

A New Application of Orthogonal Range Searching for Computing Giant Graph Diameters

Guillaume Ducoffe

National Institute for Research and Development in Informatics,
The Research Institute of the University of Bucharest ICUB,
University of Bucharest, Faculty of Mathematics and Computer Science, Romania
guillaume.ducoffe@ici.ro

Abstract

A well-known problem for which it is difficult to improve the textbook algorithm is computing the graph diameter. We present two versions of a simple algorithm (one being Monte Carlo and the other deterministic) that for every fixed h and unweighted undirected graph G with n vertices and m edges, either correctly concludes that $\text{diam}(G) < hn$ or outputs $\text{diam}(G)$, in time $\mathcal{O}(m + n^{1+o(1)})$. The algorithm combines a simple randomized strategy for this problem (Damaschke, *IWOCA'16*) with a popular framework for computing graph distances that is based on range trees (Cabello and Knauer, *Computational Geometry'09*). We also prove that under the Strong Exponential Time Hypothesis (SETH), we cannot compute the diameter of a given n -vertex graph in truly subquadratic time, even if the diameter is an $\Theta(n/\log n)$.

2012 ACM Subject Classification Theory of computation \rightarrow Shortest paths, Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases Graph diameter, Orthogonal Range Queries, Hardness in P, FPT in P

Digital Object Identifier 10.4230/OASICS.SOSA.2019.12

Funding This work was supported by the Institutional research programme PN 1819 "Advanced IT resources to support digital transformation processes in the economy and society - RESINFO-TD" (2018), project PN 1819-01-01 "Modeling, simulation, optimization of complex systems and decision support in new areas of IT&C research", funded by the Ministry of Research and Innovation, Romania. This work was also supported by a grant of Romanian Ministry of Research and Innovation CCCDI-UEFISCDI, project no. 17PCCDI/2018.

Acknowledgements We wish to thank the referees for their careful reading of the first version of this manuscript, and their useful comments.

1 Introduction

We refer to [5] for any undefined terminology. Graphs in this study are finite, simple, connected and unweighted. For every graph $G = (V, E)$, let $n := |V|$ and $m := |E|$. The distance $\text{dist}_G(u, v)$ between any two vertices $u, v \in V$ is defined as the minimum number of edges on a uv -path in G . A *layer* is any set $L_i(v) := \{u \in V \mid \text{dist}_G(u, v) = i\}$ for some $v \in V$ and integer $i \geq 0$. Finally, the eccentricity of vertex v , denoted $\text{ecc}_G(v)$, is equal to $\max_{u \in V} \text{dist}_G(u, v)$, and the diameter of G , denoted $\text{diam}(G)$, is equal to $\max_{v \in V} \text{ecc}_G(v)$.

Computing the diameter of a graph is a fundamental problem with countless applications in computer science and other domains. As every undergraduate student (should) know, this problem can be solved in roughly quadratic time by running a single-source shortest-path algorithm from every vertex of the graph. It has been asked repeatedly whether one could improve on this textbook algorithm, in order to achieve truly subquadratic-time computation



© Guillaume Ducoffe;
licensed under Creative Commons License CC-BY
2nd Symposium on Simplicity in Algorithms (SOSA 2019).

Editors: Jeremy Fineman and Michael Mitzenmacher; Article No. 12; pp. 12:1–12:7

OpenAccess Series in Informatics



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

of the diameter. Unfortunately, the answer to that question seems to be ‘No’ [15]. Specifically, under SETH it is already hard to recognize *split graphs* with diameter 2, and this holds even if their clique-number is an $\Theta(\log n)$ [6]¹. Small diameter graphs are usually said to be the hardest case for the problem, in the sense that as the diameter gets polynomial in n we can obtain an almost optimal $(1 - n^{-\mathcal{O}(1)})$ -approximation in truly subquadratic time [3, 10]. However, the *exact* computation of such “giant” graph diameters has been less studied.

In [11], Damaschke asked whether one can compute the graph diameter in nearly linear time assuming it is a large fraction of the number of vertices. His question seems relevant to the study of chain-like structures, e.g., in road networks or chain molecules. Specifically, we consider the following problem in this note:

► **Problem 1** (*h*-DIAMETER).

Input: A graph $G = (V, E)$; a constant $h \in (0; 1)$.

Output: The exact diameter of G if it is at least hn (otherwise, any value $< hn$).

As a partial answer to Problem 1, Damaschke presented a deterministic linear-time algorithm for the special case $h > 1/2$. The latter is based on the decomposition of a graph by its biconnected components and a delicate removal procedure of irrelevant subgraphs. As noted by Damaschke himself, qualitatively different methods are needed to solve the general case. In [11], he also presented a Monte Carlo $\mathcal{O}(m + n \log n)$ -time algorithm for the case $h > 1/3$. The probability of a correct result depends on the simultaneity of several random events. Recently in a master’s thesis [4] some of Damaschke’s students generalized his ideas to any h , thereby obtaining a (not so simple) $\mathcal{O}(n^2)$ -time algorithm for the general case. The “big-oh” notation hides a large constant-factor in h .

1.1 Our results

We answer positively to the open question of [11]. Specifically, we first present a Monte Carlo algorithm that runs in time $\mathcal{O}(\frac{1}{h} \cdot (m + 2^{\mathcal{O}(\frac{1}{h})} n^{1+o(1)}))$ and that, given as entry a graph G such that $\text{diam}(G) \geq hn$, outputs with constant probability its diameter (Theorem 5).

- For that, we follow the same general (and quite intuitive) strategy as in [11, 4]: finding a separator of size $\mathcal{O}(1/h)$ that disconnects the two ends of some arbitrary diametral path. Such a separator can be easily computed, with constant probability, by performing a BFS from a random vertex v and choosing a smallest layer $L_i(v) := \{u \in V \mid \text{dist}_G(u, v) = i\}$ in the range $i \in \{hn/3, \dots, 2hn/3\}$. Although a similar approach was used in [11, 4] our proofs are, in our opinion, cleaner and more direct than the ones given in these previous works. In particular, the correctness of our algorithm only depends on the random choice of the starting vertex v .
- Then, once we computed a separator as described above, we are left with computing the maximum distance between two vertices it disconnects. Previous works [11, 4] reduce this computation to a new problem called LARGEST MIXED SUM. Instead, we can directly apply a popular framework for computing graph distances that is based on a *range tree* [8]. Doing so, we give a new simple application of this textbook data-structure.
- Finally, we remark that in order to make our algorithm deterministic, it suffices to run a BFS from every vertex v into some (hn/c) -distance dominating set, for some small constant c (instead of picking this vertex randomly). We end up observing that such a distance dominating set of size $(\frac{1}{h})^{\mathcal{O}(1)}$ can be computed by using a few BFS (Theorem 7).

¹ The logarithmic bounds on the clique-number are not explicitly stated in [6]. Nevertheless, they can be deduced from an easy application of the Sparsification Lemma, as noted e.g., in [1] for similar constructions.

For nonconstant h , our algorithm still outperforms the textbook algorithm for DIAMETER provided $h = \omega(1/\log n)$. Perhaps surprisingly, we prove this is (conditionally) optimal: any truly subquadratic algorithm for computing the diameter of a given n -vertex graph would falsify SETH, even if the diameter is an $\Theta(n/\log n)$ (Theorem 8).

2 Preliminaries

In what follows is a simple observation that is the cornerstone of our algorithms: in any consecutive subsequence of $\Theta(n)$ layers in a BFS-tree, there must be one of size $\mathcal{O}(1)$.

► **Lemma 1.** *Let $G = (V, E)$ be a graph of order n , let $0 < p < q < r < 1$, and let $v \in V$ have $\text{ecc}_G(v) \geq rn$. There exists $i \in \{\lceil pn \rceil, \dots, \lfloor qn \rfloor\}$ such that the layer $L_i(v) := \{u \in V \mid \text{dist}_G(u, v) = i\}$ contains $< \frac{1-r}{q-p} + 1$ vertices.*

Proof. The number of layers $L_i(v)$ to be considered is:

$$\lfloor qn \rfloor - \lceil pn \rceil + 1 = \lfloor (q-p)n \rfloor + 1 > (q-p)n.$$

We also know that there are $\geq (rn - \lfloor qn \rfloor) + \lceil pn \rceil \geq (r - q + p)n$ vertices that are not contained into any of these layers. Therefore, the maximum number of vertices a smallest such layer can contain is:

$$< \frac{(1-r+q-p)n}{(q-p)n} = \frac{1-r}{q-p} + 1. \quad \blacktriangleleft$$

Then, assume one such layer disconnects the two ends of an arbitrary diametral path of the graph. In order to compute the diameter of the graph, we only need to compute the maximum distance between two vertices this layer disconnects. This is a routine that naturally appears in the computation of the diameter and the Wiener index of *bounded treewidth* graphs [1, 7, 8], and as such there already exists a standard method for solving this problem:

► **Proposition 2** (implicit in [7]). *Let $G = (V, E)$ be a graph and $S \subseteq V$ be a separator, where $|S| \leq k$. We can compute $D_S := \max\{\text{dist}_G(x, y) \mid S \text{ is an } xy\text{-separator}\}$ in time $\mathcal{O}(k \cdot (m + \binom{k+1+\lceil \log n \rceil}{k+1})2^k n)$, that is in $\mathcal{O}(k \cdot (m + 2^{\mathcal{O}(k)} n^{1+o(1)}))$.*

For the sake of completeness, let us give some intuition of how this above problem relates to *range trees*. Let $S = \{s_1, s_2, \dots, s_k\}$ and C_1, C_2, \dots, C_ℓ be the connected components of $G \setminus S$. For every $i \in \{1, \dots, k\}$ we search for the furthest pair x, y such that: (a) S is an xy -separator, and (b) $\text{dist}_G(x, s_i) + \text{dist}_G(s_i, y) = \text{dist}_G(x, y)$, as follows. We first map every $v \in C_j$ to the $(k+1)$ -dimensional point $P_v^i = (p_{v,0}^i, p_{v,1}^i, \dots, p_{v,k}^i)$, where $p_{v,0}^i = j$, $p_{v,t}^i = \text{dist}_G(v, s_t) - \text{dist}_G(v, s_i)$ for every $t \geq 1$, and we associate to this point the value $f_i(v) = \text{dist}_G(v, s_i)$. To understand why, note that if $x \in C_j, y \in C_{j'}$ are such that $j < j'$ and $\text{dist}_G(x, s_i) + \text{dist}_G(s_i, y) = \text{dist}_G(x, y)$, then we have $\text{dist}_G(x, s_i) + \text{dist}_G(s_i, y) \leq \text{dist}_G(x, s_t) + \text{dist}_G(s_t, y) \iff -p_{x,t}^i \leq p_{y,t}^i$ for any $t \geq 1$. Therefore, given $x \in C_j$, in order to find a furthest vertex $y \in \bigcup_{j' > j} C_{j'}$ from x , it suffices to compute a point P_y^i such that:

$$p_{x,0}^i < p_{y,0}^i; \quad -p_{x,t}^i \leq p_{y,t}^i \text{ for every } t \in \{1, \dots, k\}; \text{ and } f_i(y) \text{ is maximized.}$$

The $(k+1)$ -dimensional range tree is a classical data-structure that can be used in order to solve this above computation. Specifically, given that there are $|V \setminus S| = \mathcal{O}(n)$ points P_v^i to store, we can construct such a range tree in time $\mathcal{O}(k \binom{k+1+\lceil \log n \rceil}{k+1} n)$ in such a way that for every $x \notin S$, the corresponding query (computation of P_y^i) can be answered in time $\mathcal{O}(2^k \binom{k+1+\lceil \log n \rceil}{k+1})$ [14]. We stress that the analysis of this construction is involved, but its implementation is quite straightforward (e.g., see [7] for details).

ALGORITHM 1: GIANTDIAMETER.**Input:** graph $G = (V, E)$, h .**Output:** a lower-bound on $\text{diam}(G)$.

- 1: Let $v \in V$ picked u.a.r.
- 2: **if** $\text{ecc}_G(v) < 2hn/3$ **then**
- 3: **return** $\text{ecc}_G(v)$. // this occurs with proba. $\leq 1 - 2h/3$ if $\text{diam}(G) \geq hn$.
- 4: Find a layer $i \in \{\lceil hn/3 \rceil, \dots, \lfloor 2hn/3 \rfloor\}$ s.t. $|L_i(v)| \leq 3/h$.
- 5: Compute $D_i := \max\{\text{dist}_G(x, y) \mid L_i(v) \text{ is an } xy\text{-separator}\}$.
- 6: **return** $\max\{D_i\} \cup \{\text{ecc}_G(u) \mid u \in L_i(v)\}$.

Comparison with previous work. Damaschke's students solved a variant of the above problem in $\mathcal{O}(kn^2)$ -time by reducing it to a new problem they called LARGEST MIXED SUM [4]. Roughly, their solution consists in a brute force range searching. We use range trees in order to improve their running time, although in doing so we sacrifice analytical simplicity. A potential drawback of our algorithms compared to [4] is that the range tree data-structure has relatively high preprocessing and storage costs, that make it less practical for moderate values of k [2]. Some way to address this issue could be the use of alternative data-structures for range searching [2, 9].

3 Monte Carlo algorithm

We present in this section a simple algorithm for the computation of graph diameters that are at least a fixed fraction of the number of vertices. Unlike previous works [11, 4], we use randomization only to choose the starting vertex of our BFS run (Algorithm GIANTDIAMETER). Our algorithm is correct assuming this starting vertex is sufficiently close to an end of some (arbitrary) diametral path. We prove next that it happens with constant probability.

► **Lemma 3.** *Let $G = (V, E)$ be a graph and assume $\text{diam}(G) \geq hn$. For every vertex $v \in V$ that is drawn u.a.r., the following holds with probability $\geq 2h/3$: There exists a diametral pair $x, y \in V$ such that $\text{dist}_G(x, v) \leq hn/3$ (and so, $\text{dist}_G(y, v) \geq 2hn/3$).*

Proof. Fix an arbitrary diametral path P with ends $x, y \in V$. Every vertex $v \in V(P)$ such that either $\text{dist}_G(v, x) \leq hn/3$ or $\text{dist}_G(v, y) \leq hn/3$ satisfies the desired property, and there are exactly $2hn/3$ such vertices. ◀

► **Proposition 4.** *Let $G = (V, E)$ be a graph and assume $\text{diam}(G) \geq hn$. Algorithm GIANTDIAMETER correctly computes $\text{diam}(G)$ with probability $\geq 2h/3$.*

Proof. By Lemma 1 (applied with $p = h/3$ and $q = r = 2h/3$), a small-size layer $L_i(v)$ as requested by the algorithm always exists if $\text{ecc}_G(v) \geq 2hn/3$. Then, the algorithm is correct if there exists a diametral pair $x, y \in V$ such that there is no connected component of $G \setminus L_i(v)$ that both contains x, y (possibly, $x \in L_i(v)$ or $y \in L_i(v)$). This is always the case if $\min\{\text{dist}_G(v, x), \text{dist}_G(v, y)\} \leq hn/3$ (and so, $\max\{\text{dist}_G(v, x), \text{dist}_G(v, y)\} \geq 2hn/3$), and by Lemma 3 the latter happens with probability $\geq 2h/3$. ◀

The bottleneck of Algorithm GIANTDIAMETER is the computation of D_i . Using Proposition 2, we can conclude as follows:

► **Theorem 5.** *Let h be a fixed constant. In time $\mathcal{O}(\frac{1}{h} \cdot (m + 2^{\mathcal{O}(\frac{1}{h})} n^{1+o(1)}))$, we can either conclude a given graph G has diameter $< hn$, or compute its diameter, with probability of correctness $\geq 2h/3$.*

Proof. We run Algorithm GIANTDIAMETER, whose output is correct with probability $\geq 2h/3$ by Proposition 4. The dominant step for the algorithm is the computation of D_i , that can be done in time $\mathcal{O}(\frac{1}{h} \cdot (m + 2^{\mathcal{O}(\frac{1}{h})} n^{1+o(1)}))$ by Proposition 2. ◀

As usual, the probability of correctness can be increased to $1 - n^{-\mathcal{O}(1)}$ by running Algorithm GIANTDIAMETER $\mathcal{O}(\log n/h)$ times and outputting the maximum distance we obtained.

4 Deterministic algorithm

Next, we show how to derandomize our algorithm from the previous section. We recall that we use randomization only to choose the starting vertex v of some BFS run. Furthermore, the latter vertex v is correctly chosen if it is at a distance $\leq hn/3$ from an end of some arbitrary diametral path. Hence, instead of choosing v at random, we can try *all* the vertices contained into some (hopefully small) $(hn/3)$ -distance dominating set. We prove next that there always exists such a set of size polynomial in $1/h$.

► **Lemma 6.** *Let $G = (V, E)$ be a graph. In $\mathcal{O}(m + n)$ -time, we can output a set S where $|S| = \mathcal{O}(1/h)$ and such that $\text{dist}_G(v, S) \leq hn/3$ for every vertex $v \in V$.*

Proof. We use the constructive proof of Meir and Moon on k -distance dominating sets in trees [13]. Specifically, let T be an arbitrary spanning tree of G . Such a tree can be computed in $\mathcal{O}(n + m)$ -time by using, say, a breadth-first search. For any integer $k \geq 0$, we will explain next how to construct a k -distance dominating set of size $\lceil \frac{n}{k+1} \rceil$ for T (and so, also for G) in time $\mathcal{O}(n)$. By setting $k = \lfloor hn/3 \rfloor$, this will prove the lemma.

For that, we first compute the two ends of a diametral path in T , that can be easily done in time $\mathcal{O}(n)$ using two BFS [12]. Let x be any one of these two ends. If $n \leq k + 1$ or more generally, $\text{diam}(T) \leq k$ then, we can output $S = \{x\}$. Otherwise, we compute in time $\mathcal{O}(n)$ a breadth-first search from x in T . For every $i \in \{0, 1, 2, \dots, k\}$, let $S_i := \{v \in V \mid \text{dist}_T(x, v) \equiv i \pmod{k+1}\}$. Note that $S_i \neq \emptyset$ for every i since we assume $\text{diam}(T) > k$. Furthermore as proved in [13, Theorem 5], S_i is a k -distance dominating set of T for any fixed i . Since there are exactly $k + 1$ possibilities for i , there exists a i_0 such that $|S_{i_0}| \leq \lceil \frac{n}{k+1} \rceil$. We output $S = S_{i_0}$. ◀

► **Theorem 7.** *Let h be a fixed constant. In time $\mathcal{O}(\frac{1}{h^2} \cdot (m + 2^{\mathcal{O}(\frac{1}{h})} n^{1+o(1)}))$, we can either conclude a given graph G has diameter $< hn$, or compute its diameter.*

Proof. We apply Lemma 6 in order to compute an $(hn/3)$ -distance dominating set S of size $\mathcal{O}(1/h)$. Then, we apply Algorithm GIANTDIAMETER for every $v \in S$, and we output the maximum distance we obtain after these $|S|$ runs. ◀

5 Conditional Lower-bound

Our GIANTDIAMETER algorithm still runs in truly subquadratic time if $h = \omega(1/\log n)$. It would be interesting to compute in truly subquadratic time the *exact* diameter of a graph when it is an $\Theta(n/\log n)$. We prove that it cannot be done under standard complexity assumptions.

► **Theorem 8.** *Under SETH, there exists a function $h(n) = \Theta(1/\log n)$ such that the following problem cannot be solved in time $\mathcal{O}(n^{2-\varepsilon})$, for any $\varepsilon > 0$: Given an n -vertex graph G , either correctly decide $\text{diam}(G) < h(n)n$, or compute $\text{diam}(G)$.*

Proof. Let $G = (K \cup S, E)$ be a n -vertex split graph such that K induces a clique of size $|K| = \mathcal{O}(\log n)$ and S induces a stable set. We construct a graph $G' = (V', E')$ as follows. The vertex-set of G' is partitioned into two disjoint copies S^0, S^1 of the stable set S and n disjoint copies K_1, K_2, \dots, K_n of the clique K . For every $s \in S$ we denote by s^0 and s^1 the respective copies of s in S^0 and S^1 ; in the same way, for every $v \in K$ and $i \in \{1, 2, \dots, n\}$ we denote by v^i the copy of vertex v in K_i . Furthermore, every copy K_i induces a clique, we add an edge $\{v^i, v^{i+1}\}$ for every $v \in K, 1 \leq i < n$, and the two edges $\{s_0, v^1\}, \{s_1, v^n\}$ for every $s \in S, v \in K$ such that $\{s, v\} \in E$. By construction, G' has order $\mathcal{O}(n \log n)$ and size $\mathcal{O}(n \log^2 n)$. We also have:

- For $u, v \in K$, $\text{dist}_{G'}(u^i, v^j) = |j - i|$ if $u = v$, and $|j - i| + 1$ otherwise.
- For $s \in S, v \in K$, $\text{dist}_{G'}(s_0, v^i) = \text{dist}_{G'}(s_1, v^{k-1+i}) = i$ if $\{s, v\} \in E$, and $i+1$ otherwise.
- For $s, s' \in S$, $\text{dist}_{G'}(s_0, s'_0) = \text{dist}_{G'}(s_1, s'_1) = \text{dist}_G(s, s') \leq 3$.
- For $s, s' \in S$, $\text{dist}_{G'}(s_0, s'_1) = n - 1 + \max\{2, \text{dist}_G(s, s')\}$.

As a result, $n + 1 \leq \text{diam}(G') \leq n - 1 + \text{diam}(G) \leq n + 2$. In particular, $\text{diam}(G') = n - 1 + \text{diam}(G)$. This implies that computing $\text{diam}(G')$ is $\tilde{\mathcal{O}}(n)$ -time equivalent to computing $\text{diam}(G)$, that cannot be done in truly subquadratic time under SETH [6, 15]. ◀

References

- 1 A. Abboud, V. Vassilevska Williams, and J. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *SODA*, pages 377–391. SIAM, 2016. URL: <https://arxiv.org/abs/1506.01799>.
- 2 J. Bentley and J. Friedman. Data structures for range searching. *ACM Computing Surveys (CSUR)*, 11(4):397–409, 1979.
- 3 P. Berman and S. Kasviswanathan. Faster approximation of distances in graphs. In *WADS*, pages 541–552. Springer, 2007.
- 4 B. Block and M. Milakovic. Computing Diameters in Slim Graphs. Master’s thesis, Chalmers University of Technology, University of Gothenburg, Sweden, 2018. URL: <http://publications.lib.chalmers.se/records/fulltext/255208/255208.pdf>.
- 5 J. A. Bondy and U. S. R. Murty. *Graph theory*. Grad. Texts in Math., 2008.
- 6 M. Borassi, P. Crescenzi, and M. Habib. Into the square: On the complexity of some quadratic-time solvable problems. *Electronic Notes in Theoretical Computer Science*, 322:51–67, 2016. URL: <https://arxiv.org/abs/1407.4972>.
- 7 K. Bringmann, T. Husfeldt, and M. Magnusson. Multivariate Analysis of Orthogonal Range Searching and Graph Distances Parameterized by Treewidth. In *IPEC. LIPIcs*, 2018. to appear. URL: <https://arxiv.org/abs/1805.07135>.
- 8 S. Cabello and C. Knauer. Algorithms for graphs of bounded treewidth via orthogonal range searching. *Computational Geometry*, 42(9):815–824, 2009.
- 9 T. Chan. Orthogonal Range Searching in Moderate Dimensions: k-d Trees and Range Trees Strike Back. In *33rd International Symposium on Computational Geometry (SoCG 2017)*, volume 77 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/7226>.
- 10 S. Chechik, D. Larkin, L. Roditty, G. Schoenebeck, R. Tarjan, and V. Vassilevska Williams. Better approximation algorithms for the graph diameter. In *SODA*, pages 1041–1052. SIAM, 2014.

- 11 P. Damaschke. Computing Giant Graph Diameters. In *IWOCA*, pages 373–384. Springer, 2016.
- 12 C. Jordan. Sur les assemblages de lignes. *J. Reine Angew. Math*, 70(185):81, 1869.
- 13 A. Meir and J. Moon. Relations between packing and covering numbers of a tree. *Pacific Journal of Mathematics*, 61(1):225–233, 1975.
- 14 L. Monier. Combinatorial solutions of multidimensional divide-and-conquer recurrences. *Journal of Algorithms*, 1(1):60–74, 1980.
- 15 L. Roditty and V. Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *STOC*, pages 515–524. ACM, 2013.