



Local Algorithms for Sparse Spanning Graphs

Reut Levi¹ · Dana Ron² · Ronitt Rubinfeld^{3,4}

Received: 27 April 2018 / Accepted: 22 July 2019 / Published online: 3 August 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Constructing a spanning tree of a graph is one of the most basic tasks in graph theory. We consider a relaxed version of this problem in the setting of local algorithms. The relaxation is that the constructed subgraph is a *sparse spanning subgraph* containing at most $(1 + \epsilon)n$ edges (where n is the number of vertices and ϵ is a given approximation/sparsity parameter). In the local setting, the goal is to quickly determine whether a given edge e belongs to such a subgraph, without constructing the whole subgraph, but rather by inspecting (querying) the local neighborhood of e . The challenge is to maintain consistency. That is, to provide answers concerning different edges according to the *same* spanning subgraph. We first show that for general bounded-degree graphs, the query complexity of any such algorithm must be $\Omega(\sqrt{n})$. This lower bound holds for constant-degree graphs that have high expansion. Next we design an algorithm for (bounded-degree) graphs with high expansion, obtaining a result that roughly matches the lower bound. We then turn to study graphs that exclude a fixed minor (and are hence non-expanding). We design an algorithm for such graphs, which may have an unbounded maximum degree. The query complexity of this algorithm is $\text{poly}(1/\epsilon, h)$ (independent of n and the maximum degree), where h is the number of vertices in the excluded minor. Though our two algorithms are designed for very different types of graphs (and have very different complexities), on a high-level there are several similarities, and we highlight both the similarities and the differences.

This work is based on part of the extended abstract [32] and on the extended abstract [33]. Other parts of [32] have appeared in the journal publication [30].

Reut Levi: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 819702). It was also supported by ERC-CoG Grant 772839. Dana Ron: Supported by the Israel Science Foundation Grant Nos. 246/08 and 671/13. Ronitt Rubinfeld: Supported by the Israel Science Foundation Grant No. 1147/09, and by NSF Grants CCF-1217423 and CCF-1065125.

✉ Reut Levi
reut.levi@weizmann.ac.il; reuti.levi@gmail.com

Extended author information available on the last page of the article

Keywords Sublinear algorithms · Local algorithms · Sparse spanning subgraphs · Minor-free graphs · Property testing · Bounded degree model

1 Introduction

When dealing with large graphs, it is often important to work with a sparse subgraph that maintains essential properties, such as connectivity, bounded diameter and other distance properties (e.g., small *stretch* [48,49]), of the original input graph. Can one provide fast random access to such a sparsified approximation of the original input graph? In this work, we consider the property of connectivity: Given a connected graph $G = (V, E)$ over n vertices, find a sparse subgraph of G' that spans G . This task can be accomplished by constructing a spanning tree in linear time. However, it may be crucial to *quickly* determine whether a particular edge e belongs to such a subgraph G' , where by “quickly” we mean in time much faster than constructing all of G' . The hope is that by inspecting only some small local neighborhood of e , one can answer in such a way that maintains consistency with the same G' for all edges e . We focus on such algorithms (referred to as (centralized) *local* algorithms [53]), which are of use when we do not need to know the answer for every edge at any single point in time, or if there are several independent processes that want to determine the answer for edges of their choice, possibly in parallel.

If we insist that G' should have the *minimum* number of edges sufficient for spanning G , namely, that G' be a spanning tree, then it is easy to see that the task cannot be performed in general without inspecting almost all of G . To verify this, observe that if G consists of a single path, then the algorithm must answer positively on all edges, while if G consists of a cycle, then the algorithm must answer negatively on one edge. However, the two cases cannot be distinguished without inspecting a linear number of edges. On the other hand, suppose we allow the algorithm some more slackness. Namely, rather than requiring that G' be a tree, require that it be relatively sparse, i.e., contains at most $(1 + \epsilon)n$ edges, (allowing it to contain more than $n - 1$ edges). Now the algorithm may answer positively on all cycle edges, so distinguishing between these two cases is no longer necessary.

Given the above observation, the question we consider is whether a relaxed version of this task can be solved by inspecting a sublinear number of edges. Specifically, the relaxation is that the spanning subgraph G' may contain $(1 + \epsilon) \cdot n$ edges, for a given approximation/sparsity parameter ϵ . To be more precise we consider the problem as formulated next.

Definition 1 An algorithm \mathcal{A} is a Local Sparse Spanning Graph (LSSG) algorithm if, given $n \geq 1$, $\epsilon > 0$, a sequence of random bits $r \in \{0, 1\}^*$ and query access to the incidence-lists representation of a connected graph $G = (V, E)$ over n vertices,¹ it provides oracle access to a subgraph $G' = (V, E')$ of G such that:

1. G' is connected.

¹ Namely, the algorithm can query, for any vertex $v \in V$ and index i , who is the i th neighbor of v (where if v has less than i neighbors, then a special symbol is returned).

2. $|E'| \leq (1 + \epsilon) \cdot n$ with probability at least $1 - 1/\Omega(n)$, where E' is determined by G and r .

By “providing oracle access to G' ” we mean that on input $(u, v) \in E$, \mathcal{A} returns whether $(u, v) \in E'$, and for any sequence of edges, \mathcal{A} answers consistently with the same G' . We consider the following performance measures of the algorithm.

- *Query complexity* Number of graph-queries it makes in order to answer a single oracle-query.
- *Running time* the time it takes to answer a single oracle-query (in the word-RAM model).²
- Total number of random bits the algorithm uses (i.e. $|r|$).

An algorithm \mathcal{A} is an LSSG algorithm for a family of graphs \mathcal{C} if the above conditions hold, provided that the input graph G belongs to \mathcal{C} .

We note that the randomness of the algorithm does not affect its query complexity and running time, but only the number of edges in the sparse spanning graph.

The complexity of the algorithm may depend on a bound, d , on the maximum degree. We note that if $|E| = O(n)$ (i.e., the maximum or average degree is a constant), then clearly one can get a constant factor approximation (to the minimum number of edges necessary for connectivity) by simply taking $E' = E$. However, we are interested in a much better approximation of $(1 + \epsilon)$ for any given approximation parameter ϵ .

Two comments are in place regarding the success probability of an LSSG algorithm. First, Definition 1 can be modified so as to allow a higher probability of failure in terms of the sparsity (that is, higher than $1/\Omega(n)$ in Item 2). Second, the definition can be modified so as to allow a probability of failure in terms of connectivity (that is, in Item 1). Since our algorithms ensure the required sparsity with probability $1 - 1/\Omega(n)$ and always ensure connectivity, for the sake of succinctness, we have chosen not to introduce additional parameters in the definition.

We are interested in LSSG algorithms that for each given edge, perform as few queries as possible to G . We are also interested in bounding the total amount randomness used by such algorithms, and their running time (in the word-RAM model). Note that Item 2, which states that the set of edges E' is determined by the graph G and the randomness r , implies that the answers of an LSSG algorithm to queries cannot depend on previously asked queries.

1.1 Our Results

We start by showing that even after the aforementioned relaxation, local algorithms for sparse spanning graphs require the inspection of $\Omega(n^{1/2})$ edges for each given edge e and for a constant ϵ . This lower bound holds for constant-degree graphs with strong expansion properties, and in our first positive result we provide an almost matching upper bound for such graphs. Namely, the first algorithm we present, the Sparse-

² In this model (which is attributed to Fredman and Willard [19]), memory is divided into words of length w bits each. Every basic operation on words (in particular, comparison and addition), is assumed to take unit time. In our setting a word is $\Theta(\log n)$ bits.

Centers algorithm³ gives meaningful results for graphs in which the neighborhoods of almost all the vertices in the graph expand in a similar rate. In particular, for graphs with high expansion we get complexity that grows approximately like $n^{1/2}$. More precisely, if the expansion of small sets (of size roughly $O(n^{1/2})$) is $\Omega(d)$, where d is the maximum degree in the graph, then the complexity of the algorithm is $n^{1/2+O(1/\log d)}$. The number of random bits the algorithm uses, in terms of n , is $\tilde{O}(\sqrt{n})$.⁴

We then turn to consider non-expanding graphs. A family of non-expanding graphs that is of wide interest, is the family of *minor-free* graphs, which can be parameterized by the size, h , of the excluded minor. In particular, this family includes planar graphs. Note that although the average degree in graphs with an excluded fixed minor is constant, the maximum degree might be large as $\Omega(n)$. Our second algorithm, the Dense-Centers algorithm, is designed for minor-free graphs, and has complexity polynomial in h and $1/\epsilon$ (with no dependence on n or on the maximum degree,⁵ so that it works for unbounded-degree graphs). Another parameter of interest is the *stretch*⁶ of G' [48,49]. The spanning subgraph G' has *stretch* $\tilde{O}(h/\epsilon)$. The length of the random seed the algorithm uses is $\text{poly}(h/\epsilon) \cdot \log n$.

Observe that while the query complexity of our algorithm for expanding graphs has relatively high dependence on n , our algorithm for minor-free graphs has no dependence on n . This raises the natural question regarding the relation between expansion properties of graphs and the query complexity of the LSSG problem (in terms of the dependence on n). This question was addressed in [30], as we further elaborate in Sect. 1.3.

1.2 Our Algorithms

Interestingly, though our two algorithms are designed for very different types of graphs (and have very different complexities), on a high-level there are several similarities. In particular, underlying each of the two algorithms is the existence of a (global) partition of the graph vertices where the partition determines the sparse spanning subgraph, and the local algorithm decides whether a given edge belongs to the spanning subgraph by constructing parts of the partition. The differences between the algorithms are in the way the partition is defined, the precise way the spanning subgraph is determined by the partition (in particular, the choice of edges between parts) and the local (partial) construction of the partition. We further discuss the similarities and differences between the algorithms in Sect. 1.2.3 following their descriptions, presented next.

³ In the extended abstract [32], the algorithm is referred to simply as the Centers algorithm. Here we name it the Sparse-Centers algorithm so as to distinguish it from the other (centers based) algorithm that is presented in this work.

⁴ We use \tilde{O} for $\tilde{O}(x) = O(x) \cdot \text{poly}(\log x)$.

⁵ The extension to unbounded degrees does not appear in the extended abstract [33] and appears only in this version.

⁶ Given a connected graph $G = (V, E)$, a subgraph $G' = (V, E')$ is a t -spanner of G if for every $u, v \in V$, $\frac{\Delta_{G'}(u,v)}{\Delta_G(u,v)} \leq t$ (where for a graph H and two vertices u, v in H , $\Delta_H(u, v)$ denotes the distance between u and v in H). We refer to t as the *stretch factor* of G' .

1.2.1 The Sparse-Centers Algorithm (for Expanding Graphs)

This algorithm is based on the following idea. Suppose we can partition the graph vertices into $\sqrt{\epsilon n}$ disjoint parts where each part is connected. If we now take a spanning tree over each part and select one edge between every two parts that have an edge between them, then we obtain a connected subgraph with at most $(1 + \epsilon)n$ edges.

The Partition The partition is defined based on $\sqrt{\epsilon n}$ special *center* vertices, which are selected uniformly at random. Each vertex is assigned to the closest center (breaking ties according to a fixed order over the centers), unless it is further than some threshold t from all centers, in which case it is a singleton in the partition. This definition ensures that for each center the subgraph induced on the vertices assigned to that center is connected.

The Local Algorithm Given an edge (u, v) , the local algorithm finds the centers to which u and v are assigned (or determines that they are singletons). This is simply done by performing a Breadth First Search (BFS) to distance at most t from each of the two vertices, until a center is encountered (or no center is encountered, in which case we get a singleton). If one of the two vertices is a singleton, then (u, v) is always taken to the spanning subgraph. Otherwise, if u and v are assigned to the same center, σ , then the algorithm determines whether the edge between them belongs to the BFS-tree rooted at σ that spans the vertices assigned to σ . (Since there may be many such trees, we consider a fixed choice, based on the vertices' ids.)

To this end it is actually not necessary to determine for each vertex at distance at most t from σ whether it is assigned to σ or not (which would be too costly in terms of the local algorithm's query complexity). Rather, given that it is known that both u and v are assigned to σ , it suffices to compare the distances of u and v to σ , and if they differ by 1, to determine the distances between the neighbors of the further vertex, and σ .

If u and v belong to different centers, $\sigma(u)$ and $\sigma(v)$, respectively, then the algorithm determines whether (u, v) is the single edge allowed to connect the parts corresponding to the centers $\sigma(u)$ and $\sigma(v)$. This is done according to a prespecified rule which, once again, can be decided without identifying exactly which vertices belong to the two parts. Specifically, consider the shortest path between $\sigma(u)$ and $\sigma(v)$ that is first in lexicographic order, denoted $P(\sigma(u), \sigma(v))$. Observe that due to the edge (u, v) , this path has length at most $2t + 1$. Therefore, it can be found efficiently by performing a BFS to depth t from both $\sigma(u)$ and $\sigma(v)$. If the edge between u and v belongs to $P(\sigma(u), \sigma(v))$, then it is taken to the spanning graph, and otherwise it is not taken. Indeed it may be the case that no edge between a vertex assigned to $\sigma(u)$ and a vertex assigned to $\sigma(v)$ is included in the spanning graph, despite the existence of such edges. This occurs if $P(\sigma(u), \sigma(v))$ passes through vertices that are assigned to other centers. As a consequence, the argument that the subgraph G' is connected is required to be more subtle. Essentially, we show by induction that $\sigma(u)$ and $\sigma(v)$ will be connected, but possibly by a path that passes through other centers.

The Complexity of the Algorithm From the above description one can see that there is a certain tension between the complexity of the algorithm and the number of edges

in the spanning subgraph G' . On one hand, the threshold t should be such that with high probability (over the choice of the centers) almost all vertices have a center at distance at most t from them. Thus we need a lower bound (of roughly $\sqrt{n/\epsilon}$) on the size of the distance- t neighborhood of (most) vertices. On the other hand, we also need an upper bound on the size of such neighborhoods so that we can efficiently determine which edges are selected between parts. Hence, this algorithm works for graphs in which the sizes of the aforementioned local neighborhoods do not vary by too much and its complexity (in terms of the dependence on n) is nearly $\tilde{O}(\sqrt{n})$. In particular this property holds for good expander graphs.

The Threshold t If the threshold t is not given to the algorithm, then as a first step, the algorithm approximates t . Namely, it finds a threshold t' that gives similar guarantees as t , both in terms of the sparsity of the output subgraph and in terms of the query complexity of the algorithm.

1.2.2 The Dense-Centers Algorithm (for Minor-Free Graphs)

In what follows we assume that the graph has a bounded degree, d , which is provided to the algorithm. We then explain how to remove this assumption.

Similarly to the Sparse-Centers algorithm, the algorithm for families of graphs with a fixed excluded minor is based on a global partition that is determined by a random set of centers. However, the number of centers in this algorithm is linear in n (hence we call it the Dense-Centers algorithm), and the definition of the partition is more involved. Also similarly to what was described in Sect. 1.2.1, given such a partition, the edge set E' of the spanning subgraph G' is defined by taking a spanning tree in each part, and a single edge between every pair of parts that have at least one edge between them (though the choice of the edge is different). Here, the minor-freeness of the graph, together with a bound on the number of parts, ensure that the number of pairs of parts that have at least one edge between them is at most ϵn , and therefore $|E'| \leq (1 + \epsilon)n$.

We now provide a high level description of the partition and the corresponding local algorithm.

Properties of the Partition We define a partition that has each of the following four properties (which for simplicity are not precisely quantified in this introductory text): (1) the number of parts in the partition is not too large (though it is linear in n); (2) the size of each part is not too large (it is polynomial in d , h and $1/\epsilon$); (3) each part induces a connected subgraph; (4) for every vertex v it is possible to efficiently find (that is, in time polynomial in d , h and $1/\epsilon$), the subset of vertices in the part that v belongs to.

An Initial Centers-Based Partition Initially, the partition is defined (as in the Sparse-Centers algorithm) by a selection of *centers*, where the centers are selected randomly (though not completely independently). Each vertex is initially assigned to the closest selected center. For an appropriate setting of the probability that a vertex is selected to be a center, with high probability, this initial partition has the first property. Namely,

the number of parts in the partition is not too large. By the definition of the partition, each part is connected, so that the partition has the third property as well. However, the partition is not expected to have the second property. That is, some parts may be too large. Even for a simple graph such as the line, we expect to see parts of size logarithmic in n . In addition, the same example shows that the fourth property may not hold, i.e., there may be many vertices for which it takes super-constant time to find the part that the vertex belongs to. To deal with these two problems, we give a procedure for refining the partition in two phases, as explained next.

Refining the Partition, Phase 1 (and a Structural Lemma) A first refinement is obtained by the separation of vertices that in order to reach their assigned center in a BFS, must observe a number of vertices that is above a certain predetermined threshold, k . Each of these *remote* vertices becomes a singleton subset in the partition. We prove that with probability $1 - 1/\Omega(n)$, the number of these remote vertices is not too large, so that the first property (concerning the number of parts) is maintained with high probability.

The probabilistic analysis builds on a structural lemma, which may be of independent interest. The lemma upper bounds, for any given vertex v , the number of vertices u such that v can be reached from u by performing a BFS until at most k vertices are observed. This number in turn upper bounds the size of each part in the partition after the aforementioned refinement. While this upper bound is not polynomial in k and d , it suffices for the purposes of our probabilistic analysis (and we also show that there is no polynomial upper bound). Specifically, we use this bound in order to obtain a bound on the variance of the number of remote vertices.

In addition to the first property, the third property (connectivity of parts) is also maintained in the resulting refined partition, and the refinement partially addresses the fourth property, as remote vertices can quickly determine that they are far from all centers. In addition, the new parts of the partition will be of size 1 and thus will not violate the second property.

Refining the Partition, Phase 2 After the the first phase of the refinement, there might be some large parts remaining so that the second and fourth properties are not necessarily satisfied. We further partition large parts into smaller (connected) parts by a certain decomposition based on the BFS-tree of each large part.

The Local Algorithm Given an edge $(u, v) \in E$, the main task of the local algorithm is to find the two parts to which the vertices belong in the final partition. Once these parts are found, the algorithm can easily decide whether the edge between u and v should belong to E' or not (since each part is small, and every vertex has a bounded degree). In order to find the part that u (similarly, v) belongs to, the local algorithm does the following. First it performs a BFS until it finds a center, or it sees at least k vertices (none of which is a center). In the latter case, u is a singleton (remote) vertex. In the former case, the algorithm has found the center that u is assigned to in the initial partition, which we denote by $\sigma(u)$. The algorithm next performs a BFS starting from $\sigma(u)$. For each vertex that it encounters, it checks whether this vertex is assigned to $\sigma(u)$ (in the initial partition). If the number of vertices that are found to be assigned to

$\sigma(u)$ is above a certain threshold, then the decomposition of the part assigned to $\sigma(u)$ needs to be emulated locally. We show that this can be done recursively in an efficient manner.

The Unbounded-Degree Case We next describe how to remove the dependence on the given degree bound d and obtain an algorithm that works for unbounded-degree minor-free graphs. For a degree threshold $\tilde{d} = \tilde{d}(h, \epsilon)$, we say that a vertex u is a *heavy* vertex if the degree of u is greater than \tilde{d} , and otherwise it is a *light* vertex. We denote the set of heavy vertices by \mathcal{H} , and the set of light vertices by \mathcal{L} . The degree threshold \tilde{d} is set so that $|\mathcal{H}|$ is sufficiently small, and (since the graph is minor-free) the total number of edges between vertices in \mathcal{H} is $\epsilon n/4$. This already implies that we can take to the sparse spanning graph all edges between vertices in \mathcal{H} .

Consider the graph $G[\mathcal{L}]$ induced by \mathcal{L} . Observe that given a vertex v , we can determine with a single query whether it is heavy or light (simply by checking if it has a neighbor at index \tilde{d} or not). Furthermore, we can run a BFS on $G[\mathcal{L}]$ starting from any light vertex, by performing at most $\tilde{d} \cdot q$ queries, where q is the number of light vertices reached in the BFS. Also note that $G[\mathcal{L}]$ is not necessarily connected.

Suppose we run a slight variant of the algorithm for bounded-degree graphs on $G[\mathcal{L}]$ with the degree d set to \tilde{d} . The variant is that within each connected component of size less than k (where the parameter k is as defined for the bounded-degree case), we simply construct a (BFS) spanning tree rooted at the vertex with the smallest rank (id) in the component (this can be easily implemented in a local manner). For larger connected components, the bounded-degree algorithm is run as is. We thus obtain a graph that spans each connected component of $G[\mathcal{L}]$, and contains at most $|\mathcal{L}| + \epsilon n/4$ edges with high probability.

It remains to decide which edges between heavy and light vertices should be taken to the spanning graph. For each part P of the partition computed on $G[\mathcal{L}]$ (where each small connected component constitutes a part), we can select a single edge connecting the part to a heavy vertex (if such an edge exists). We say in such a case that the vertices in P are *assigned* to u , and refer to the set of vertices assigned to u , together with u as the *cluster* of u , denoted $\mathcal{C}(u)$.

The remaining decisions are with respect to edges between vertices in different clusters: a light vertex v in one cluster and a heavy vertex u in another. This is the challenging part in the extension of the algorithm to unbounded-degree graphs. Ideally, we would like to select one edge between every two clusters that have an edge between them. The difficulty is that due to the high degree of heavy vertices, clusters may be very large, so we cannot afford to determine for a given vertex who are all the vertices in its cluster. On the other hand, if we take all edges between clusters, then sparsity may not be maintained.⁷ We address this difficulty by selecting only a small fraction of edges between adjacent clusters, or all of them if there aren't too many (see more details in Sect. 5.5.3).

⁷ To illustrate this, consider the following (planar) graph. It has two vertices, u_1 and u_2 whose degree is $D = (n - 2)/2$ each. Let us denote the neighbors of u_1 by $v_{1,1}, \dots, v_{1,D}$, and the neighbors of u_2 by $v_{2,1}, \dots, v_{2,D}$, where these two subsets of vertices are disjoint and there is an edge between $v_{1,j}$ and $v_{2,j}$ for each $j = 1, \dots, D$. Observe that for this graph, all connected components of $G[\mathcal{L}]$ are small (with 2 vertices). For any assignment of the light vertices to the heavy ones, the number of edges between clusters is $n/2 - 1$.

1.2.3 A Discussion of the Similarities and Differences Between the Algorithms

As noted previously, our two local algorithms share several common features (but differ in others). We next emphasize both the similarities and the differences.

1. The first and most basic similarity is that they are both based on a partition of the vertices into connected subsets. Furthermore, the partition is (either fully or partially) determined by Voronoi cells of centers, where the centers are chosen randomly, and (relatively few) cells of single vertices.
2. The second similarity is that the sparse spanning subgraph is defined by the edges of BFS trees, one for each part of the partition (rooted at the centers), and (at most) one edge between every pair of parts, (or relatively few in the unbounded-degree variant of the Dense-Centers algorithm).
3. The first difference is in the number of centers (and hence the number of parts). For the Sparse-Centers algorithm this number grows like \sqrt{n} , and for the Dense-Centers algorithm it is linear in n . The bound on the number of parts in the case of the Sparse-Centers algorithm immediately gives a bound on the sparsity of the resulting spanning subgraph, while in the case of the Dense-Centers algorithm we rely on the minor-freeness of the graph.
4. The second difference is in the way singletons are defined. For the Sparse-Centers algorithm these are vertices that in their bounded-distance neighborhood there is no center, and for the Dense-Centers algorithm, these are vertices that in their bounded-size neighborhood there is no center (where the neighborhood is allowed to have relatively large diameter).
5. The third difference is in the way the edges between parts are selected, and identified (locally), and the related need to refine the partition for the Dense-Centers algorithm. Specifically, given a vertex u , the Dense-Centers algorithm finds all the vertices that belong to u 's part. (Here we are not referring yet to the clusters, which are defined in the unbounded-degree case, but rather only to the parts defined in the bounded-degree case.) It is then easy to decide for an edge (u, v) where u and v belong to different parts, whether the edge (u, v) belongs to the spanning subgraph. However, this requires to refine the partition so as to obtain sufficiently small parts, while being able to construct a part in a local manner. On the other hand, the Sparse-Centers algorithm decides whether an edge (u, v) belongs to the spanning subgraph by identifying the centers that u and v are assigned to, but without determining the vertices in the corresponding parts (and the spanning subgraph might not include an edge between some pairs of parts even though there are edges between them).
6. The fourth difference is that the Dense-Centers algorithm extends to unbounded-degree (minor-free) graphs, by applying an additional edge-sampling technique.

Given the above, the main challenge in the design and analysis of the Sparse-Centers algorithm is in maintaining connectivity in a sufficiently efficient manner by avoiding the need to construct complete parts of the partition. The main challenge in the design and analysis of the Dense-Centers algorithm is in proving that sparsity is obtained with high probability (at least $1 - \Omega(1/n)$). This requires a careful analysis of the number

of singletons (which relies on structural properties of the graph), and in bounding the number of edges between clusters in the unbounded-degree variant of the algorithm.

1.3 Related Work

The work that is most closely related to the current work is [30]. The upper bound in [30] builds on a part of the extended abstract [32] (while the current work builds on a different part). It was shown in [30] that for families of graph that are, roughly speaking, sufficiency non-expanding, one can provide an LSSG algorithm with query complexity that is independent of n (however, super-exponential in $1/\epsilon$). This is achieved by simulating a localized version of Kruskal's algorithm. On the negative side, it was also shown in [30] that for graphs with expansion properties that are a little better, there is *no local algorithm* that inspects a number of edges that is independent of n .

A couple of additional LSSG algorithms for families of minor-free graphs, which we do not include in this paper, are presented in [32]. The first algorithm is based on a reduction to the partition oracle in [31] and the second one is designed for weighted graphs. For the unweighted case, the result presented here for minor-free graphs significantly improves on what was shown in [32]. Specifically, the algorithm presented in [32] has complexity $(d/\epsilon)^{\text{poly}(h) \log(1/\epsilon)}$ (as well as higher stretch, which in particular depends polynomially on d). Furthermore, the success probability of the algorithm presented in this work (that is, the probability that G' has at most $(1 + \epsilon)n$ edges), is $1 - 1/\Omega(n)$, while it was only ensured to be a high constant in the [32] algorithm.

Graph partitioning for graphs with an excluded minor has been studied extensively (see e.g., [4,26]), and was found useful in constructing spanners and distance oracles (see e.g., [25]). Compared to previous results, the main disadvantage of our partition is that the edge cut is not guaranteed to be small. However, its main benefit is that it is designed to be constructed locally in an efficient manner.

1.3.1 Followup Work on LSSG

Recently, Lenzen and Levi [28] provided an LSSG algorithm for general bounded degree graphs with query complexity $\tilde{O}(n^{2/3}) \cdot \text{poly}(1/\epsilon, d)$, where d is the bound on the maximum degree. Parter et al. [46] provide local algorithms for constructing spanners in graphs with unbounded degree, part of their techniques are inspired by the algorithms for LSSG on bounded degree graphs.

1.3.2 Local Algorithms for Other Graph Problems

The model of *local computation algorithms* as used in this work, was defined by Rubinfeld et al. [53] (see also [3] and survey in [29]). Such algorithms for maximal independent set, hypergraph coloring, k -CNF, approximated maximum matching and approximated minimum vertex cover for bipartite graphs are given in [3,17,18,34,37,38,53]. This model generalizes other models that have been studied in various contexts, including locally decodable codes (e.g., [57]), local decompression [15],

and local filters/reconstructors [1,11,14,23,24,54]. Local computation algorithms that give approximate solutions for various optimization problems on graphs, including vertex cover, maximal matching, and other packing and covering problems, can also be derived from sublinear time algorithms for parameter estimation [21,39,43,45,58].

Campagna et al. [12] provide a local reconstructor for connectivity. Namely, given a graph which is almost connected, their reconstructor provides oracle access to the adjacency matrix of a connected graph which is close to the original graph. We emphasize that our model is different from theirs, in that they allow the addition of new edges to the graph, whereas our algorithms must provide spanning subgraphs whose edges are present in the original input graph.

1.3.3 Distributed Algorithms

The term *local algorithms* was originally coined in the distributed context [35,40,41]. As observed by Parnas and Ron [45], local distributed algorithms can be used to obtain local computation algorithms as defined in this work, by simply emulating the distributed algorithm on a sufficiently large subgraph of the graph G . However, while the main complexity measure in the distributed setting is the number of rounds (and messages may be required to be of bounded size), our main complexity measure is the number of queries performed on the graph G . By this standard reduction, the bound on the number of queries (and hence running time) depends on the size of the queried subgraph and may grow exponentially with the number of rounds. Therefore, this reduction gives meaningful results only when the number of rounds is significantly smaller than the diameter of the graph.

Ram and Vicari [52] studied the problem of constructing sparse spanning subgraphs in the distributed LOCAL model [35] and provided an algorithm that runs in $\min\{D(G), O(\log n)\}$ number of rounds where $D(G)$ denotes the diameter of G .

In the next paragraphs, unless stated otherwise, we refer to algorithms in the CONGEST model (in which the size of the messages is bounded by $O(\log n)$) [50]. The problem of computing a minimum weight spanning tree in this model is a central one.

Kutten and Peleg [27] provided an algorithm that works in $O(\sqrt{n} \log^* n + D)$ rounds, where D denotes the diameter of the graph. Their result is nearly optimal in terms of the complexity in n , as shown by Peleg and Rubinfeld [47] who provided a lower bound of $\Omega(\sqrt{n}/\log n)$ rounds (when the length of the messages must be bounded).

Another problem studied in the distributed setting that is related to the one studied in this paper, is finding a sparse spanner. The requirement for spanners is much stronger since the distortion of the distance should be as minimal as possible. Thus, to achieve this property, it is usually the case that the number of edges of the spanner is super-linear in n . Spanners with only $O(n)$ edges are usually referred to as *ultra-sparse spanners*. Pettie [51] was the first to provide a distributed algorithm for finding ultra-sparse spanners without requiring messages of unbounded length or $O(D)$ rounds. The number of rounds of his algorithm is $\log^{1+o(1)} n$. Recently, Elkin and Neiman [16] introduced an algorithm that produces a spanner of size $n(1 + o(1))$, as long as the number of rounds is $\omega(\log n)$. However, in both cases, the standard reduction of [45] yields a local algorithm with a trivial linear bound on the query complexity.

1.3.4 Parallel Algorithms

The problems of computing a spanning tree and a minimum weight spanning tree were studied extensively in the parallel computing model (see, e.g., [8], and the references therein). However, these parallel algorithms have time complexity that is at least logarithmic in n and therefore do not yield an efficient algorithm in the local computation model. See [3,53] for further discussion on the relationship between the ability to construct local computation algorithms and the parallel complexity of a problem.

1.3.5 Local Cluster Algorithms

Local algorithms for graph theoretic problems have also been given for PageRank computations on the web graph [5,6,9,22,55]. Local graph partitioning algorithms have been presented in [6,7,44,56,59], which find subsets of vertices whose internal connections are significantly richer than their external connections in time that depends on the size of the cluster that they output. Even when the size of the cluster is guaranteed to be small, it is not obvious how to use these algorithms in the local computation setting where the cluster decompositions must be consistent among queries to all vertices.

1.3.6 Other Related Sublinear-Time Approximation Algorithms for Graphs

The problem of estimating the weight of a minimum weight spanning tree in sublinear time was considered by Chazelle, Rubinfeld and Trevisan [13]. They describe an algorithm whose running time depends on the approximation parameter, the average degree and the range of the weights, but does not directly depend on the number of nodes.

1.3.7 Organization

We start with some preliminaries in Sect. 2. In Sect. 3 we prove our lower bound (an alternative, direct proof can be found in Appendix B). In Sects. 4 and 5 we describe and analyze our two algorithms. Since we introduce many notations throughout the paper, for the aid of the reader we provide two tables with these notations in Appendix A.

2 Preliminaries

Let $G = (V, E)$ be a graph over n vertices. Each vertex $v \in V$ has an id, $id(v)$, where there is a full order over the ids. We assume we have query access to the incidence-lists representation of G . Namely, for any vertex v and index $i \geq 1$ it is possible to obtain the i^{th} neighbor of v by performing a query to the graph (if v has less than i neighbors, then a special symbol is returned).⁸

⁸ Graphs are allowed to have self-loops and multiple edges, but for our problem we may assume that there are no self-loops and multiple-edges (since the answer on a self-loop can always be negative, and the same is true for all but at most one among a set of parallel edges).

We denote the distance between two vertices u and v in G by $\Delta_G(u, v)$. For vertex $v \in V$ and an integer r , let $\Gamma_r(v, G)$ denote the set of vertices at distance at most r from v , referred to as the *distance- r neighborhood* of v , and let $C_r(v, G)$ denote the subgraph of G induced by $\Gamma_r(v, G)$. Let $n_r(G) \stackrel{\text{def}}{=} \max_{v \in V} |\Gamma_r(v, G)|$. When the graph G is clear from the context, we shall use the shorthand $\Delta(u, v)$, $\Gamma_r(v)$ and $C_r(v)$ for $\Delta_G(u, v)$, $\Gamma_r(v, G)$ and $C_r(v, G)$, respectively. We let $N(v)$ denote the set of neighbors of v .

For a subset of vertices X , we let $G[X]$ denote the subgraph of G induced by X .

The total order over the vertices induces a total order (ranking) ρ over the edges of the graph in the following straightforward manner: $\rho((u, v)) < \rho((u', v'))$ if and only if $\min\{u, v\} < \min\{u', v'\}$ or $\min\{u, v\} = \min\{u', v'\}$ and $\max\{u, v\} < \max\{u', v'\}$ (recall that $V = [n]$). The total order over the vertices also induces an order over those vertices visited by a Breadth First Search (BFS) starting from any given vertex v , and whenever we refer to a BFS, we mean that it is performed according to this order.

Recall that a graph H is called a *minor* of a graph G if H is isomorphic to a graph that can be obtained by zero or more edge contractions on a subgraph of G . A graph G is *H -minor free* if H is not a minor of G .

The next definition and theorem are used to measure the sparsity of minor-free graphs.

Definition 2 The *arboricity* of a graph is the minimum number of forests into which its edges can be partitioned.

The next theorem follows from [36] and [42].

Theorem 1 *Let H be a fixed graph over h vertices and let $G = (V, E)$ be an H -minor free graph. Then the arboricity of G is at most $c(h)$, where $c(h) = O(h \log h)$. In particular, $|E| \leq c(h) \cdot |V|$.*

We shall also make use of the following theorem regarding t -wise independent random variables. Recall that a discrete random variable $\chi = (\chi_1, \dots, \chi_m) \in [q]^m$ (where $[q] \stackrel{\text{def}}{=} \{1, \dots, q\}$) is *t -wise independent* if for every subset $I = (i_1, \dots, i_s) \subseteq [m]$ such that $s \leq t$ and every set of values $(a_1, \dots, a_s) \in [q]^s$, $\Pr\left[\bigwedge_{j \in [s]} (\chi_{i_j} = a_j)\right] = \prod_{j \in [s]} \Pr[\chi_{i_j} = a_j]$.

The following theorem is implied by the proof of Proposition 6.3 in [2].

Theorem 2 (Implied by [2]) *For every $1 \leq t \leq m$, there exists an explicit construction of a t -wise independent random variable $\chi = (\chi_1, \dots, \chi_m)$ for the following two cases:*

1. $(\chi_1, \dots, \chi_m) \in \{0, 1\}^m$ and for every $i \in [m]$, $\Pr[\chi_i = 1] = a_i/b_i$ where $\{a_i\}_{i=1}^m, \{b_i\}_{i=1}^m$ are integers and $b_i \leq m$ for every $i \in [m]$.
2. $(\chi_1, \dots, \chi_m) \in [q]^m$ and the distribution of χ_i , for every $i \in [m]$, is the uniform distribution over $[q]$, where $q \geq m$ is any prime power.

The seed length (the number of independent random bits used in the construction) is $O(t \log m)$. Moreover, for any $1 \leq i \leq m$, X_i can be computed in $O(t)$ operations in the word-RAM model.

3 A Lower Bound for General Bounded-Degree Graphs

Theorem 3 *Any LSSG algorithm must have query complexity $\Omega(\sqrt{n})$. This result holds for graphs with a constant degree bound d and for constant $0 \leq \epsilon \leq 2d/3$. Furthermore, the result holds even if the algorithm is allowed a constant failure probability (i.e., the number of edges in the subgraph G' may be larger than $(1 + \epsilon)n$ with small constant probability).*

The theorem, with a slightly lower upper bound on ϵ , can be proved based on a reduction from one-sided error testing of cycle-freeness. We provide the reduction below. In the appendix we give a direct proof for Theorem 3, where in fact, we prove that a similar statement holds even when G' is allowed to be disconnected with some constant probability. In both proofs of the lower bound the construction is of constant-degree graphs with strong expansion properties and hence the lower bound holds even when restricted to these families of graphs.

A Reduction from Testing Cycle-Freeness with One-Sided Error In [20, Proposition 3.4] Goldreich and Ron showed that any one-sided error algorithm for testing cycle-freeness in bounded-degree graphs must perform $\Omega(\sqrt{n})$ queries. Such a tester must accept every cycle-free graph with probability 1, and must reject with probability at least $2/3$ every graph that is ϵ -far from being cycle free. A graph is said to be ϵ -far from being cycle-free, if more than ϵdn edges must be removed in order to make it cycle-free (where d is the degree bound). In other words, the graph contains at least $(1 + \epsilon d)n$ edges. The lower bound in [20] holds even if the input graph is promised to be connected, and for constant ϵ and d .

We observe that a lower bound of $\Omega(\sqrt{n})$ queries for the LSSG problem is implied by this lower bound for testing cycle-freeness. To verify this, consider the following reduction. Given an LSSG algorithm \mathcal{A} , one can obtain a tester \mathcal{B} for cycle-freeness as follows. On input graph $G = (V, E)$, \mathcal{B} samples uniformly at random $\Theta(1/\epsilon)$ pairs (v, i) where $v \in V$ and $i \in [d]$. It queries G to obtain each of the corresponding edges (where some may not exist), and then queries \mathcal{A} on all sampled edges, with sparsity parameter set to $\epsilon d/2$. The algorithm \mathcal{B} accepts if and only if \mathcal{A} returns YES on all the edges in the sample.

It remains to analyze \mathcal{B} . **Completeness:** If G is cycle-free, then \mathcal{A} must return YES on all edges in E , and hence \mathcal{B} accepts with probability 1. **Soundness:** if G is ϵ -far from being cycle-free, then \mathcal{A} must, with probability $1 - 1/\Omega(n)$ return NO on at least $(\epsilon d - \epsilon d/2)n = \epsilon d/2$ edges. Let \mathcal{E} denote the event that indeed \mathcal{A} returns NO on at least $\epsilon d/2$ edges of G (recall that these edges are determined only by G and the randomness of \mathcal{A}). Thus, given that \mathcal{E} occurs (which happens with probability $1 - 1/\Omega(n)$), each pair (v, i) that \mathcal{B} samples is a witness for rejection with probability at least ϵ . Hence, with probability $1 - (1 - \epsilon)^{\Theta(1/\epsilon)}$, \mathcal{B} samples one of these edges and rejects. Thus, \mathcal{B} rejects with high constant probability, as desired.

4 Graphs with High Expansion: The Sparse-Centers Algorithm

In this section we describe an algorithm that gives meaningful results for graphs in which, roughly speaking, the local neighborhood of almost all vertices expands in a similar rate and d is either a constant or grows very moderately as a function of n . In particular this includes graphs with high expansion. In fact we only require that the graph expands quickly for small sets: A graph G is an (s, h) -vertex expander if for all sets S of size at most s , $N^+(S)$ is of size at least $h \cdot |S|$, where $N^+(S)$ denotes the set of vertices in S and vertices adjacent to vertices in S . Define $h_s(G)$ to be the maximum h such that G is an (s, h) -vertex expander (hence, $1 \leq h_s(G) \leq d$).

Recall that for a vertex $v \in V$ and an integer r , $\Gamma_r(v, G)$ denotes the set of vertices at distance at most r from v , and $n_r(G) \stackrel{\text{def}}{=} \max_{v \in V} |\Gamma_r(v, G)|$. We shall make use of the following simple claim.

Claim 1 *For every two integers t and s , $n_t(G) \geq \min\{(h_s(G))^t, s\}$.*

Proof By definition, for every integer s , and graph G is an $(s, h_s(G))$ -vertex expander. Thus, for every vertex v and for every integer r for which $|\Gamma_r(v, G)| \leq s$, it holds that $|\Gamma_{r+1}(v, G)| \geq h_s(G) \cdot |\Gamma_r(v, G)|$ and so the claim follows. \square

We shall prove the following theorem (where Algorithm 1 appears in Sect. 4.2).

Theorem 4 *Algorithm 1 is an LSSG algorithm. For any graph G over n vertices, and for every $\epsilon > 0$, its query complexity is $O(d^2 \cdot s^{\log_{h_s(G)} d})$ for $s = s(n, \epsilon) \stackrel{\text{def}}{=} 2\sqrt{2n/\epsilon} \ln n$. The number of random bits the algorithm uses is $O(\sqrt{\epsilon n} \log n)$, and its running time is the same as its query complexity up to a factor of $O(\log n)$.*

By Theorem 4, for bounded degree graphs with high expansion we get query and running time complexity nearly $O(n^{1/2})$. In particular, if $h_s(G) = \Omega(d)$ then the complexity is $d^2 \cdot n^{1/2+O(1/\log d)}$. In fact, even for $h_s(G) \geq d^{1/2+1/\log n}$ the complexity is $o(n)$. We remark that in the construction of our lower bound of $\Omega(n^{1/2})$, that appears in Sect. B, we construct a pair of families of d -regular random graphs. In both families, the expansion (of small sets) is $\Omega(d)$, implying that for these families the gap between our lower bound and upper bound is at most $n^{O(1/\log d)}$.

Our algorithm, the Sparse-Centers algorithm (which appears as Algorithm 1 in Sect. 4.2), is based on a global algorithm, which is presented in Sect. 4.1. The (local) Sparse-Centers algorithm appears in Sect. 4.2 and it is analyzed in the proof of Theorem 5.

4.1 The Global Algorithm

For a given parameter t , the global algorithm first defines a global partition of (part or all of) the graph vertices in the following randomized manner.

1. Select $\ell \stackrel{\text{def}}{=} \sqrt{\epsilon n/2}$ centers uniformly and independently at random from V , and denote them $\sigma_1, \dots, \sigma_\ell$.
2. Initially, all vertices are *unassigned*.

3. For $i = 0, \dots, t$, for $j = 1, \dots, \ell$:
 Let L_j^i denote the vertices in the i^{th} level of the BFS tree of σ_j (where $L_j^0 = \{\sigma_j\}$).
 Assign to σ_j all vertices in L_j^i that were not yet assigned to any other $\sigma_{j'}$.

Let $S(\sigma_j)$ denote the set of vertices that are assigned to the center σ_j . By the above construction, the subgraph induced by $S(\sigma_j)$ is connected. To verify this, consider any vertex $x \in S(\sigma_j)$ and observe that the parent of x in the BFS tree of σ_j has to be in $S(\sigma_j)$ as well (otherwise we get a contradiction to the fact that $x \in S(\sigma_j)$).

The subgraph $G' = (V, E')$ is defined as follows.

1. For each center σ , let $E'(\sigma)$ denote the edges of a BFS-tree that spans the subgraph induced by $S(\sigma)$ (where the BFS-tree is determined by the order over the ids of the vertices in $S(\sigma)$). For each center σ , put in E' all edges in $E'(\sigma)$.
2. For each vertex w that does not belong to any $S(\sigma)$ for a center σ , put in E' all edges incident to w .
3. For each pair of centers σ and σ' , let $P(\sigma, \sigma')$ be the shortest path between σ and σ' that has minimum lexicographic order among all shortest paths (as determined by the ids of the vertices on the path). If all vertices on this path belong either to $S(\sigma)$ or to $S(\sigma')$, then add to E' the single edge $(x, y) \in P(\sigma, \sigma')$ such that $x \in S(\sigma)$ and $y \in S(\sigma')$, where we denote this edge by $e(\sigma, \sigma')$.

In what follows we shall prove that G' is connected and that for t that is sufficiently large, G' is sparse with high probability as well. We begin by proving the latter claim. To this end we define a parameter that determines the minimum distance needed for most vertices to see roughly \sqrt{n} vertices. More formally, define $t_{\epsilon,d}(G)$ to be the minimum distance t ensuring that all but an $\epsilon/(2d)$ -fraction of the vertices have at least s vertices in their t -neighborhood. That is,

$$t_{\epsilon,d}(G) \stackrel{\text{def}}{=} \min_t \{|\{v : |\Gamma_t(v)| \geq s\}| \geq (1 - \epsilon/(2d)) |V|\} . \tag{1}$$

Lemma 2 *For $t \geq t_{\epsilon,d}(G)$ it holds that $|E'| \leq (1 + \epsilon)n$ with probability at least $1 - 1/n$ over the random choice of centers.*

Proof Since for $j = 1, \dots, \ell$ the sets $E'(\sigma_j)$ are disjoint, we have that $\left| \bigcup_{j=1}^{\ell} E'(\sigma_j) \right| < n$. Since there is at most one edge (x, y) added to E' for each pair of centers σ, σ' and the number of centers is $\ell (= \sqrt{\epsilon n}/2)$, the total number of these edges in E' is at most $\epsilon n/2$. Finally, Let $V_{t,s} \subseteq V$ denote the subset of the vertices, v , such that $|\Gamma_t(v)| \geq s$. Since the centers are selected uniformly, independently at random, for each $w \in V_{t,s}$ the probability that no vertex in $\Gamma_t(w)$ is selected to be a center, is at most $(1 - s/n)^\ell < 1/n^2$. By taking a union bound over all vertices in $V_{t,s}$, with probability at least $1 - 1/n$, every $w \in V_{t,s}$ is assigned to some center σ . Since the number of vertices in $V \setminus V_{t,s}$ is at most $\epsilon n/(2d)$ and each contributes at most d edges to E' , we get the desired upper bound on $|E'|$. □

Lemma 3 *G' is connected.*

Proof It suffices to prove that there is a path in G' between every pair of centers σ and σ' . This suffices because for each vertex w that is assigned to some center σ , there

is a path between w and σ (in the BFS-tree of σ), and for each vertex w that is not assigned to any center, all edges incident to w belong to E' . The proof proceeds by induction on $\Delta(\sigma, \sigma')$ and the sum of the ids of σ and σ' as follows.

For the base case, consider a pair of centers σ and σ' for which $\Delta(\sigma, \sigma') = 1$. In this case, the shortest path $P(\sigma, \sigma')$ consists of a single edge (σ, σ') (where $\sigma \in S(\sigma)$ and $\sigma' \in S(\sigma')$ by the definition of $S(\cdot)$), implying that $(\sigma, \sigma') \in E'$.

For the induction step, consider a pair of centers σ and σ' for which $\Delta(\sigma, \sigma') > 1$, and assume by induction that the claim holds for every pair of centers (τ, τ') such that either $\Delta(\tau, \tau') < \Delta(\sigma, \sigma')$ or $\Delta(\tau, \tau') = \Delta(\sigma, \sigma')$ and $id(\tau) + id(\tau') < id(\sigma) + id(\sigma')$. Similarly to base case, if the set of vertices in $P(\sigma, \sigma')$ is contained entirely in $S(\sigma) \cup S(\sigma')$, then σ and σ' are connected by construction. Namely, $P(\sigma, \sigma') = (\sigma, x_1, \dots, x_q, x'_s, \dots, x'_{q'}, \sigma')$ where $x_1, \dots, x_q \in S(\sigma)$ and $x'_1, \dots, x'_{q'} \in S(\sigma')$. The edge $(x_q, x'_{q'})$ was added to E' and there are paths in the BFS-trees of σ and σ' between σ and x_q and between σ' and $x'_{q'}$, respectively.

Otherwise ($P(\sigma, \sigma')$ is not contained entirely in $S(\sigma) \cup S(\sigma')$), we consider two cases.

1. There exists a vertex x in $P(\sigma, \sigma')$, and a center (different from σ and σ'), τ , such that $x \in S(\tau)$. Note that this must be the case when $\Delta(\sigma, \sigma') \leq 2t + 1$. This implies that either $\Delta(x, \tau) < \Delta(x, \sigma')$ or that $\Delta(x, \tau) = \Delta(x, \sigma')$ and $id(\tau) < id(\sigma')$. Hence, either

$$\Delta(\sigma, \tau) \leq \Delta(\sigma, x) + \Delta(x, \tau) < \Delta(\sigma, x) + \Delta(x, \sigma') = \Delta(\sigma, \sigma')$$

or $\Delta(\sigma, \tau) = \Delta(\sigma, \sigma')$ and $id(\sigma) + id(\tau) < id(\sigma) + id(\sigma')$. In either case we can apply the induction hypothesis to obtain that σ and τ are connected. A symmetric argument gives us that σ' and τ are connected.

2. Otherwise, all the vertices on the path $P(\sigma, \sigma')$ that do not belong to $S(\sigma) \cup S(\sigma')$ are vertices that are not assigned to any center. Since E' contains all edges incident to such vertices, σ and σ' are connected in this case as well.

The proof of Lemma 3 is completed. □

4.2 The Local Algorithm

Theorem 5 *Algorithm 1, when run with $t \geq t_{\epsilon,d}(G)$, is an LSSG algorithm. The query complexity and the running time of the algorithm are $O(d \cdot n_t(G))$ and $O(d \cdot n_t(G) \cdot \log n)$, respectively. The random seed it uses is of size $O(\sqrt{\epsilon n} \log n)$.*

Proof We prove the theorem by showing that Algorithm 1 is a local emulation of the global algorithm that appears in Sect. 4.1. Given u and v , by performing a BFS to depth t from each of the two vertices, Algorithm 1 either finds the centers $\sigma(u)$ and $\sigma(v)$ that u and v are (respectively) assigned to (by the global algorithm, for the same selection of centers), or for at least one of them it finds no center in its distance- t neighborhood. In the latter case, the edge (u, v) belongs to E' , and Algorithm 1 returns a positive answer, as required. In the former case, there are two subcases.

Algorithm 1 – Sparse-Centers Algorithm (with known t)

For a random choice of $\ell = \sqrt{\epsilon n/2}$ centers, $\sigma_1, \dots, \sigma_\ell$ in V (which is fixed for all queries), and for a given parameter t , on query (u, v) :

1. Perform a BFS to depth t in G from u and from v .
2. If either $\Gamma_t(u) \cap \{\sigma_1, \dots, \sigma_\ell\} = \emptyset$ or $\Gamma_t(v) \cap \{\sigma_1, \dots, \sigma_\ell\} = \emptyset$, then return YES.
3. Otherwise, let $\sigma(u)$ be the center closest to u and let $\sigma(v)$ be the center closest to v (if there is more than one such center, break ties according to the order $\sigma_1, \dots, \sigma_\ell$).
4. If $\sigma(u) = \sigma(v) = \sigma$ then do the following:
 - If $\Delta(u, \sigma) = \Delta(v, \sigma)$, then return NO.
 - If $\Delta(v, \sigma) = \Delta(u, \sigma) + 1$, then consider all neighbors w of v on a shortest path between v and σ . If there exists such a neighbor w for which $id(w) < id(u)$, then return NO, otherwise, return YES. (The case $\Delta(u, \sigma) = \Delta(v, \sigma) + 1$ is handled analogously).
5. If $\sigma(u) \neq \sigma(v)$, then perform a BFS to depth t from both of the centers $\sigma(u)$ and $\sigma(v)$. Query also all neighbors of vertices at distance t from $\sigma(u)$ and $\sigma(v)$ (note that for every vertex at distance less than t from one of these vertices, all its neighbors were queried in the course of the corresponding BFS). Find the shortest path between $\sigma(u)$ and $\sigma(v)$ that has the smallest lexicographic order and denote it by $P(\sigma(u), \sigma(v))$. Finding this path can be implemented by considering the subgraph induced by all queried edges. This subgraph contains all shortest paths between $\sigma(u)$ and $\sigma(v)$, from which the one having smallest lexicographic order can be selected. Return YES if both u and v belong to $P(\sigma(u), \sigma(v))$. Otherwise, return NO.

1. If u and v are assigned to the same center, that is, $\sigma(u) = \sigma(v) = \sigma$, then Algorithm 1 checks whether the edge (u, v) is an edge in the BFS-tree of σ (i.e., $(u, v) \in E'(\sigma)$). If u and v are on the same level of the tree (i.e., are at the same distance from σ), then Algorithm 1 returns a negative answer, as required. If v is one level further than u , then Algorithm 1 checks whether v has another neighbor, w , that is also assigned to σ , is on the same level as u and has a smaller id than u . Namely, a neighbor of v that is on a shortest path between v and σ and has a smaller id than u . If this is the case, then the edge (u, v) does not belong to the tree (but rather the edge (w, v)) so that the algorithm returns a negative answer. If no such neighbor of v exists, then the algorithm returns a positive answer (as required).
2. If u and v are assigned to different centers, that is, $\sigma(u) \neq \sigma(v)$, then Algorithm 1 determines whether $(u, v) = e(\sigma(u), \sigma(v))$ exactly as defined in the global algorithm: The algorithm finds $P(\sigma(u), \sigma(v))$ and returns a positive answer if and only if (u, v) belongs to $P(\sigma(u), \sigma(v))$. Notice that since $u \in S(\sigma(u))$ and $v \in S(\sigma(v))$, if (u, v) belongs to $P(\sigma(u), \sigma(v))$ then all the vertices on $P(\sigma(u), \sigma(v))$ belong to either $S(\sigma(u))$ or $S(\sigma(v))$. This is implied by the fact that for every center σ and a vertex w that is assigned to σ , it holds that every vertex on a shortest path between σ and w is also assigned to σ .

The bound on the query complexity of Algorithm 1 follows directly by inspection of the algorithm. The running time of performing Steps 2–4 is upper bounded by the query complexity of performing the BFS, i.e., $O(d \cdot n_t(G))$, times the time sufficient for determining if a vertex is a center. The latter can be done in $O(\log n)$ time. To see this we explain next how the algorithm access the set of centers. In order to select ℓ random centers, the algorithm uses a random seed of length $O(\sqrt{\epsilon n} \log n)$, which is interpreted as a sequence of $O(\sqrt{\epsilon n})$ uniformly selected random centers (with

replacement). If we first sort this sequence of random centers in $O(\sqrt{\epsilon n} \log n)$ time, then we can decide for each vertex if it is a center in $O(\log n)$ time. Hence, the running time of performing Steps 2–4 is bounded by the query complexity of the algorithm up to a factor of $O(\log n)$. The running time of Step 5 is linear in the size of the subgraph induced on the BFS trees (from both centers) and therefore is linear in the query complexity. Overall, we obtain that the running time is $O(\log n)$ times the query complexity, as desired. \square

4.3 The Parameter t

Recall that Algorithm 1 is given a parameter t that determines the depth of the BFS that the algorithm performs. By Theorem 5 it suffices to require that $t \geq t_{\epsilon,d}(G)$ in order to ensure that the spanning graph obtained by the algorithm is sparse. For the case that t is not given in advance we describe next how to compute t such that with probability at least $1 - 1/n$ it holds that

$$t_{\epsilon,d}(G) \leq t \leq t'_{\epsilon,d}(G), \tag{2}$$

where $t'_{\epsilon,d}(G) = \min_t \{|\{v : \Gamma_t(v) \geq s\}| \geq (1 - \epsilon/(4d)) |V|\}$. Select uniformly at random $q = \Theta(\log n/\epsilon^2)$ vertices from V . Let x_1, \dots, x_q denote the selected vertices. For each vertex x_j in the sample, let $t_j = \min_t \{\Gamma_t(x_j) \geq s\}$. Assume without loss of generality that $t_1 \leq \dots, \leq t_q$ and set $t = t_{\lceil 1 - \frac{3\epsilon}{8d} \rceil}$. By Chernoff’s inequality we obtain that with probability greater than $1 - 1/n$ Equation (2) holds.

We are now ready to prove Theorem 4.

Proof of Theorem 4 Assume without loss of generality that $t_{\epsilon,d}(G)$ is unknown and we run Algorithm 1 with t such that $t_{\epsilon,d}(G) \leq t \leq t'_{\epsilon,d}(G)$. Recall that $n_t(G)$ is the maximum size of a distance- t neighborhood in G and that $h_s(G)$ is the maximum h such that G is an (s, h) -vertex expander. By Claim 1, $n_t(G) \geq \min\{(h_s(G))^t, s\}$. By the definition of $t'_{\epsilon,d}(G)$ we obtain that $(h_s(G))^{t'_{\epsilon,d}-1} \leq s$, since otherwise we get a contradiction to the minimality of $t'_{\epsilon,d}$. Thus, we obtain that $(h_s(G))^{t-1} \leq s$ and so $t \leq \frac{\log s}{\log h_s(G)} + 1$. On the other hand, since the degree is bounded by d , it holds that $n_t(G) \leq 1 + d^t$. By Theorem 5, the query complexity and the running time of the algorithm are $O(d \cdot n_t(G))$ and $O(d \cdot n_t(G) \cdot \log n)$, respectively. Hence, we get a bound of $O(d \cdot (1 + d^{\frac{\log s}{\log h_s(G)} + 1}))$ and $O(d \cdot (1 + d^{\frac{\log s}{\log h_s(G)} + 1}) \log n)$ on the query complexity and the running time of the algorithm, as desired. \square

5 The Algorithm for Graphs with an Excluded Minor

In this section we prove the following theorem. We say that a family of graphs is *minor-free* if the graphs in the family exclude some fixed minor.

Theorem 6 *There exists an LSSG algorithm for any family of minor-free graphs. The query complexity and running time of the algorithm are poly(h/ϵ), where h is the*

number of vertices in the excluded minor. The algorithm uses $\text{poly}(h/\epsilon) \cdot \log(n)$ random bits, and the stretch factor of the sparse spanning graph G' is $\tilde{O}(h/\epsilon)$.

For the sake of simplicity, we did not try to optimize the various complexity measures of the LSSG algorithm (query complexity, running time, and number of random bits it uses). Hence, the exponents in the polynomial expressions for these complexities are quite high, and can probably be improved.

In order to prove Theorem 6, we first design an LSSG algorithm for (minor-free) graphs whose maximum degree is upper bounded by a given integer d . Thereafter, we show how to adapt this algorithm to graphs with unbounded degree. We begin with proving the following theorem, where Algorithm 3 is provided in Sect. 5.4.

Theorem 7 *Algorithm 3 is an LSSG algorithm for the family of graphs that exclude a minor over h vertices. The query complexity of the algorithm is $\tilde{O}((h/\epsilon)^4 d^5)$ and its running time is $\tilde{O}((h/\epsilon)^5 d^5)$. The algorithm uses $\tilde{O}((h/\epsilon)d \log n)$ random bits, and the stretch factor of the sparse spanning graph G' is $\tilde{O}(h/\epsilon)$.*

In Sects. 5.1–5.4 we describe our LSSG algorithm (Algorithm 3) for graphs with degree bounded by d . In Sect. 5.5 we explain how to use this algorithm as a subroutine in an LSSG algorithm for graphs with unbounded degree. For this purpose we require that Algorithm 3 work also for graphs that are not connected. Namely, the guarantee for the algorithm is that every connected component in G remains connected in G' , and that $|E'| \leq (1 + \epsilon) \cdot n$. More precisely, since we are interested in bounding the total number of edges in the final sparse spanning subgraph by $(1 + \epsilon) \cdot n$, we shall actually ensure (with high probability) that Algorithm 3 answers positively on at most $(1 + \epsilon/4) \cdot n$ edges.

We start by describing a global partition of the vertices. We later explain how to locally generate this partition and design our algorithm (and the sparse subgraph it defines), based on this partition.

5.1 The Partition \mathcal{P}

The partition described in this subsection, denoted by \mathcal{P} , is a result of a random process. We describe how the partition is obtained in three steps where in each step we refine the partition from the previous step. The partition is defined based on two parameters: $\gamma \in (0, 1)$, and an integer $k > 1$, which are set subsequently (as a function of d , ϵ and h). Another central ingredient in the definition of the partition is a subset $W \subseteq V$ of vertices. This subset is selected randomly as follows: each vertex $v \in V$ draws a γ -biased bit, denoted b_v , and $v \in W$ if and only if $b_v = 1$. The joint-distribution of $(b_v)_{v \in V}$ is t -wise independent where $t \stackrel{\text{def}}{=} 2kd$. (The reason that the choice of W is determined by a t -wise independent distribution rather than an n -wise independent distribution is so as to bound the number of random bits used by the local emulation of the global algorithm.) In the next three Sects. (5.1.1–5.1.3) we explain how the partition \mathcal{P} is defined, given the choice of W .

5.1.1 First Step

We begin with some notation. For a vertex $v \in V$, let $cc(v)$ denote the subset of vertices in the connected component of v in G . Let Z denote the set of vertices v such that $|cc(v)| \geq k$. Given a vertex $v \in V$, we define the *center* of v with respect to W , denoted $\sigma(v)$ as follows.

1. If $v \notin Z$, or if $v \in Z$ and $cc(v) \cap W = \emptyset$, then $\sigma(v)$ is the vertex with the minimum id (identifier) in $cc(v)$.
2. If $v \in Z$ and $cc(v) \cap W \neq \emptyset$, then $\sigma(v)$ is the vertex in W that is closest to v , breaking ties using vertex ids. That is, $\sigma(v)$ is the vertex with the minimum identifier in the subset $\left\{y \in W : \Delta(y, v) = \min_{w \in W} \{\Delta(w, v)\}\right\}$.

For each $w \in W$ we define the *cell* of w with respect to W as $C(w) \stackrel{\text{def}}{=} \{v \in V : \sigma(v) = w\}$. Namely, the set of vertices in $C(w)$ consists of the vertices that are closer to w than to any other vertex in W (where ties are broken according to the ids of the vertices). For every vertex $w \in V$ such that w has the minimum identifier in $cc(w)$, and $cc(w) \cap W = \emptyset$, $C(w)$ is defined the same way (namely, as the subset of all vertices v for which $\sigma(v) = w$). Notice that these cells form a partition of V .

5.1.2 Second Step

In this step we identify a subset of *special* vertices, which we call the *remote vertices* and make these vertices singletons in the partition \mathcal{P} . The set of remote vertices, R , which is a subset of Z , is defined with respect to W and the integer parameter k as described next. For every vertex $v \in Z$, let $\ell_k(v)$ be the minimum integer ℓ such that the BFS tree rooted at v of depth ℓ has size at least k . Let $B_k(v)$ be the set of vertices in the BFS tree rooted at v of depth $\ell_k(v)$. We define $R = \{v \in Z : B_k(v) \cap W = \emptyset\}$, i.e., those vertices v for which $B_k(v)$ does not contain a vertex (center) in W . Clearly, a vertex can identify efficiently if it is in R by probing at most kd vertices and checking whether they intersect W . In Sect. 5.3 we obtain a bound on the size of R .

5.1.3 Third Step

In this step we decompose cells that are too big. We first argue that the cells are still connected (after the removal of the vertices in R from all original cells defined in Sect. 5.1.1). Thereafter we will use a procedure of tree decomposition in order to break the cells into smaller parts.

Claim 4 *For every $w \in W$, the subgraph induced by $C(w) \setminus R$ is connected. Furthermore, for every $v \in C(w) \setminus R$, the subgraph induced by $C(w) \setminus R$ contains all vertices that belong to the shortest paths between v and w .*

Proof Fix $w \in W$, and consider any $v \in C(w) \setminus R$. We will prove that the subgraph induced by $C(w) \setminus R$ contains all vertices on the shortest paths between v and w and this will imply the connectivity as well. The proof is by induction on the distance to

w . In the base case $v = w$. In this case $C(w) \setminus R$ clearly contains a path between v to itself because it contains v . Otherwise, we shall show that for any $u \in N(v)$ for which $\Delta(u, w) < \Delta(v, w)$ it holds that $u \in C(w) \setminus R$. The proof will then follow by induction. Let P be a shortest path between v and w and let $(v, u) \in E$ denote the first edge in P . We first observe that $\sigma(u) = w$ and thus $u \in C(w)$. Assume otherwise and conclude that there is a vertex in $w' \in W$ for which either $\Delta(v, w') < \Delta(v, w)$ or $\Delta(v, w) = \Delta(v, w')$ and $id(w') < id(w)$, in contradiction to the fact that $\sigma(v) = w$ (see the definition of $\sigma(\cdot)$ in Sect. 5.1.1). Since u is on a shortest path between v and w it follows that

$$\Gamma_{\Delta(u,w)-1}(u) \subseteq \Gamma_{\Delta(v,w)-1}(v). \quad (3)$$

From the fact that $v \notin R$ it follows that $|\Gamma_{\Delta(v,w)-1}(v)| \leq k$ and hence from Equation (3) it follows that $|\Gamma_{\Delta(u,w)-1}(u)| \leq k$ and so $u \notin R$ as well. We conclude that $u \in C(w) \setminus R$ and $\Delta(u, w) = \Delta(v, w) - 1$ as desired. \square

We shall use the following notation. For each $w \in W$ let $\mathcal{T}(w)$ denote the BFS tree rooted at w of the subgraph induced by $C(w) \setminus R$ (recall that the BFS is performed by exploring the vertices according to the order of their identifiers). To obtain the final refinement of our partition, for each $w \in W$ such that $|\mathcal{T}(w)| > k$, we run Algorithm 2 on $\mathcal{T}(w)$, w and k .

Algorithm 2 – Recursive Tree decomposition

Input: A tree \mathcal{T} , the root of the tree v and an integer k .

Output: A decomposition of \mathcal{T} into subtrees, where each subtree is assigned a (sub-)center.

1. Initialize the set of vertices of the current part $Q := \emptyset$.
 2. Perform a BFS starting from v and stop at level $\ell \stackrel{\text{def}}{=} \ell_k(v)$ (see the definition of $\ell_k(\cdot)$ in the Second Step). Add to Q all the vertices explored in the BFS.
 3. Let S denote the set of all the children of the vertices in the ℓ^{th} level of the BFS (namely, all the vertices in level $\ell + 1$).
 4. For each vertex $u \in S$:
 - (a) If the subtree rooted at u , \mathcal{T}_u , has size at least k , then disconnect this subtree from \mathcal{T} and continue to decompose by recursing on input \mathcal{T}_u , u and k .
 - (b) Otherwise, add the vertices of \mathcal{T}_u to Q .
 5. Set v to be the *sub-center* of all the vertices in Q .
-

5.2 The Edge Set

Given the partition \mathcal{P} defined in the previous Sect. 5.1, we define the edge set, E' , of our sparse spanning graph in the following simple manner. In each part of \mathcal{P} that is not a singleton, we take a spanning tree. Specifically, we take the BFS-tree rooted at the sub-center of that part (see Algorithm 2, Step 5). For every pair of parts of $X, Y \in \mathcal{P}$, if the set $E(X, Y) \stackrel{\text{def}}{=} \{(x, y) \in E : x \in X \text{ and } y \in Y\}$ is not empty, then we add to E' the edge $e \in E(X, Y)$ with minimal ranking (where the ranking over edges is as

defined in Sect. 2). Clearly every connected component of G is spanned by E' . We would like to bound the size of E' . To this end we will use the following claim.

Claim 5 *Let \mathcal{P}' denote the partition \mathcal{P} when restricted to Z . The number of parts in \mathcal{P}' is bounded as follows: $|\mathcal{P}'| \leq |W| + |R| + \frac{n}{k}$.*

Proof Consider the three steps of refinement of the partition. Clearly, after the first step the number of cells for which some $w \in W$ is the center is at most $|W|$. Therefore, after the second step, the size of the partition when restricted to Z , is at most $|W| + |R|$. Finally, since in the last step we break only parts whose size is greater than k into smaller parts that are of size at least k , we obtain that the number of new parts that are introduced in this step is at most n/k . The claim follows. \square

The next lemma establishes the connection between the size of \mathcal{P}' and the sparsity of G' .

Claim 6 *For an input graph G that is H -minor free for a graph H over h vertices,*

$$|E'| < n + |\mathcal{P}'| \cdot c(h) ,$$

where $c(h)$ is a bound on the arboricity of G (as defined in Theorem 1) and \mathcal{P}' is as defined in Claim 5.

Proof Since for each $X \in \mathcal{P}$ the subgraph induced by X is connected, we can contract each part in \mathcal{P} and obtain an H -minor free graph. The number of vertices in this graph is $|\mathcal{P}|$. If we replace multi-edges with single edges, then by the bound on the number of edges of H -minor free graph (see Theorem 1), we obtain that the number of edges in this graph is at most $|\mathcal{P}'| \cdot c(h)$ (observe that every part that is in \mathcal{P} but not in \mathcal{P}' becomes an isolated vertex in the contracted graph). Finally, since the total number of edges in the union of spanning trees of each part is $n - |\mathcal{P}| < n$, we obtain the desired result. \square

5.3 Bounding the Number of Remote Vertices

In this subsection we prove the following lemma.

Lemma 7 *If $k = \Omega((\log^2(1/\gamma) + \log d)/\gamma)$, then with probability at least $1 - \frac{1}{\Omega(n)}$ over the choice of W , it holds that $|R| \leq \gamma n$.*

We note that by Lemma 7 it follows that there exists a distributed algorithm in the CONGEST model for the LSSG problem with round complexity $O(k)$ which proceeds as follows. At the first step the set W is selected distributively. Afterwards, each vertex in W sends a unique token which is forwarded to neighbors by all vertices that did not receive a token so far. This goes on for k steps. By the end of this process each vertex in the network knows to which part it belongs to (if it did not receive a token then it is in R which means it is a singleton).

On this note, we also remark that the partition-oracle of Hassidim et al. [21], which can be used to obtain an LSSG algorithm, can be implemented in the LOCAL model in time which is also $\text{poly}(\epsilon^{-1})$.

In order to establish Lemma 7 we start by defining for every $v \in Z$,

$$Y_k(v) \stackrel{\text{def}}{=} \{u \in V : v \in B_k(u)\} . \tag{4}$$

Informally, $Y_k(v)$ is the set of vertices that encounter v while performing a BFS that stops after the first level in which the total number of explored vertices is at least k . We first establish the following simple claim.

Claim 8 *For every vertex $u \in Y_k(v)$ and for every vertex w that is on a shortest path between u and v , we have that $w \in Y_k(v)$.*

Proof Let $\ell = \Delta(u, v)$ and $\ell' = \Delta(w, v)$, so that $\Delta(u, w) = \ell - \ell' \geq 1$. Assume, contrary to the claim, that $w \notin Y_k(v)$. This implies that $|\Gamma_{\ell'-1}(w)| \geq k$. But since $\Gamma_{\ell'-1}(w) \subseteq \Gamma_{\ell-1}(u)$, we get that $|\Gamma_{\ell-1}(u)| \geq k$, contrary to the premise of the claim that $u \in Y_k(v)$. □

We now turn to upper bound the size of $Y_k(v)$.

Lemma 9 *For every graph $G = (V, E)$ with degree bounded by d , and for every $v \in Z$,*

$$|Y_k(v)| \leq d^3 \cdot k^{\log k+1} .$$

Proof Fix a vertex $v \in Z$. For every $0 \leq j \leq k$, define $Y_k^j(v) \stackrel{\text{def}}{=} \{u \in Y_k(v) : \Delta(v, u) = j\}$. Observe that $Y_k(v) = \bigcup_{j=0}^k Y_k^j(v)$. Therefore, if we bound the size of $Y_k^j(v)$, for every $0 \leq j \leq k$, we will get a bound on the size of $Y_k(v)$. Consider first any $3 \leq j < k$ and any vertex $u \in Y_k^j(v)$. Recall that $\ell_k(u)$ is the minimum integer ℓ such that the BFS tree rooted at v of depth ℓ has size at least k . Since $j \leq \ell_k(u)$, it follows that $|\Gamma_{j-1}(u)| < k$. Now consider a shortest path between u and v and let w be the vertex on this path for which $\Delta(u, w) = \lfloor (j - 1)/2 \rfloor$. Denote $q \stackrel{\text{def}}{=} \Delta(w, v)$. By Claim 8, $w \in Y_k(v)$, and by the definition of q , $w \in Y_k^q(v)$. Therefore,

$$|\Gamma_{q-1}(w)| \leq k . \tag{5}$$

From the fact that w is on the shortest path between u and v it also follows that

$$\begin{aligned} q &= \Delta(v, u) - \Delta(u, w) = j - \lfloor (j - 1)/2 \rfloor \\ &= \lceil (j - 1)/2 \rceil + 1 \\ &\geq \lfloor (j - 1)/2 \rfloor + 1 = \Delta(u, w) + 1 . \end{aligned} \tag{6}$$

Therefore $q - 1 \geq \Delta(u, w)$ and so $u \in \Gamma_{q-1}(w)$. It follows that

$$Y_k^j(v) \subseteq \bigcup_{w \in Y_k^q(v)} \Gamma_{q-1}(w) . \tag{7}$$

From Eqs. (5) and (7) we get that $|Y_k^j(v)| \leq k \cdot |Y_k^q(v)|$. For every $j \leq 3$ we have the trivial bound that $|Y_k^j(v)| \leq d^3$. By combining with Equation (6) we get that $|Y_k^j(v)| \leq d^3 \cdot k^{\log j}$. Since $Y_k(v) = \bigcup_{j=0}^k Y_k^j(v)$ we obtain the desired bound. \square

While the bound on $|Y_k(v)|$ in Lemma 9 may not be tight, it suffices for our purposes. One might conjecture that it is possible to prove a polynomial bound (in k and d). We show that this is not the case (see Lemma 12 in the appendix).

We now use the bound in Lemma 9 in order to bound the number of remote vertices.

Proof of Lemma 7 Let χ_w denote the characteristic vector of the set W . For a subset $S \subseteq V$, let $\chi_w(S)$ denote the projection of χ_w onto S . That is, $\chi_w(S)$ is a vector of length $|S|$ indicating for each $x \in S$ whether $\chi_w(x) = 1$.

For each vertex $v \in V$ define a random variable ζ_v indicating whether it is a remote vertex with respect to W . Recall that v is remote if and only if $v \in Z$ and $B_k(v) \cap W = \emptyset$. Recalling that W is selected according to a t -wise independent distribution where $t = 2kd$ and that $k \leq |B_k(v)| < k \cdot d$, we get that $\Pr[\zeta_v = 1] \leq (1 - \gamma)^k$. We also define $S_v \stackrel{\text{def}}{=} \{u \in Z : B_k(u) \cap B_k(v) \neq \emptyset\}$. Let $v \in V$. If $v \notin Z$ then the value of ζ_v is fixed and equals 0. Otherwise, observe that the value of ζ_v is determined by $\chi_w(B_k(v))$. Furthermore, since for every $v \in Z$ and $u \in Z \setminus S_v$, $\chi_w(B_k(v))$ and $\chi_w(B_k(u))$ are independent it follows that ζ_u and ζ_v are independent as well. Hence, in this case $\text{Cov}[\zeta_v, \zeta_u] = 0$, and we obtain the following upper bound on the variance of the number of remote vertices.

$$\begin{aligned} \text{Var} \left[\sum_{v \in V} \zeta_v \right] &= \sum_{(v,u) \in V} \text{Cov}[\zeta_v, \zeta_u] = \sum_{v \in V} \sum_{u \in S_v} (\text{Exp}[\zeta_v \cdot \zeta_u] - \text{Exp}[\zeta_u] \cdot \text{Exp}[\zeta_v]) \\ &\leq \sum_{v \in Z} \sum_{u \in S_v} \text{Exp}[\zeta_v \cdot \zeta_u | \zeta_v = 1] \cdot \Pr[\zeta_v = 1] \leq \sum_{v \in Z} |S_v| \cdot (1 - \gamma)^k. \end{aligned} \tag{8}$$

By the definition of $Y_k(\cdot)$ in Equation (4) it follows that $S_v \subseteq \bigcup_{u \in B_k(v)} Y_k(u)$. By Lemma 9, $\max_{v \in V} \{|Y_k(v)|\} \leq d^3 \cdot k^{\log k+1}$. Therefore

$$|S_v| \leq |B_k(v)| \cdot d^3 \cdot k^{\log k+1} \leq d^4 \cdot k^{\log k+2}. \tag{9}$$

Hence, by Equations (8) and (9) we get $\text{Var} \left[\sum_{v \in V} \zeta_v \right] \leq nd^4 \cdot k^{\log k+2} \cdot (1 - \gamma)^k$. Since $(1 - \gamma)^k \leq e^{-\gamma k}$ we obtain that $\text{Var} \left[\sum_{v \in V} \zeta_v \right] \leq \gamma^2 n$ for $k = \Omega((\log^2(1/\gamma) + \log d)/\gamma)$. Since for every $v \in V$, $\Pr[\zeta_v = 1] \leq (1 - \gamma)^k \leq \gamma$, we get that $\text{Exp} \left[\sum_{v \in V} \zeta_v \right] \leq \gamma n/2$. By applying Chebyshev’s inequality we get that

$$\Pr \left[\sum_{v \in V} \zeta_v \geq \text{Exp} \left[\sum_{v \in V} \zeta_v \right] + \gamma n/2 \right] \leq \frac{4 \text{Var} \left[\sum_{v \in V} \zeta_v \right]}{\gamma^2 n^2} \leq \frac{4}{n}.$$

Since $|R| = \sum_{v \in V} \zeta_v$ it follows that $|R| < \gamma n$ with probability at least $1 - (4/n)$, as desired. \square

5.4 The Local Algorithm

In this subsection we provide Algorithm 3, which on input $e \in E$, locally decides whether $e \in E'$, as defined in Sect. 5.2, based on the (random, but not completely independent) choice of W . Given an edge $e = (u, v)$ we distinguish between two cases. In the first case, $v \notin Z$. In this case, the center of u is identical to the center of v and can be found by performing a BFS from v until the entire connected component $cc(v)$ is revealed. This is done in Step 1 of Algorithm 4. In this case, (u, v) belongs to E' if and only if it belongs to the BFS tree rooted at $\sigma(v)$.

In the second case, $v \in Z$. In this case, the algorithm first finds, for each $y \in \{u, v\}$, the center, $\sigma(y)$, that y is assigned to by the initial partition, under the condition that $\sigma(y) \in B_k(y)$. This is done in Step 2 of Algorithm 4, which simply performs a BFS starting from y until it encounters a vertex in W , or it reaches level $\ell_k(y)$ without encountering such a vertex (in which case y is a remote vertex). Algorithm 4 assumes access to χ_w , which is implemented using the random seed that defines the t -wise independent distribution, and hence determines W . If y is not found to be a remote vertex, then Algorithm 3 next determines to which sub-part of $C(\sigma(y)) \setminus R$ does y belong. This is done by emulating the tree decomposition of the BFS tree rooted at $\sigma(y)$ and induced by $C(\sigma(y)) \setminus R$. A central procedure that is employed in this process is Algorithm 5. This algorithm is given a vertex $v \in W$, and a vertex u in the BFS subtree rooted at v and induced by $C(v) \setminus R$. It returns the edges going from u to its children in this tree, by performing calls to Algorithm 4.

Algorithm 3 – Dense-Centers Algorithm

Input: An edge $(u, v) \in E$.

Output: YES if $(u, v) \in E'$ and NO otherwise.

1. For each $y \in \{u, v\}$ find the part that y belongs to as follows:
 - (a) Use Algorithm 4 to obtain $\sigma(y)$.
 - (b) If $\sigma(y)$ is 'null', then the part that y belongs to is the singleton set $\{y\}$.
 - (c) Otherwise, let \mathcal{T} denote the BFS tree rooted at $\sigma(y)$ in the subgraph induced by $C(\sigma(y)) \setminus R$. By Claim 4 every shortest path between $\sigma(y)$ and $v \in C(\sigma(y)) \setminus R$ is contained in the subgraph induced on $C(\sigma(y)) \setminus R$. Therefore the edges of \mathcal{T} can be explored via Algorithm 5. Reveal the part (the subset of vertices) that y belongs to in \mathcal{P} (as defined in Sect. 5.1). Recall that the part of y is the subtree of \mathcal{T} that contains y after running Algorithm 2 on input \mathcal{T} , $\sigma(y)$ and k . This part can be revealed locally as follows.
 - i. Reveal the path between $\sigma(y)$ and y in \mathcal{T} , denoted by $P(y)$. Since $P(y)$ is the shortest path between y and $\sigma(y)$ with the lexicographically smallest order, it can be revealed by performing a BFS from y until $\sigma(y)$ is encountered.
 - ii. Run Algorithm 2 on \mathcal{T} while recursing in Step 4a only on the subtrees in which the root is contained in $P(y)$.
 2. If u and v are in the same part, then return YES iff the edge (u, v) belongs to the BFS tree of that part.
 3. Otherwise, return YES iff the edge (u, v) is the edge with minimum rank connecting the two parts.
-

Proof of Theorem 7 Let $c(h)$ as defined in Theorem 1 (namely, a bound on the arboricity of H -minor free graphs). We set $\gamma = \epsilon/(16c(h))$ and $k = c' \cdot (\log^2(1/\gamma) + \log d)/\gamma$, where c' is a sufficiently large constant. We start by bounding the size of W (with

Algorithm 4 – Find Center

Input: A vertex v and an integer k . Query access to χ_w .

Output: $\sigma(v)$ if $\sigma(v) \in B_k(v)$ or $|\text{cc}(v)| < k$, otherwise return ‘null’.

1. Perform a BFS from v until at most k vertices are revealed. If the entire connected component of v is revealed, then return the vertex with the minimum id from $\text{cc}(v)$.
2. Otherwise, perform a BFS from v until the first level that contains a vertex in W or until at least k vertices are reached. That is, defining $W_j \stackrel{\text{def}}{=} \Gamma_j(v) \cap W$, stop at level ℓ where $\ell \stackrel{\text{def}}{=} \min\{\ell_k(v), \min_j\{W_j \neq \emptyset\}\}$.
3. If $W_\ell = \emptyset$ then return ‘null’ (v is remote).
4. Otherwise, return the vertex with the minimum id from W_ℓ .

high probability). Recall that W is selected according to t -wise independent random variable over $\{0, 1\}^n$ where $\Pr[\chi_w(i) = 1] = \gamma$ for each $i \in [n]$ (and recall that $t = 2kd$). By the definition of W we have that $\text{Exp}[|W|] = \text{Exp}[\sum_{i \in [n]} \chi_w(i)] = \gamma n$. Since for every $1 \leq i < j \leq n$, $\chi_w(i)$ and $\chi_w(j)$ are pairwise-independent, we obtain that

$$\text{Var} \left[\sum_{i \in [n]} \chi_w(i) \right] = \sum_{i \in [n]} \text{Var} [\chi_w(i)] = \sum_{i \in [n]} (\text{Exp} [\chi_w^2(i)] - \text{Exp} [\chi_w(i)]^2) = n\gamma(1 - \gamma) .$$

Therefore, by Chebyshev’s inequality $\Pr [|W| \geq 2\gamma n] \leq \frac{1-\gamma}{\gamma n}$. By Lemma 7 (and the setting of k), with probability $1 - 1/\Omega(n)$, $|R| \leq \gamma n$. By Claims 5 and 6 and the settings of γ and k , we get that

$$|E'| \leq n + (|W| + |R| + n/k) \cdot c(h) \leq n + 4\gamma \cdot c(h) \leq (1 + \epsilon/4)n$$

with probability $1 - 1/\Omega(n)$.

The connectivity in G' of each connected component of G follows by construction of E' . The claim about the stretch of G' follows from the fact that the diameter of every part of \mathcal{P} is bounded by $2k$.

The number of vertices that Algorithm 4 traverses while performing a BFS (for any vertex v it is called with) is at most $(k - 1)d < kd$. To verify this, observe that the algorithm continues to explore another level of the BFS only if strictly less than k vertices were encountered so far. Since the degree of each vertex is bounded by d , the query complexity of performing the BFS is bounded by kd^2 . Algorithm 4 makes at most $O(kd)$ accesses to χ_w . By Theorem 2, each access to χ_w takes $O(kd)$ time (as the distribution of χ_w is $O(kd)$ -wise independent). Hence, the running time of the algorithm is bounded by $O(k^2d^2)$. The query complexity of Algorithm 5 is upper bounded by the total query complexity of at most d calls to Algorithm 4, plus d executions of a BFS until at most kd vertices are encountered (Steps 2b and 2c, which for simplicity we account for separately from Algorithm 4). A similar bound holds for the running time. Hence, the total query complexity and running time of Algorithm 5 are $O(kd^3)$ and $O(k^2d^3)$, respectively.

The size of any subtree returned by Algorithm 2 is upper bounded by k^2d^2 . To verify this, recall that at the end of Step 2 of Algorithm 2, at most kd vertices were

explored. Hence, the number of vertices that are incident to the explored vertices is at most kd^2 . Thus, due to Step 4a, the total number of vertices in each part is at most k^2d^2 . Recall that the distance of every vertex in \mathcal{T} from the root v is at most k . Hence, Step 4a can be implemented locally by calling Algorithm 5 at most k times. Thus, we obtain that the total query complexity and running time of constructing each part is at most $O(k^3d^5)$ and $O(k^4d^5)$, respectively. The query complexity and running time of Algorithm 3 are dominated by those of Step 1c. Observe that in Step 1c at most $|P(y)|$ parts are constructed. Since $|P(y)| \leq k$, we obtain that the overall query complexity and running time of Algorithm 3 are bounded by $O(k^4d^5)$ and $O(k^5d^5)$, respectively. By the setting of k we obtain the final result. \square

Algorithm 5 - Get BFS outgoing-edges endpoints

Input: $v \in W$ and a vertex $u \in \mathcal{T}(v)$ where $\mathcal{T}(v)$ denotes the BFS tree induced by $C(v) \setminus R$ and rooted at v .

Output: The outgoing edges from u in $\mathcal{T}(v)$ (the orientation of the edges is from the root to the leaves).

1. Initialize $S := \{v\}$
 2. For each $u' \in N(u)$, if the following three conditions hold, then add u' to S :
 - (a) Algorithm 4 on input u' returns v . Namely, $\sigma(u') = v$.
 - (b) u is on a shortest path between u' and v . Namely, $\Delta(u', v) = \Delta(u, v) + 1$.
 - (c) u is the vertex with the minimum id among all vertices in $\{w \in N(u') : \Delta(w, v) = \Delta(u', v) - 1\}$.
 3. Return S .
-

5.5 Removing the Dependence on the Maximum Degree

In this subsection we show how to adapt Algorithm 3 to work with unbounded degrees. Let $c(h)$ be a bound on the arboricity of H -minor free graphs (as defined in Theorem 1), where h denotes the number of vertices in the excluded minor H . We say that a vertex is *heavy* if its degree is greater than $\tilde{d} \stackrel{\text{def}}{=} 8(c(h))^2/\epsilon$. Let \mathcal{H} denote the set of *heavy* vertices and let $\mathcal{L} \stackrel{\text{def}}{=} V \setminus \mathcal{H}$ denote the set of *light* vertices. As we explain next, we deal separately (and differently) with edges in $\mathcal{H} \times \mathcal{H}$, $\mathcal{L} \times \mathcal{L}$, and $\mathcal{H} \times \mathcal{L}$.

5.5.1 The Edges Between Pairs of Heavy Vertices

We add all the edges in $G[\mathcal{H}]$, namely, all the edges between pairs of heavy vertices, to the spanning graph. Since $G[\mathcal{H}]$ excludes H as a minor, by Theorem 1, the number of edges in $G[\mathcal{H}]$ is at most $c(h) \cdot |\mathcal{H}|$. On the other hand, the size of \mathcal{H} is at most $\epsilon n/(4c(h))$ (otherwise, we obtain a contradiction to the size of E). Hence, $G[\mathcal{H}]$ has at most $(\epsilon/4)n$ edges.

5.5.2 The Edges Between Pairs of Light Vertices

Consider $G[\mathcal{L}]$, namely, the subgraph that remains after removing the vertices in \mathcal{H} (and the edges incident to them). By its definition, the maximum degree in $G[\mathcal{L}]$ is

at most \tilde{d} . For an edge (u, v) such that $u, v \in \mathcal{L}$, consider emulating the execution of Algorithm 3 on $G[\mathcal{L}]$ given the query (u, v) . Observe that such an emulation can be easily implemented on G at a multiplicative cost of at most \tilde{d} . The edge (u, v) belongs to the spanning graph if and only if the emulation of Algorithm 3 returns YES on query (u, v) .

5.5.3 The Edges Between Heavy and Light Vertices

For every light vertex v , let $\mathcal{P}(v)$ denote the part that v belongs to in the partition \mathcal{P} of $G[\mathcal{L}]$ as computed in Step 1c of Algorithm 3. If there is at least one edge (u, v') between a heavy vertex u and a (light) vertex $v' \in \mathcal{P}(v)$, then we assign $\mathcal{P}(v)$ to the adjacent heavy vertex with the minimum id. For a heavy vertex u , the cluster of u , denoted by $\mathcal{C}(u)$, includes u and all the vertices in the parts that are assigned to u . We refer to u as the leader of the vertices in $\mathcal{C}(u)$.

On query (u, v) where u is heavy and v is light, the query is answered as described in Algorithm 6. We next discuss how the algorithm works, and give some high-level intuition for its correctness (a formal proof appears in Sect. 5.5.4).

If v belongs to the cluster of u , i.e., $v \in \mathcal{C}(u)$, then the answer to the query (u, v) is YES if and only if v has the minimum id among the vertices in $\mathcal{P}(v)$ that are adjacent to u . Otherwise (v does not belong to $\mathcal{C}(u)$), let w denote the leader of v (so that $w \neq u$). Consider all the edges in $\mathcal{H} \times \mathcal{L}$ between $\mathcal{C}(u)$ and $\mathcal{C}(w)$, whose set we denote by $\tilde{E}(u, w)$, so that $(u, v) \in \tilde{E}(u, w)$. Observe that all edges in $\tilde{E}(u, w)$ are incident to either u or w . We would have liked to include in E' only a single edge from $\tilde{E}(u, w)$, say the edge with minimum rank (recall that we consider a ranking over the edges based on the ids of their endpoints). However, u and w may have a high degree, so that we cannot afford querying all their incident edges in order to find the one with minimum rank in $\tilde{E}(u, w)$ (and answer the query on (u, v) accordingly).

In order to save in the number of queries, we allow for more than a single edge in $\tilde{E}(u, w)$ to be included in E' . Suppose we had access to uniformly selected edges in $\tilde{E}(u, w)$ (at unit cost per edge). Then we could take a sample of such edges and answer YES on (u, v) if and only if no edge in the sample has rank smaller than (u, v) . Roughly speaking, this would ensure (with high probability) that only edges in $\tilde{E}(u, w)$ with relatively low rank will be included in E' .

Since we do not have direct access to uniformly selected edges in $\tilde{E}(u, w)$ we do the following. We take a random sample of edges incident to u and a random sample of edges incident to w . For each selected edge (u, y) (and analogously for selected edges (w, y)), we check what vertex is the leader of y . If the leader of y is w , then we obtained an edge (u, y) in $\tilde{E}(u, w)$. If the leader of y is u , then we find all vertices in $\mathcal{P}(y)$, and if some $z \in \mathcal{P}(y)$ is a neighbor of w (recall that z has low degree), then we obtained an edge (z, w) in $\tilde{E}(u, w)$. (If the leader of y is neither w nor u , then the sampled edge is not useful for us.) If we obtained an edge in $\tilde{E}(u, w)$ that has rank smaller than (u, v) , then we answer NO on (u, v) , otherwise we answer YES.

Conditioned on an edge in $\tilde{E}(u, w)$ being obtained, while the distribution on these edges is not uniform, we can lower bound the probability that each edge is selected, which suffices for our purposes. If the probability that we obtain an edge in $\tilde{E}(u, w)$ is sufficiently large, then we are done (as we can upper bound the total fraction of

edges in $\tilde{E}(u, w)$ that are included in E'). Otherwise, we may include in E' all edges in $\tilde{E}(u, w)$, where we upper bound the total number of such edges by using the fact that G is H -minor free. For further details, see the proof of Theorem 6.

For each queried edge (u, v) (where u is heavy, v is light, and $v \in \mathcal{C}(w)$ for $w \neq u$), the sampled edges incident to u and w are distributed uniformly (among the edges incident to u and w , respectively). However, similarly to the choice of the subset W in the definition of the partition \mathcal{P} (see Sect. 5.1), in order to save in randomness, the selection of edges that is performed for different queries is not done independently. Rather, the samples are all determined by a random seed of $\text{poly}(h/\epsilon) \log(n)$ bits.

Algorithm 6 - Heavy-Light Edges

Input: query (u, v) where u is heavy and v is light.

Output: YES if $(u, v) \in E'$ and NO otherwise.

1. Find the part of v , $\mathcal{P}(v)$ (as described in Step 1 of Algorithm 3). Find the leader of v , denoted w (by querying all edges incident to vertices in $\mathcal{P}(v)$).
 2. If $u = w$, then return YES if and only if v has minimum id among all vertices in $\mathcal{P}(v)$ that are adjacent to u .
 3. Sample a set X_u of $q = \tilde{\Theta}((c^3(h)k^2\tilde{d}^3)/\epsilon^3)$ vertices that are adjacent to u , where k is as defined in the proof of Theorem 7, and similarly sample a set X_w of q vertices that are adjacent to w .
 4. For each light vertex y in X_u :
 - (a) If its leader is w and $\rho(u, y) < \rho(u, v)$, then return NO.
 - (b) If its leader is u and there is an edge (z, w) such that $\rho(z, w) < \rho(u, v)$ and $z \in \mathcal{P}(y)$, then return NO.
 5. For each light vertex y in X_w :
 - (a) If its leader is u and $\rho(w, y) < \rho(u, v)$, then return NO.
 - (b) If its leader is w and there is an edge (z, u) such that $\rho(z, u) < \rho(u, v)$ and $z \in \mathcal{P}(y)$, then return NO.
 6. Return YES.
-

5.5.4 Correctness of the Algorithm—Proof of Theorem 6

The connectivity of G' follows directly from the construction of E' . Namely: (1) Pairs of vertices that belong to the same connected components of $G[\mathcal{L}]$ are connected in G' by the correctness of Algorithm 3. (2) Pairs of vertices that belong to the same cluster are connected either within their part in the partition \mathcal{P} of $G[\mathcal{L}]$ or via their leader in the cluster. (3) Pairs of vertices in different clusters (that are not in the same connected components of $G[\mathcal{L}]$) are connected since there is at least one edge in E' between every pair of clusters for which such an edge exists.

By construction, the stretch factor is at most a constant factor greater than the stretch factor of Algorithm 3. To verify this, observe that the diameter of each cluster is bounded by $2k + 2$ where k is a bound on the stretch factor of Algorithm 3. Since for every pair of adjacent clusters there exists at least one edge in G' which connects these clusters, we obtain that for each edge we remove from $(u, v) \in \mathcal{H} \times \mathcal{L}$ there exists in G' a path of length at most $2k + 2$ from u to v .

We now turn to bound the total number of edges in E' . Let F denote the set of edges in $E' \cap (\mathcal{H} \times \mathcal{L})$ that are incident to two different clusters (namely, each endpoint belongs to a different cluster).

Claim 10 *With probability $1 - 1/\Omega(1/n)$, $|E' \setminus F| \leq (1 + (\epsilon/2))n$.*

Proof For a subset of vertices U , let $E'(U)$ denote the subset of edges in E' between vertices in U . For each part X in the partition \mathcal{P} of $G[\mathcal{L}]$ we add to E' a spanning tree and so we have that $|E'(X)| = |X| - 1$. For a heavy vertex u , let X_1, \dots, X_s be the parts assigned to it (so that $\mathcal{C}(u) = \{u\} \cup \bigcup_{j=1}^s X_j$). Since each part connects to its leader by a single edge (in G'), it follows that $|\bigcup_{j=1}^s E'(\{u\} \cup X_j)| = |\mathcal{C}(u)| - 1$. As for the rest of the edges in $(\mathcal{L} \times \mathcal{L}) \cap E'$, by the proof of Theorem 7, with probability $1 - 1/\Omega(n)$ (over the choice of W), their number is upper bounded by $(\epsilon/4)n$. Since the number of edges in $\mathcal{H} \times \mathcal{H}$ is at most $(\epsilon/4)n$ (see details in Sect. 5.5.1), the claim follows. \square

Claim 11 *With probability $1 - 1/\Omega(n)$, $|F| \leq (\epsilon/2)n$.*

Proof For each cluster B , we charge to B a subset of the edges incident to B so that the union of all the charged edges (over all clusters) contains F . Our goal is to show that with probability $1 - 1/\Omega(n)$, the total number of charged edges is $(\epsilon/2)n$.

Consider the auxiliary graph, denoted \tilde{G} , that results from contracting each cluster B in G into a *mega-vertex* in \tilde{G} , which we denote by $v(B)$. For each pair of clusters B and B' such that $E(B, B')$ is non-empty, there is an edge $(v(B), v(B'))$ in \tilde{G} , which we refer to as a *mega edge*, and whose weight is $|E(B, B')|$. Since G is H -minor free, so is \tilde{G} . By Theorem 1, which bounds the arboricity of H -minor free graphs, we can partition the mega-edges of \tilde{G} into $c(h)$ forests. Consider orienting the mega-edges of \tilde{G} according to this partition (from children to parents in the trees of these forests), so that each mega-vertex has at most $c(h)$ outgoing mega-edges. For cluster B and a cluster B' such that $(v(B), v(B'))$ is an edge in \tilde{G} that is oriented from $v(B)$ to $v(B')$, we shall charge to B a subset of the edges in $E(B, B')$, as described next.

Recall first that each part in the partition \mathcal{P} of $G[\mathcal{L}]$ has size at most $k^2 \tilde{d}^2$. Let $E^b(B, B')$ denote the subset of edges in $E(B, B')$ whose rank is in the bottom $(\epsilon/(8c(h))) \cdot |E(B, B')|$ edges of $E(B, B')$. We charge all the edges in $E^b(B, B')$ to B . The rationale is that for these edges it is likely that the algorithm won't sample an edge in $E(B, B')$ with lower rank. The total number of such edges is at most $(\epsilon/(8c(h))) \cdot |E| \leq (\epsilon/8)n$. Let u be the leader of the cluster B , and let $N^b(u, B')$ be the set of vertices, $y \in N(u)$, such that:

$$\left(y \in B' \text{ and } (u, y) \in E^b(B, B') \right) \text{ or } \left(\exists (y', z) \in E^b(B, B') \text{ s.t. } y' \in \mathcal{P}(y) \right) .$$

That is, $N^b(u, B')$ is the subset of neighbors of u such that if Algorithm 3 selects one of them in Step 3, then it obtains an edge in $E^b(B, B')$. We consider two cases.

First case $|N^b(u, B')|/|N(u)| < \epsilon^2/(2^7 c^2(h)k^2 \tilde{d}^3)$. In this case we charge all edges in $E(B, B')$ to B . For each part $\mathcal{P}(y)$ such that $y \in N^b(u, B')$ there are at most $|\mathcal{P}(y)| \cdot \tilde{d}$ edges $(y', z) \in E^b(B, B')$ for $y' \in \mathcal{P}(y)$. Therefore, in this case

$|E^b(B, B')| \leq k^2 \tilde{d}^3 \cdot |N^b(u, B')| \leq N(u) \cdot \epsilon^2 / (2^7 c^2(h))$ and hence $|E(B, B')| < (8c(h)/\epsilon) \cdot \epsilon^2 / (2^7 c^2(h)) \cdot N(u)$. It follows that the total number of charged edges of this type is at most $(\epsilon / (2^4 c(h))) \cdot 2|E| \leq (\epsilon/8)n$.

Second case $|N^b(u, B')|/|N(u)| \geq \epsilon^2 / (2^7 c^2(h)k^2 \tilde{d}^3)$. For each u and B' that fall under this case we define the set of edges $Y(u, B') = \{(u, v) : v \in N^b(u, B')\}$ and denote by Y the union of all such sets (over all such pairs u and B'). Edges in Y are charged to B if and only if they belong to F and are incident to a vertex in B . Fix an edge in Y that is incident to B , and note that the selection of neighbors of u is done according to a t -wise independent distribution for $t > 4q$, where q is the sample size set in Step 3 of the algorithm. Therefore, the probability that the edge belongs to F is upper bounded by $(1 - \epsilon^2 / (2^7 c^2(h)k^2 \tilde{d}^3))^q$, which by the setting of q , is at most $p = \epsilon / (8c(h))$.

We next show, using Chebyshev’s inequality, that with high probability, the number of edges in Y that are in F is at most $2p|E|$. For $y \in Y$, define J_y to be an indicator variable that is 1 if and only if $y \in F$. Then for a fixed $y \in Y$, $E[J_y] \leq p$ and $\{J_y\}$ are pairwise independent (this is due to the fact that the samples of every pair of edges are pairwise independent). Therefore, by Chebyshev’s inequality,

$$\begin{aligned} \Pr \left[\sum_{y \in Y} J_y \geq 2p|E| \right] &\leq \frac{\text{Var}[\sum_{y \in Y} J_y]}{(p|E|)^2} = \frac{\sum_{y \in Y} \text{Var}(J_y)}{(p|E|)^2} \leq \frac{p(1-p)|E|}{(p|E|)^2} \\ &= \frac{1-p}{p|E|} = \frac{1}{\Omega(n)}, \end{aligned}$$

and the proof of Claim 11 is completed. □

By combining Claims 10 and 11 we get that $|E'| \leq (1 + \epsilon) \cdot n$ with probability $1 - 1/\Omega(n)$, as required.

The random seed that Algorithm 6 uses consists of two parts. The first part is for running Algorithm 3, whose random seed has length $\tilde{O}((h/\epsilon)d \log n)$, and hence for $d = \tilde{d}$, its length is $\tilde{O}((h/\epsilon)^2 \log n)$. The second part is for selecting random neighbors in Step 3. Since for each edge in E we pick a random sample of $\tilde{\Theta}((c^3(h)k^2 \tilde{d}^3)/\epsilon^3) = \tilde{O}((h^{11}/\epsilon^8))$ neighbors independently, we obtain by Theorem 2 that a random seed of length $\tilde{O}(h^{11}/\epsilon^8) \log n$ is sufficient.

The running time of Algorithm 6 is upper bounded by a constant times the running time of Algorithm 3 times the number of edges sampled in Step 3, which by Theorem 7 and the setting of q in the algorithm is $\text{poly}(h/\epsilon)$. This completes the proof of Theorem 6.

Acknowledgements We would like to thank Oded Goldreich for his helpful suggestions. We would also like to thank the anonymous reviewers of this paper for their comments.

A Tables of Notations and Parameters

Table of notations and parameters for Sect. 4

Notation	Meaning
$\Gamma_r(v, G)$	The set of vertices at distance at most r from v (distance- r neighborhood of v).
$C_r(v, G)$	The subgraph of G induced by $\Gamma_r(v, G)$.
$n_r(G)$	$\max_{v \in V} \Gamma_r(v, G) $.
$h_s(G)$	Maximum h such that G is an (s, h) -vertex expander.
$s = s(n, \epsilon)$	$2\sqrt{2n/\epsilon} \ln n$.
$\ell = \ell(n, \epsilon)$	$\sqrt{\epsilon n/2}$ (number of centers).
L_j^i	Vertices in the i^{th} level of the BFS tree of the center σ_j .
$S(\sigma_j)$	Vertices that are assigned to the center σ_j .
$E'(\sigma)$	Edges of the BFS-tree that spans the subgraph induced on $S(\sigma)$.
$P(\sigma, \sigma')$	The shortest path between σ and σ' that has minimum lexicographic order among all shortest paths.
$e(\sigma, \sigma')$	If all vertices on $P(\sigma, \sigma')$ belong either to $S(\sigma)$ or to $S(\sigma')$, then $e(\sigma, \sigma')$ is the single edge $(x, y) \in P(\sigma, \sigma')$ such that $x \in S(\sigma)$ and $y \in S(\sigma')$.
$t_{\epsilon, d}(G)$	The minimum distance t ensuring that all but an $\epsilon/(2d)$ -fraction of the vertices have at least s vertices in their t -neighborhood.
$t'_{\epsilon, d}(G)$	The minimum distance t ensuring that all but an $\epsilon/(4d)$ -fraction of the vertices have at least s vertices in their t -neighborhood.
\mathcal{P}	The partition of the graph which the spanning graph construction is based on.
γ	The probability of a vertex to be a center.
k	$\Theta((\log^2(1/\gamma) + \log d)/\gamma)$ – size parameter used to bound number of vertices in each part of \mathcal{P}
W	The set of centers.
t	$2kd$ (the random string used by the algorithm is t -wise independent).
$\text{cc}(v)$	Subset of vertices in the connected component of v in G .
Z	The set of vertices v such that $ \text{cc}(v) \geq k$.
$\sigma(v)$	The center of v with respect to W .
$C(w)$	$\{v \in V : \sigma(v) = w\}$ (the cell of w with respect to W).
$\ell_k(v)$	The minimum integer ℓ such that the BFS tree rooted at v of depth ℓ has size at least k .
$B_k(v)$	Vertices in the BFS-tree rooted at v of depth $\ell_k(v)$.
R	The set of vertices v for which $B_k(v)$ does not contain a vertex (center) in W .
$\mathcal{T}(w)$	The BFS tree rooted at w of the subgraph induced by $C(w) \setminus R$
\mathcal{P}'	The partition \mathcal{P} when restricted to Z .
$Y_k(v)$	$\{u \in V : v \in B_k(u)\}$
χ_w	The characteristic vector of the set W .
$\chi_w(S)$	The projection of χ_w onto S (where S is a subset of V).
ζ_v	The random variable indicating whether v is a remote vertex with respect to W .
\hat{S}_v	$\{u \in Z : B_k(u) \cap B_k(v) \neq \emptyset\}$.
\tilde{d}	$8(c(h))^2/\epsilon$.
\mathcal{H}	The set of heavy vertices (i.e., with degree greater than \tilde{d}).
\mathcal{L}	$V \setminus \mathcal{H}$, the set of light vertices.
$\mathcal{C}(u)$	The cluster of u , includes u and all the vertices in the parts that are assigned to u .
F	The subset of edges in $E' \cap (\mathcal{H} \times \mathcal{L})$ that are incident to two different clusters (namely, each endpoint belongs to a different cluster).
$E^b(B, B')$	The subset of edges in $E(B, B')$ whose rank is in the bottom $(\epsilon/(8c(h))) \cdot E(B, B') $ edges of $E(B, B')$.
$N^b(u, B')$	For u that is the leader of a cluster B , it is the set of vertices in $N(u)$ such that $(y \in B'$ and $(u, y) \in E^b(B, B')$) or $(\exists(y', z) \in E^b(B, B')$ s.t. $y' \in \mathcal{P}(y)$)

B Proof of Theorem 3

Proof of Theorem 3 Let V be a set of vertices and let v_0 and v_1 be a pair of distinct vertices in V . In order to prove the lower bound we construct two families of random d -regular graphs over V , $\mathcal{F}_{(v_0, v_1)}^+$ and $\mathcal{F}_{(v_0, v_1)}^-$. $\mathcal{F}_{(v_0, v_1)}^+$ is the family of d -regular graphs, $G = (V, E)$, for which $(v_0, v_1) \in E$. $\mathcal{F}_{(v_0, v_1)}^-$ is the family of d -regular graphs for which $(v_0, v_1) \in E$ and the removal of (v_0, v_1) leaves the graph with two connected components,⁹ each of size $n/2$. We prove that given (v_0, v_1) , any algorithm that performs at most \sqrt{n}/c queries for some sufficiently large constant $c > 1$ cannot distinguish the case in which the graph is drawn uniformly at random from $\mathcal{F}_{(v_0, v_1)}^+$ from the case in which the graph is drawn uniformly at random from $\mathcal{F}_{(v_0, v_1)}^-$. Essentially, if the number of queries is at most \sqrt{n}/c , then with high constant probability, each new query to the graph returns a new random vertex in both families. By “new vertex” we mean a vertex that neither appeared in the query history nor in the answers history. Since the algorithm must answer consistently with a connected graph G' , for every graph in the support of $\mathcal{F}_{(v_0, v_1)}^-$ it must answer positively on the query (v_0, v_1) , with probability 1 (or with high constant probability in case we allow G' to be disconnected with some small constant probability). But since the distributions on query-answer histories in both cases are very close statistically, this can be shown to imply that there exist graphs for which the algorithm answers positively on a large fraction of the edges.

We consider a pair of oracles that generate a random graph from $\mathcal{F}_{(v_0, v_1)}^+$ and $\mathcal{F}_{(v_0, v_1)}^-$, respectively, on the fly. On query (w, i) the oracle returns (v, j) where v is i -th neighbor of w , denoted $N(w, i)$, and w is the j -th neighbor of v . If w does not have an i -th neighbor then the oracle returns \emptyset . Both oracles construct the graph on demand, i.e. on query (w, i) , if $N(w, i)$ was not determined in a previous step then the oracle determines $N(w, i)$.

First Oracle Recall that a random d -regular graph can be generated as follows. Define a matrix $M = [n] \times [d]$ where the i -th row corresponds to the i -th vertex. Partition the cells of M into $nd/2$ pairs, namely, find a random perfect matching on the cells of M . Now let (w, i) and (v, j) be a pair of matched cells, then in the corresponding graph it holds that $N(w, i) = v$ and $N(v, j) = w$. Thus, to obtain a random d -regular graph from the family of d -regular graphs $G = (V, E)$, for which $(v_0, v_1) \in E$ proceed as follows.

1. Match a random cell in v_0 's row to a random cell with v_1 's row.
2. Find a random perfect matching of the remaining cells.

Our first oracle, $\mathcal{O}_{(v_0, v_1)}^+$, determines the perfect matching on demand as we describe next. The oracle keeps a matrix, M , with n rows which correspond to the n vertices and d columns, which correspond to d neighbors. At the initialization step:

⁹ Although a graph that is drawn uniformly from $\mathcal{F}_{(v_0, v_1)}^+$ (or $\mathcal{F}_{(v_0, v_1)}^-$) might be disconnected, this event happens with negligible probability [10]. Hence, the proof of the lower bound remains valid even if we consider $\mathcal{F}_{(v_0, v_1)}^+ \cap \mathcal{C}$ and $\mathcal{F}_{(v_0, v_1)}^- \cap \mathcal{C}$ where \mathcal{C} is the family of connected graphs.

1. All the cells are initialized to \emptyset .
2. A single edge is determined: $M(v_0, t_0) := (v_1, t_1)$ and $M(v_1, t_1) := (v_0, t_0)$ where t_0 and t_1 are chosen uniformly and independently over $[d]$.

On query (w, i) , the oracle $\mathcal{O}_{(v_0, v_1)}^+$, proceeds as follows:

1. Checks if $M(w, i)$ was determined in a previous step, i.e. if $M(w, i) \neq \emptyset$, if so it returns the matched cell (u, j) .
2. Otherwise, it picks uniformly one of the empty cell in the matrix, (u, j) .
3. Sets $M(w, i) := (u, j)$, $M(u, j) := (w, i)$. Namely, adds the edge (u, w) to the graph.
4. Returns (u, j) .

Without loss of generality assume that $n \cdot d$ is even.¹⁰ Thus, if the oracle is queried on all the entries of the matrix, then the resulting graph is a random d -regular graph that contain the edge (v_0, v_1) (where self-loops and parallel edges are allowed). Notice that the resulting graph is drawn uniformly (and independently of the order of the queries) from the family of graphs $\mathcal{F}_{(v_0, v_1)}^+$.

Second Oracle Our second oracle, $\mathcal{O}_{(v_0, v_1)}^-$ generates a random graph from the family of graphs $\mathcal{F}_{(v_0, v_1)}^-$. The oracle $\mathcal{O}_{(v_0, v_1)}^-$ keeps a pair of matrices M_0, M_1 . The matrix, M_0 , has $\lfloor n/2 \rfloor$ rows and d columns and the matrix, M_1 , has $\lceil n/2 \rceil$ rows and d columns.¹¹ At the initialization step:

1. All the cells of both matrices are initialized to \emptyset .
2. The rows of the matrices are not allocated to any vertex.
3. A random row in M_b, i_b , is allocated to v_b for each $b \in \{0, 1\}$.
4. A single edge is determined: $M_0(i_0, t_0) := M_1(i_1, t_1)$ and $M_1(i_1, t_1) := M_0(i_0, t_0)$ where t_0 and t_1 are chosen uniformly and independently over $[d]$.

For a vertex v , let $r(v)$ denote the index of the row that is allocated to v and the corresponding matrix by M_v . Given a row j and a matrix M , let $v(M, j)$ denote the vertex that the j -th row in M is allocated to. On query (w, i) , $\mathcal{O}_{(v_0, v_1)}^-$ proceeds as follows:

1. If a row was not allocated to w in previous steps then a random row is picked uniformly (from the set of rows that are free) and is allocated to w .
2. If the i -th neighbor of w was determined previously, i.e., the cell $M_w(r(w), i)$ was matched in previous steps to another cell $M_w(\ell, j)$, then return $(v(M_w, \ell), j)$.
3. Otherwise, select randomly and uniformly an empty cell in $M_w, (\ell, j)$. If the ℓ -th row of M_w is free, then allocate this row to a vertex which is picked uniformly from the set of vertices without an allocated row.
4. Add an edge between (w, i) and $(v(M_w, \ell), j)$, i.e., set $M_w(r(w), i) := M_w(\ell, j)$ and $M_w(\ell, j) = M_w(r(w), i)$. Return $(v(M_w, \ell), j)$.

¹⁰ In case that $d \cdot n$ is odd we add an extra entry to the matrix, $(0, 0)$, so that $M(w, i) = (0, 0)$ means that w does not have an i -th neighbor.

¹¹ Here too, if the number of entries in a matrix is odd we add an entry $(0, 0)$.

Let \mathcal{A} be an algorithm that interacts with $\mathcal{O}_{(v_0, v_1)}^+$ and let $\pi_r^+ = (q_1^+, a_1^+, \dots, q_r^+, a_r^+)$ be the random variable describing the communication between \mathcal{A} and $\mathcal{O}_{(v_0, v_1)}^+$. Namely, π_r^+ is a list of r queries of \mathcal{A} and r corresponding answers of $\mathcal{O}_{(v_0, v_1)}^+$. Similarly, define π_r^- for $\mathcal{O}_{(v_0, v_1)}^-$. We claim that for $r = c\sqrt{n}$, the distribution of π_r^+ is statistically close to the distribution of π_r^- . Recall that every query q_i is a pair $(q_{i,1}, q_{i,2})$ where the first entry denotes a name of a vertex and the second entry denotes an index in $[d]$. Similarly, every answer a_i is a pair $(a_{i,1}, a_{i,2})$. Let A^+ be the event that for every $i \in [r]$ it holds that, $a_{i,1}^+ \neq a_{j,1}^+$ for every $i \neq j \in [r]$ and $a_{i,1}^+ \neq q_{j,1}^+$ for every $j \in [r]$. Similarly, define A^- . Without loss of generality we assume that \mathcal{A} does not make queries that does not reveal any new information on the graph. Hence, in words, A^+ (respectively, A^-) is the event that whenever $\mathcal{O}_{(v_0, v_1)}^+$ (respectively $\mathcal{O}_{(v_0, v_1)}^-$) selects an empty (un-matched) cell, the corresponding row is empty as well. It is easy to verify that the distribution of π_r^+ conditioned on A^+ is identical to the distribution of π_r^- conditioned on A^- . Hence, the statistical distance between the distribution of π_r^+ and the distribution of π_r^- is bounded by $2|\Pr(A^+) - \Pr(A^-)|$. Since both $\Pr(A^+)$ and $\Pr(A^-)$ are upper bounded by $2r \cdot \frac{2r}{n/2} = 8r^2/n = 8c^2$, we obtain that the statistical distance is bounded by $16c^2 = 16/49 < 1/3$ for $c = 1/7$.

From this the lower bound is implied as follows. Assume towards contradiction that there exists an LSSG algorithm, \mathcal{A} , with query complexity $c\sqrt{n}$. When \mathcal{A} interacts with $\mathcal{O}_{(v_0, v_1)}^-$ it must answer yes with probability 1 (over the random coins of \mathcal{A} and $\mathcal{O}_{(v_0, v_1)}^-$) when queried on (v_0, v_1) . Since $\mathcal{O}_{(v_0, v_1)}^-$ and $\mathcal{O}_{(v_0, v_1)}^+$ are statistically close, \mathcal{A} that interacts with $\mathcal{O}_{(v_0, v_1)}^+$ must answer yes with probability at least $2/3$ (over the random coins of $\mathcal{O}_{(v_0, v_1)}^-$ and for every setting of random coins of \mathcal{A}) when queried on (v_0, v_1) . This implies that for every setting of random coins of \mathcal{A} , \mathcal{A} answers yes on at least $2/3$ of the d -regular graphs that contain (v_0, v_1) . By an averaging argument it follows that there exists a d -regular graph, G , such that \mathcal{A} answers yes on at least $2/3$ of the edges of G . A contradiction.

Now consider the setting in which we allow G' to be disconnected with probability $1 - \delta_1$ and we require that $|E'| \leq (1/3)nd$ with probability at least $1 - \delta_2$ for some constants $0 < \delta_1, \delta_2 \leq 1$. In this case \mathcal{A} must answer yes with probability $1 - \delta_1$, over the random coins of \mathcal{A} , when queried on (v_0, v_1) and interacting with $\mathcal{O}_{(v_0, v_1)}^-$ (for any setting of random coins of $\mathcal{O}_{(v_0, v_1)}^-$). Therefore, by an averaging argument, there exists a d -regular graph G , such that in expectation, \mathcal{A} answers yes on at least $(2/3)(1 - \delta_1)$ of the edges of G . However, since the expected size of E' is bounded by $(1 - \delta_2)(1/3)nd + \delta_2nd$ we reach a contradiction whenever $\delta_1 + \delta_2 < 1/2$. \square

C A Non-polynomial Relation Between k and $Y_k(v)$

Recall that $Y_k(v) \stackrel{\text{def}}{=} \{u \in V : v \in B_k(u)\}$. In the following lemma we show that $|Y_k(v)|$ can be super polynomial.

Lemma 12 *There exists a graph $G = (V, E)$ with degree bounded by d and $v \in V$ such that $|Y_k(v)| = k^{\Omega(\log \log d)}$.*

Proof The graph G is a tree, rooted at v , and defined as follows. For simplicity, we let the degree bound be $d + 1$ (so that a vertex may have d children, and hence degree $d + 1$). We partition the levels of the tree into consecutive subsets: $L_0 = \{1, \dots, \ell_0\}$ (where the root is at level 1), $L_1 = \{\ell_0 + 1, \dots, \ell_1\}, \dots, L_r = \{\ell_{r-1} + 1, \dots, \ell_r\}$. For each subset L_i , and for each level j in the subset, all vertices in level j have the same number of children, which is $d_i \stackrel{\text{def}}{=} d^{2^{-i}}$. We set $r = \log \log d$, so that all vertices in levels belonging to L_r have two children. Finally we set $s_i = |L_i| = \log_{d_i} g^{1/(r+1)}$, where g determines the size of the tree, as well as the minimum k that ensures that all vertices in the tree belong to $Y_k(v)$.

By the construction of the tree, the number of vertices in it is of the order of $\prod_{i=0}^r d_i^{s_i} = g$. In order to upper-bound k (such that all vertices belong to $Y_k(v)$), consider any vertex u in some level $j \in L_i$, where $0 \leq i \leq r$. Since $d_i = d_{i-1}^{1/2}$, so that $s_i = 2s_{i-1}$ for each i , we get that $s_i \leq \sum_{i' < i} s_{i'}$. Therefore, $\Delta(u, v) \leq 2s_i$. It follows that the number of vertices in the subtree rooted at u that are at distance at most $\Delta(u, v)$ from u is upper bounded by $d_i^{s_i} \cdot d_{i+1}^{s_{i+1}} < g^{3/2(r+1)}$. Since this is true for every vertex in the tree, we get that $|\Gamma_{\Delta(u,v)}(u)| = O(g^{3/2(r+1)} \cdot s_r) = O(g^{3/2(r+1)} \cdot \log g)$, which gives us an upper bound on k , from which the lemma follows. \square

References

1. Ailon, N., Chazelle, B., Comandur, S., Liu, D.: Property-preserving data reconstruction. *Algorithmica* **51**(2), 160–182 (2008)
2. Alon, N., Babai, L., Itai, A.: A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms* **7**(4), 567–583 (1986)
3. Alon, N., Rubinfeld, R., Vardi, S., Xie, N.: Space-efficient local computation algorithms. In: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1132–1139 (2012)
4. Alon, N., Seymour, P.D., Thomas, R.: A separator theorem for graphs with an excluded minor and its applications. In: Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13–17, 1990, Baltimore, Maryland, USA, pp. 293–299 (1990)
5. Andersen, R., Borgs, C., Chayes, J., Hopcroft, J., Mirrokni, V., Teng, S.: Local computation of pagerank contributions. *Internet Math.* **5**(1–2), 23–45 (2008)
6. Andersen, R., Chung, F., Lang, K.: Local graph partitioning using pagerank vectors. In: Proceedings of the Forty-Seventh Annual Symposium on Foundations of Computer Science (FOCS), pp. 475–486 (2006)
7. Andersen, R., Peres, Y.: Finding sparse cuts locally using evolving sets. In: Proceedings of the Forty-First Annual ACM Symposium on the Theory of Computing (STOC), pp. 235–244 (2009)
8. Bader, D.A., Cong, G.: A fast, parallel spanning tree algorithm for symmetric multiprocessors (smgs). *J. Parallel Distrib. Comput.* **65**(9), 994–1006 (2005)
9. Berkhin, P.: Bookmark-coloring algorithm for personalized pagerank computing. *Internet Math.* **3**(1), 41–62 (2006)
10. Bollobás, B.: *Random Graphs*. Cambridge University Press, Cambridge (2001)
11. Brakerski, Z.: Local property restoring. Unpublished manuscript (2008)
12. Campagna, A., Guo, A., Rubinfeld, R.: Local reconstructors and tolerant testers for connectivity and diameter. In: Proceedings of the Seventeenth International Workshop on Randomization and Computation (RANDOM), pp. 411–424 (2013)
13. Chazelle, B., Rubinfeld, R., Trevisan, L.: Approximating the minimum spanning tree weight in sub-linear time. *SIAM J. Comput.* **34**(6), 1370–1379 (2005)
14. Chazelle, B., Seshadhri, C.: Online geometric reconstruction. In: Proceedings of the Twenty-Second Annual ACM Symposium on Computation Geometry (SoCG), pp. 386–394 (2006)

15. Dutta, A., Levi, R., Ron, D., Rubinfeld, R.: A simple online competitive adaptation of lempel-ziv compression with efficient random access support. In: Proceedings of the Data Compression Conference (DCC), pp. 113–122 (2013)
16. Elkin, M., Neiman, O.: Efficient algorithms for constructing very sparse spanners and emulators. CoRR, [arXiv:1607.08337](https://arxiv.org/abs/1607.08337), (2016)
17. Even, G., Medina, M., Ron, D.: Deterministic stateless centralized local algorithms for bounded degree graphs. In: Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8–10, 2014. Proceedings, pp. 394–405 (2014)
18. Feige, U., Mansour, Y., Schapire, R.E.: Learning and inference in the presence of corrupted inputs. In: Proceedings of the 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3–6, 2015, pp. 637–657 (2015)
19. Fredman, M.L., Willard, D.E.: BLASTING through the information theoretic barrier with FUSION TREES. In: Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13–17, 1990, Baltimore, Maryland, USA, pp. 1–7 (1990)
20. Goldreich, O., Ron, D.: Property testing in bounded degree graphs. *Algorithmica* **32**(2), 302–343 (2002)
21. Hassidim, A., Kelner, J.A., Nguyen, H.N., Onak, K.: Local graph partitions for approximation and testing. In: Proceedings of the Fiftieth Annual Symposium on Foundations of Computer Science (FOCS), pp. 22–31 (2009)
22. Jeh, G., Widom, J.: Scaling personalized web search. In: Proceedings of the 12th International Conference on World Wide Web, pp. 271–279 (2003)
23. Jha, M., Raskhodnikova, S.: Testing and reconstruction of Lipschitz functions with applications to data privacy. In: Proceedings of the Seventeenth Annual Symposium on Foundations of Computer Science (FOCS), pp. 433–442 (2011)
24. Kale, S., Peres, Y., Seshadhri, C.: Noise tolerance of expanders and sublinear expander reconstruction. In: Proceedings of the Forty-Ninth Annual Symposium on Foundations of Computer Science (FOCS), pp. 719–728 (2008)
25. Kawarabayashi, K., Klein, P.N., Sommer, C.: Linear-space approximate distance oracles for planar, bounded-genus and minor-free graphs. In: Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4–8, 2011, Proceedings, Part I, pp. 135–146 (2011)
26. Kawarabayashi, K., Reed, B.A.: A separator theorem in minor-closed classes. In: 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23–26, 2010, Las Vegas, Nevada, USA, pp. 153–162 (2010)
27. Kutten, S., Peleg, D.: Fast distributed construction of small k -dominating sets and applications. *J. Algorithms* **28**(1), 40–66 (1998)
28. Lenzen, C., Levi, R.: 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, pp. 87:1–87:14 (2018)
29. Levi, R., Medina, M.: A (centralized) local guide. *Bull. EATCS* **122**, (2017)
30. Levi, R., Moshkovitz, G., Ron, D., Rubinfeld, R., Shapira, A.: Constructing near spanning trees with few local inspections. *Random Struct. Algorithms* **50**(2), 183–200 (2017)
31. Levi, R., Ron, D.: A quasi-polynomial time partition oracle for graphs with an excluded minor. *ACM Trans. Algorithms* **11**(3), 24:1–24:13 (2015)
32. Levi, R., Ron, D., Rubinfeld, R.: Local algorithms for sparse spanning graphs. In: Proceedings of the Eighteenth International Workshop on Randomization and Computation (RANDOM), pp. 826–842 (2014)
33. Levi, R., Ron, D., Rubinfeld, R.: A local algorithm for constructing spanners in minor-free graphs. In: Proceedings of the Twentieth International Workshop on Randomization and Computation (RANDOM), pp. 38:1–38:15 (2016)
34. Levi, R., Rubinfeld, R., Yodpinyanee, A.: Local computation algorithms for graphs of non-constant degrees. *Algorithmica* **77**, 1–24 (2016)
35. Linal, N.: Locality in distributed graph algorithms. *SIAM J. Comput.* **21**(1), 193–201 (1992)
36. Mader, W.: Homomorphiesätze für graphen. *Mathematische Annalen* **178**, 154–168 (1968)
37. Mansour, Y., Rubinfeld, R., Vardi, S., Xie, N.: Converting online algorithms to local computation algorithms. In: Automata, Languages and Programming: Thirty-Ninth International Colloquium (ICALP), pp. 653–664 (2012)

38. Mansour, Y., Vardi, S.: A local computation approximation scheme to maximum matching. In: Proceedings of the Sixteenth International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX), pp. 260–273 (2013)
39. Marko, S., Ron, D.: Distance approximation in bounded-degree and general sparse graphs. *ACM Trans. Algorithms* **5**(2), 475–486 (2009)
40. Mayer, A., Naor, S., Stockmeyer, L.: Local computations on static and dynamic graphs. In: Proceedings of the 3rd Israel Symposium on Theory and Computing Systems (ISTCS) (1995)
41. Naor, M., Stockmeyer, L.: What can be computed locally? *SIAM J. Comput.* **24**(6), 1259–1277 (1995)
42. Nash-Williams, Crispin St J.A.: Decomposition of finite graphs into forests. *J. London Math. Soc.* **s1-39**(1), 12 (1964)
43. Nguyen, H.N., Onak, K.: Constant-time approximation algorithms via local improvements. In: Proceedings of the Forty-Ninth Annual Symposium on Foundations of Computer Science (FOCS), pp. 327–336 (2008)
44. Orecchia, L., Zhu, Z.A.: Flow-based algorithms for local graph clustering. *CoRR*, [arXiv:1307.2855](https://arxiv.org/abs/1307.2855), (2013)
45. Parnas, M., Ron, D.: Approximating the minimum vertex cover in sublinear time and a connection to distributed algorithms. *Theor. Comput. Sci.* **381**(1–3), 183–196 (2007)
46. Parter, M., Rubinfeld, R., Vakilian, A., Yodpinyanee, A.: Local computation algorithms for spanners. In: 10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10–12, 2019, San Diego, California, USA, pp. 58:1–58:21 (2019)
47. Peleg, D., Rubinfeld, R.: A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction. *SIAM J. Comput.* **30**(5), 1427–1442 (2000)
48. Peleg, D., Schäffer, A.A.: Graph spanners. *J. Graph Theory* **13**, 99–116 (1989)
49. Peleg, D., Ullman, J.D.: An optimal synchronizer for the hypercube. *SIAM J. Comput.* **18**, 229–243 (1989)
50. Peleg, D.: Distributed computing. *SIAM Monogr. Discrete Math. Appl.* **5**, 1 (2000)
51. Pettie, S.: Distributed algorithms for ultrasparse spanners and linear size skeletons. *Distrib. Comput.* **22**(3), 147–166 (2010)
52. Ram, L.S., Vicari, E.: Distributed small connected spanning subgraph: Breaking the diameter bound. Technical report, Zürich (2011)
53. Rubinfeld, R., Tamir, G., Vardi, S., Xie, N.: Fast local computation algorithms. In: Proceedings of The Second Symposium on Innovations in Computer Science (ICS), pp. 223–238 (2011)
54. Saks, M.E., Seshadhri, C.: Local monotonicity reconstruction. *SIAM J. Comput.* **39**(7), 2897–2926 (2010)
55. Sarlos, T., Benczur, A., Csalogany, K., Fogaras, D., Racz, B.: To randomize or not to randomize: Space optimal summaries for hyperlink analysis. In: Proceedings of the 15th International Conference on WorldWide Web, pp. 297–306 (2006)
56. Spielman, D., Teng, S.: Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In: Proceedings of the Thirty-Sixth Annual ACM Symposium on the Theory of Computing (STOC), pp. 81–90 (2004)
57. Sudan, M., Trevisan, L., Vadhan, S.P.: Pseudorandom generators without the XOR lemma. *J. Comput. Syst. Sci.* **62**(2), 236–266 (2001)
58. Yoshida, Y., Yamamoto, M., Ito, H.: An improved constant-time approximation algorithm for maximum matchings. In: Proceedings of the Forty-First Annual ACM Symposium on the Theory of Computing (STOC), pp. 225–234 (2009)
59. Zhu, Z.A., Lattanzi, S., Mirrokni, V.: A local algorithm for finding well-connected clusters. In: Proceedings of the Thirtieth International Conference on Machine Learning (ICML) (2013)

Affiliations

Reut Levi¹  · Dana Ron² · Ronitt Rubinfeld^{3,4}

Dana Ron
danar@eng.tau.ac.il

Ronitt Rubinfeld
ronitt@csail.mit.edu

- ¹ Weizmann Institute of Science, Rehovot, Israel
- ² School of Electrical Engineering, Tel Aviv University, 69978 Tel Aviv, Israel
- ³ CSAIL, MIT, Cambridge, MA 02139, USA
- ⁴ Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel