic (randomized) CONGEST model.

Then, there is a deterministic (randomized) distributed algorithm that, given an unweighted graph G, constructs an  $\alpha$ -spanner H such that

- 1.  $|E(H)| \le n + n/t$ ,
- 2.  $\alpha = O(D(n, 10t) + t)$

in  $O(\log n) \cdot (T(n, 10t) + D(n, 10t) + t)$  CONGEST rounds.

The proof follows along the lines of the proof sketch given in Section 5, with two differences.

- 1. We start with  $\Theta(t)$ -separated clusters instead of  $\Theta(t \log n)$ -separated ones. This is because it is enough for our purposes that a large fraction of clusters needs to stop growing after O(t)steps; we do not need to wait for  $O(t \log n)$  steps after all of them finish.
- 2. Choosing an arbitrary edge between any two neighboring clusters of a given clustering is a non-trivial task. Therefore, we gradually compute a set of edges  $E^{inter}$  such that the size of  $E^{inter}$  is sufficiently small, and between any two neighboring clusters there exists at least one edge in  $E^{inter}$  connecting these two clusters.



**Figure 1:** The figure shows two clusters  $C_1, C_2$  that are 8-separated. We can enlarge  $C_1$  to  $C_1^{+1}$  and  $C_2$  to  $C_2^{+3}$  so that the new clusters  $C_1^{+1}$  and  $C_2^{+3}$  have the property that the number of nodes neighboring with them (blue for  $C_1^{+1}$  and red for  $C_2^{+3}$ ) is only a small fraction of their size. The clusters remain 4-separated, and their diameter is only increased additively by at most +2 in case of  $C_1$  and at most +6 in case of  $C_2$ .

In particular, we obtain the following lemma.

**Lemma F.2.** Assume that for any t' one can construct a t'-separated low diameter clustering of an n-vertex graph with diameter D(n, t') in T(n, t') rounds of the deterministic (randomized) CONGEST model.

Then, for any t > 0 there is a deterministic (randomized) distributed algorithm that outputs a complete clustering C together with a set of edges  $E^{inter} \subseteq E(G)$  such that

- 1.  $diam(C) \le D(n, 10t) + 10t$ ,
- 2.  $|E^{inter}| \leq n/t$ ,
- 3. for any two neighboring clusters  $C_1, C_2 \in C$ , there exists at least one edge in  $E^{inter}$  with one endpoint in  $C_1$  and one endpoint in  $C_2$ .

The round complexity of the algorithm in the CONGEST model is

$$O(\log n) \cdot (T(n, 10t) + D(n, 10t) + t).$$

*Proof.* Let  $T = \lceil 1 + \log_{10/7} n \rceil$ . The algorithm computes a sequence of clusterings of  $G \ C_0 \subseteq C_1 \subseteq \cdots \subseteq C_T$  and a sequence of edges  $E_0^{inter} \subseteq E_1^{inter} \subseteq \cdots \subseteq E_T^{inter} \subseteq E$ . We denote by  $V_i^{unclustered}$  the set of unclustered nodes in  $C_i$ . The clustering  $C_i$  and the set of edges  $E_i^{inter}$  will satisfy

- 1. diam $(\mathcal{C}_i) \leq D(n, 10t) + 10t$ ,
- 2.  $|V_i^{unclustered}| \leq (7/10)^i n$ ,
- 3.  $|E_i^{inter}| \leq |\bigcup_{C \in \mathcal{C}_i} C|/t,$

- 4. for any two neighboring clusters  $C_1, C_2 \in C_i$ , there exists at least one edge in  $E_i^{inter}$  with one endpoint in  $C_1$  and one endpoint in  $C_2$ ,
- 5. for any node  $v \in V_i^{unclustered}$  and neighboring cluster  $C \in \mathcal{C}_i$ , there exists a node  $u \in C$  with  $\{v, u\} \in E_i^{inter}$ .

Note that the invariants for i = 0 are satisfied by letting  $C_0$  be the empty clustering and  $E_0^{inter}$  be the empty set.

Moreover,  $|V_T^{unclustered}| \leq (7/10)^T n < 1$  and therefore  $\mathcal{C}_T$  is a complete clustering of G. Hence, we can output  $\mathcal{C} = \mathcal{C}_T$  and  $E^{inter} = E_T^{inter}$ .

Thus, it remains to show how to compute  $C_{i+1}$  and  $E_{i+1}^{inter}$  given  $C_i$  and  $E_i^{inter}$  while preserving the invariants.

Let  $G_i = G[V_i^{unclustered}]$ . First, we compute a 10t-separated strong-diameter clustering  $C'_i$  of  $G_i$  with diameter D(n, 10t) that clusters at least  $|V(G_i)|/2$  nodes in T(n, 10t) rounds of the deterministic (randomized) CONGEST model.

To explain the next part of the phase, let us introduce a new notation (see also Figure 1). For a cluster  $C \in \mathcal{C}'_i$ , we denote by  $C^{+j}$  its superset such that  $u \in C^{+j}$  if and only if  $d_{G_i}(u, C) \leq j$ .

For  $j \in \{0, 1, \dots, 4t - 1\}$ , we say that distance j is a good cutting distance for a cluster  $C \in \mathcal{C}'_i$  if  $C^{+j}$  has at most  $\frac{|C|}{t}$  neighboring nodes in  $G_i$ .

We say that a cluster  $C \in \mathcal{C}'_i$  is good if there exists a good cutting distance for C, and otherwise, we refer to C as bad. For a good cluster  $C \in \mathcal{C}'_i$ , let  $j_C$  denote the smallest good cutting distance for C. Then, we obtain the clustering  $\mathcal{C}_{i+1}$  from  $\mathcal{C}_i$  by adding for each good cluster  $C \in \mathcal{C}'_i$  the cluster  $C^{+j_C}$  to the clustering  $\mathcal{C}_{i+1}$ .

Moreover, we obtain  $E_{i+1}^{inter}$  from  $E_i^{inter}$  by adding for each node v in  $G_i$  and newly added neighboring cluster  $C^{+j_c}$  an arbitrary edge connecting v with  $C^{+j_c}$ .

We next show that all the invariants are preserved, and afterward discuss the distributed implementation of the procedure.

We start with the first invariant, namely that  $\operatorname{diam}(\mathcal{C}_{i+1}) \leq D(n, 10t) + 10t$ . Each cluster  $C \in \mathcal{C}'_i$  has a diameter of at most D(n, 10t). Therefore,  $C^{+j}$  has a diameter of at most D(n, 10t) + 2j and as for each good cluster  $C \in \mathcal{C}'_i$ ,  $j_C < 4k$ , it directly follows that  $\operatorname{diam}(\mathcal{C}_{i+1}) \leq D(n, 10t) + 10t$ , as needed.

Next, we have to show that  $|V_{i+1}^{unclustered}| \leq (7/10)^{i+1}n$ . As  $|V_i^{unclustered}| \leq (7/10)^i n$ , it therefore suffices to show that  $|V_{i+1}^{unclustered}| \leq (7/10)|V_i^{unclustered}|$ . Note that  $C_i'$  clusters at least  $\frac{|V_i^{unclustered}|}{2}$ nodes. Hence, it suffices to show that at most one-fifth of the nodes in  $G_i$  are contained in bad clusters. As  $C_i'$  is 10t-separated, any two clusters  $C_1 \neq C_2 \in C_i'$  satisfy  $d_{G_i}(C_1^{+4t}, C_2^{+4t}) \geq 2t$ . In particular,  $C_1^{+4t}$  and  $C_2^{+4t}$  are disjoint and therefore  $|\bigcup_{C \in C_i'} C^{+4t}| \leq |V_i^{unclustered}|$ .

Moreover, for each bad cluster C we have  $|C^{+4t}| \ge |C| + 4t \cdot \frac{|C|}{t} = 5|C|$ . Therefore,

$$|\bigcup_{C\in\mathcal{C}'_i\colon C \text{ is bad}} C| \leq \frac{1}{5} |\bigcup_{C\in\mathcal{C}'_i} C^{+4t}| \leq \frac{1}{5} |V^{unclustered}_i|,$$

as needed.

For the third invariant, we have to show that  $|E_{i+1}^{inter}| \leq |\bigcup_{C \in \mathcal{C}_{i+1}} C|/t$ . As  $|E_i^{inter}| \leq |\bigcup_{C \in \mathcal{C}_i} C|/t$ , it suffices to show that  $|E_{i+1}^{inter} \setminus E_i^{inter}| \leq |\bigcup_{C \in \mathcal{C}_{i+1} \setminus \mathcal{C}_i} C|/t$ . Note that each newly added cluster C is neighboring with at most |C|/t nodes. As  $E_{i+1}^{inter} \setminus E_i^{inter}$  contains at most one edge with one endpoint in C for each neighboring vertex v, this implies that  $E_{i+1}^{inter} \setminus E_i^{inter}$  contains at most at most |C|/t edges with one endpoint in C, which shows that the third invariant is preserved.

Next, we verify that the fourth invariant is preserved. Consider any two neighboring clusters  $C_1, C_2 \in \mathcal{C}_{i+1}$ . We have to show that there exists at least one edge in  $E_{i+1}^{inter}$  with one endpoint in  $C_1$  and one endpoint in  $C_2$ . If both  $C_1$  and  $C_2$  were already contained in  $\mathcal{C}_i$ , then this property directly follows from induction. Note that  $C_1$  and  $C_2$  cannot both be newly added clusters, as this would imply that  $d_{G_i}(C_1, C_2) > 2t$  and therefore they could not be neighboring in G. Thus, it remains to consider the case that exactly one of the two clusters was newly added to  $\mathcal{C}_{i+1}$ , let's say  $C_1$ . As  $C_1$  and  $C_2$  are neighbors, there exists by definition a vertex v in  $C_1$  that is neighboring with  $C_2$ . As  $C_1$  is newly added, v was unclustered before, i.e.,  $v \in V_i^{unclustered}$ . In particular, the fifth invariant implies that  $E_i^{inter}$  contains an edge incident to v with the other endpoint being contained in  $C_2$ , as desired.

We now check the last invariant. Consider an arbitrary  $v \in V_{i+1}^{unclustered}$  and cluster  $C \in \mathcal{C}_{i+1}$  such that v is neighboring with C. We have to show that there exists a node  $u \in C$  with  $\{v, u\} \in E_{i+1}^{inter}$ . If  $C \in \mathcal{C}_i$ , then this follows from induction. Otherwise, it directly follows from the description of how we obtain  $E_{i+1}^{inter}$  from  $E_i^{inter}$ .

It remains to discuss the distributed implementation.

Computing the clustering  $C'_i$  takes T(n, 10t) CONGEST rounds. Next, we have to decide for each cluster  $C \in C'_i$  whether it is good or bad, and if C is a good cluster, we additionally have to compute the smallest good cutting distance of C. For each  $C \in C'_i$ , let  $F_C$  be a BFS forest in  $G_i$  of depth 4t with C being the set of roots. As  $C'_i$  is 10t-separated, it follows that  $F_C$  is disjoint from  $F_{C'}$  for any other cluster  $C' \in C'_i$ . Hence, we can compute  $F_C$  for all the clusters  $C \in C'_i$  in O(t)CONGEST rounds. Moreover, each node u in  $F_C$  gets to know its distance  $d_{G_i}(u, C)$  to C in  $G_i$ .

A standard pipelining idea therefore implies that each root of  $F_C$  can learn for every  $j \in \{0, 1, 2, ..., 4t\}$  how many nodes of depth j its tree contains in O(t) CONGEST rounds. As diam $(C) \leq D(n, 10t)$ , with another standard pipelining idea, the root of the cluster C can learn in O(t + D(n, 10t)) additional CONGEST rounds for each  $j \in \{0, 1, 2, ..., 4t\}$  how many vertices  $u \in V(G_i)$  with  $d_{G_i}(u, C) = j$  exists. This also allows the root to decide whether C is good or bad, and if it is good, to compute the smallest good cutting distance of C.

Finally, the root can broadcast this information to each node having a distance of at most 4t to C in O(t + D(n, 10t)) CONGEST rounds.

Moreover, after having computed the clustering  $C_{i+1}$ , it is easy to see that we can compute  $E_{i+1}^{inter}$  in O(1) additional CONGEST rounds.

Hence, the overall round complexity of the algorithm is  $O(\log n) \cdot (T(n, 10t) + D(n, 10t) + t)$ , as desired.

#### The Spanner Construction We are now ready to prove Theorem F.1.

Proof. We apply Lemma F.2 to get a complete clustering C and a set of edges  $E^{inter}$ . For each cluster  $C \in C$ , let  $T_C$  be a spanning tree of C with diameter O(D(n, 10t) + t). Such a spanning tree can be computed by a simple BFS from an arbitrary node in the cluster, as the diameter of each cluster is O(D(n, 10t) + t). Now, let F be the union of trees  $T_C$  for  $C \in C$ .

The spanner H now includes all the edges of F and all the edges of  $E^{inter}$ .

As F is a forest, it contains at most n-1 edges and as  $|E^{inter}| \leq \frac{n}{t}$ , we get  $|E(H)| \leq n+n/t$ .

Claim F.3. *H* is an O(D(n, 10t) + t)-spanner.

*Proof.* Consider any edge  $\{u, v\} \in E(G)$ . We have to show that  $d_H(u, v) = O(D(n, 10t) + t)$ .

First, consider the case that u and v belong to the same cluster. Hence, there exists a path of length O(D(n, 10t) + t) between u and v in F and therefore  $d_H(u, v) = O(D(n, 10t) + t)$ .

It remains to consider the case that  $u \in C_1$  while  $v \in C_2 \neq C_1$ . As  $C_1$  and  $C_2$  are neighbors, there exits at least one edge  $\{x, y\} \in E^{inter}$  with  $x \in C_1$  and  $y \in C_2$ . Now, u and x are connected by a path of length O(D(n, 10t)+t) in F and v and y are connected by a path of length O(D(n, 10t)+t) in F. Thud, u and v are connected by a path of length at most  $2 \cdot O(D(n, 10t)+t)+1 = O(D(n, 10t)+t)$  and therefore  $d_H(u, v) = O(D(n, 10t)+t)$ , as desired.

It remains to analyze the complexity. The algorithm invokes Lemma F.2 with round complexity  $O(\log n) \cdot (T(n, 10t) + D(n, 10t) + t).$ 

After that, we need  $O(\operatorname{diam}(\mathcal{C})) = O(D(n, 10t) + t)$  additional rounds to compute the forest F and therefore the construction of the whole spanner takes  $O(\log n) \cdot (T(n, 10t) + D(n, 10t) + t)$ CONGEST rounds.

#### F.2 Unweighted Sparse Spanners via Weak-Diameter Clusterings

In this section, we prove Theorem 1.7. It shows that an efficient algorithm for computing a 3-separated weak-diameter clustering, defined below, implies an efficient algorithm for computing a sparse unweighted spanner.

**Definition F.4** (*t*-separated Weak Diameter Clustering). Assume we are given an unweighted graph G and parameters  $t, \xi_{AVG}, D$ . A *t*-separated weak-diameter clustering with weak-diameter D is a clustering C, with each cluster C coming with a tree  $T_C, V(T_C) \supseteq C$  such that

- 1. For each  $C \in \mathcal{C}$  we have diam $(T_C) \leq D$ .
- 2. For each  $C_1 \neq C_2$  we have  $d(C_1, C_2) \geq t$ .
- 3. We have  $|\bigcup_{C \in \mathcal{C}} C| \ge n/2$ .

Moreover, if  $\xi(v)$  is the number of trees  $T_C, C \in \mathcal{C}$  that a node  $v \in G$  is contained in, then  $\sum_{v \in V(G)} \xi(v) \leq n \cdot \xi_{AVG}$ .

Note that each strong-diameter clustering is a weak-diameter clustering with the same parameters and  $\xi_{AVG} \leq 1$ . Therefore, Theorem 1.7 can be seen as a stronger result compared to Theorem F.1 in the sense that it only assumes an efficient algorithm for computing a weak-diameter clustering instead of a strong one. On the other hand, Theorem 1.7 cannot be used to directly obtain ultra-sparse spanners, which is however not a problem due to Theorem 1.2.

**Theorem 1.7.** Suppose there exists an algorithm A which for any unweighted n-vertex graph finds a 3-separated clustering with weak-diameter D(n) with average overlap at most  $\xi_{AVG}(n)$ . Then, there is an algorithm A' that builds a  $\beta$ -spanner H on an unweighted graph such that (1)  $|E(H)| = O(\xi_{AVG}(n)n), (2) \beta = O(D(n)).$ 

Furthermore, if A is deterministic then so is A' and

- If A requires at most T(n) CONGEST rounds then A' requires  $O(\log(n)T(n))$  CONGEST rounds.
- If A requires at most T(n,r) CONGEST rounds on a r-cluster-graph, then A' requires  $O(\log(n)(T(n,r)+r))$  CONGEST rounds on a r-cluster-graph.

• If A is a PRAM algorithm with T(n) depth and W(m, n) work then A' is a PRAM algorithm with depth  $O(\log(n)T(n))$  and work  $O(\operatorname{poly}(\log(n))(W(m, n) + m + n\xi_{AVG}(n)))$ .

Proof. The proof is very similar to the proof of Theorem F.1. We start with an empty clustering. In each step, we add additional clusters to the clustering until we have a complete clustering. In particular, consider a fixed step and let  $V^{unclustered}$  denote the set of unclustered vertices with respect to the current clustering. We now use algorithm  $\mathcal{A}$  to compute a 3-separated clustering  $\mathcal{C}$  in  $G[V^{unclustered}]$  with weak-diameter D(n) and average overlap at most  $\xi_{AVG}$  clustering at least half of the nodes. For each cluster  $C \in \mathcal{C}$ , we add the cluster C to the current clustering and we add all edges in the tree  $T_C$  to the final spanner. This way, we add at most  $\xi_{AVG}(n) \cdot |V^{unclustered}|$  many edges to the spanner. Moreover, for each unclustered node v neighboring with a cluster  $C \in \mathcal{C}$  (due to the 3-separation, each node neighbors with at most one cluster), we add an arbitrary edge from v to C to the final spanner, for a total of at most  $|V^{unclustered}|$  edges. As the number of unclustered nodes decreases by a factor of 2 in each step, all the nodes are clustered after  $O(\log n)$  steps and the total number of edges of the final spanner is  $O(\xi_{AVG}(n)n)$ . With the exact same argumentation as in the proof of Claim F.3, it follows that H is a  $\beta$ -spanner with  $\beta = O(D(n))$ . The runtime bounds readily follow.

#### F.3 Weighted Work-Efficient Ultra-Sparse Spanners

In this section, we show how to efficiently compute weighted ultra-sparse spanners.

**Theorem 1.8.** There is a deterministic work-efficient CONGEST algorithm that, given any n-vertex weighted graph G and  $t \ge 1$ , computes in  $O(t \log^{10} n)$  rounds an ultra-sparse spanner with n + n/tedges and stretch  $O(t \log^4 n \log(U+1))$ , where  $U \ge 1$  is the aspect ratio of the weights. There is also a deterministic PRAM algorithm that computes such a spanner in polylog(n)-time and with  $m \cdot \text{polylog}(n)$  work.

*Proof.* The starting point is the following theorem.

**Theorem F.5** (Theorem 1.12 in [RG20]). Given an undirected graph, we can build a weak-diameter clustering C such that

- 1. Each  $T_C$  for a  $C \in \mathcal{C}$  satisfies  $diam(T_C) = O(\log^3 n)$ .
- 2. Every two different clusters  $C_1, C_2 \in \mathcal{C}$  are at least 10-separated.
- 3. Every node of V(G) is in at most  $O(\log n)$  trees  $T_C$ .
- 4.  $|\bigcup \mathcal{C}| \ge n/2$ .

The algorithm needs  $O(\log^6 n)$  deterministic CONGEST rounds. <sup>6</sup>

We remark that the algorithm of [RG20] also works when each node is actually a cluster of diameter at most r, in which case the round complexity is  $O(r \log^6 n)$ . Therefore, Theorem 1.7 implies that we can compute an unweighted spanner on such a graph with  $O(n \log n)$  edges and stretch  $O(\log^3 n)$  in  $O(r \log^7 n)$  CONGEST rounds. Hence, the folklore reduction from weighted to unweighted implies that we can compute a weighted spanner with  $O(n \log n \log(U + 1))$  edges and stretch  $O(\log^3 n)$  in  $O(r \log^8 n)$  CONGEST rounds (as we are on a cluster graph, we need to compute the  $O(\log(U + 1))$  unweighted spanners one after the other). This algorithm together

 $<sup>^{6}</sup>$ The version of the theorem in [RG20] has an additional log factor as it builds a network decomposition instead of a single clustering.

with the reduction from ultra-sparse to sparse (Theorem 1.2) then implies the CONGEST part of Theorem 1.8. The PRAM bound follows from the fact that (1) the algorithm of [RG20] can be implemented in near-linear work and polylogarithmic time, (2) the sparse spanner can therefore be computed in near-linear work and polylogarithmic depth, and (3) the reduction from sparse to ultra-sparse can be implemented in near-linear work and polylogarithmic time.

# G Connectivity Certificates

A simple and common way to find a k-connectivity certificate of a graph is to repeatedly extract a skeleton from G for k times. More concretely, in the *i*-step, we remove a skeleton  $H_i$  from  $G \setminus \bigcup_{j=1}^{i-1} H_j$ . At the end,  $H = \bigcup_{i=1}^k H_i$  should be a k-connectivity certificate of G. Indeed, the output provides a stronger guarantee: It contains either all or at least k edges from each cut in G. For the proof of this, consider an arbitrary cut. If each  $H_i$  has an edge in the cut, then H has at least k such edges as  $H_i$ s are disjoint. So suppose there is an *i* such that  $H_i$  has no edge in the cut. Since  $H_i$  is a skeleton, there should be no such edge in  $G \setminus \bigcup_{j=1}^{i-1} H_j$ , implying that all edges in this particular cut should be in  $\bigcup_{i=1}^{i-1} H_j$ .

This algorithm gives us a certificate with size at most (n-1)k if we extract a forest at each step. Finding a forest in the distributed setting needs  $\Omega(D + \sqrt{n})$  many rounds, where D denotes the diameter of network G. Instead of extracting a forest per step, we can extract an ultra-sparse spanner at each step, as any spanner is also a skeleton by definition. Using Theorem 1.6, we can find a spanner with size  $n(1 + \varepsilon)$  in  $\frac{\text{polylog}(n)}{\varepsilon}$  rounds.

**Theorem G.1.** There is a deterministic distributed algorithm that computes a k-connectivity certificate of G with at most  $nk(1 + \varepsilon)$  edges in  $\frac{k \operatorname{polylog}(n)}{\varepsilon}$  rounds. Moreover, the output has either all or at least k edges of every cut in G.

As the k-connectivity certificate of each k-edge-connected graph has at least kn/2 edges, the theorem above gives us a  $2(1 + \varepsilon)$  approximation. The problem though is that the number of rounds is linear in k. To overcome this issue, we use random sampling of Karger [Kar99] and give a randomized algorithm that runs in polylog(n) rounds for any k.

According to [Kar99, Theorem 2.1], if we sample each edge of a k-edge-connected graph with probability  $p = O(\frac{\log n}{k\varepsilon^2})$ , then with high probability, the number of sampled edges in every cut is around its expectation up to a factor  $\varepsilon$ . More concretely, for a cut with c edges, at least  $pc(1 - \varepsilon)$  and at most  $pc(1 + \varepsilon)$  edges are sampled.

Suppose G is k-edge-connected. We split its edges randomly into  $Q = \Theta(\frac{k\varepsilon^2}{\log n})$  partitions. So with high probability, every cut in G with size c has size  $\frac{c}{Q}(1\pm\varepsilon)$  edges in every partition (Since there are Q = O(n) partitions, we can take union bound on Q applications of the sampling theorem).

We compute a  $\left(k' = \frac{k(1+\varepsilon)}{Q(1-\varepsilon)}\right)$ -connectivity certificate for each partition simultaneously using Theorem G.1. This takes  $\frac{k' \operatorname{polylog}(n)}{\varepsilon} = \frac{\operatorname{polylog}(n)}{\varepsilon^3}$  many rounds. We show that the union of these certificates is a k-connectivity certificate for G. For this, consider an arbitrary cut C with |C|edges. Recall that G is k-edge-connected, so  $|C| \ge k$ . To show that at least k edges of C is in the union of these Q certificates, we consider two cases based on the size of C:

- If |C| is at most  $\frac{k}{1-\varepsilon}$ , then each partition contains at most  $\frac{|C|(1+\varepsilon)}{Q} \leq \frac{k(1+\varepsilon)}{Q(1-\varepsilon)} \leq k'$  edges of C. From Theorem G.1, all of them should appear in the certificate of that partition and so, all the edges of C is in our output. - If |C| is more than  $\frac{k}{1-\varepsilon}$ , then each partition has at least  $\frac{|C|(1-\varepsilon)}{Q} \ge \frac{k}{Q}$  edges of C. Since  $k' \ge \frac{k}{Q}$ , the k'-connectivity certificate of each partition contains at least  $\frac{k}{Q}$  from C implying that there is at least  $k = \frac{k}{Q} \cdot Q$  edges of C in our output.

This concludes that if G is k-edge-connected, then the union of these Q certificates is also k-edgeconnected. If G is not k-edge-connected, then obviously any subgraph of it is not k-edge-connected. So our output is always a k-connectivity certificate of G. The number of edges in our certificate is  $Q \cdot (nk'(1+\varepsilon)) = nk \frac{(1+\varepsilon)^2}{1-\varepsilon}$ . Assuming  $\varepsilon < \frac{1}{2}$ , we have  $\frac{1}{1-\varepsilon} \le (1+2\varepsilon)$  and  $(1+\varepsilon)^2 \le (1+3\varepsilon)$ . So the total number of edges in the output is  $nk(1+8\varepsilon)$ . Replacing  $\varepsilon$  with  $\frac{\varepsilon}{8}$  proves Theorem 1.9.