



HAL
open science

Diameter computation on H-minor free graphs and graphs of bounded (distance) VC-dimension

Guillaume Ducoffe, Michel Habib, Laurent Viennot

► To cite this version:

Guillaume Ducoffe, Michel Habib, Laurent Viennot. Diameter computation on H-minor free graphs and graphs of bounded (distance) VC-dimension. SODA 2020 - ACM-SIAM Symposium on Discrete Algorithms, Jan 2020, Salt Lake City, United States. hal-02340382

HAL Id: hal-02340382

<https://inria.hal.science/hal-02340382>

Submitted on 30 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Diameter computation on H -minor free graphs and graphs of bounded (distance) VC-dimension

Guillaume Ducoffe¹, Michel Habib², and Laurent Viennot³

¹University of Bucharest, Faculty of Mathematics and Computer Science, and National Institute for Research and Development in Informatics, Romania*

²Paris University, and IRIF CNRS, France[†]

³Inria, and Paris University, France[‡]

Abstract

Under the Strong Exponential-Time Hypothesis, the diameter of general unweighted graphs cannot be computed in truly subquadratic time. Nevertheless there are several graph classes for which this can be done such as bounded-treewidth graphs, interval graphs and planar graphs, to name a few. We propose to study unweighted graphs of constant *distance VC-dimension* as a broad generalization of many such classes – where the distance VC-dimension of a graph G is defined as the VC-dimension of its ball hypergraph: whose hyperedges are the balls of all possible radii and centers in G . In particular for any fixed H , the class of H -minor free graphs has distance VC-dimension at most $|V(H)| - 1$.

- Our first main result is a Monte Carlo algorithm that on graphs of distance VC-dimension at most d , for any fixed k , either computes the diameter or concludes that it is larger than k in time $\tilde{O}(k \cdot mn^{1-\varepsilon_d})$, where $\varepsilon_d \in (0;1)$ only depends on d . We thus obtain a *truly subquadratic-time parameterized* algorithm for computing the diameter on such graphs.
- Then as a byproduct of our approach, we get the first truly subquadratic-time randomized algorithm for *constant* diameter computation on all the *nowhere dense* graph classes. The latter classes include all proper minor-closed graph classes, bounded-degree graphs and graphs of bounded expansion.
- Finally, we show how to remove the dependency on k for *any* graph class that excludes a fixed graph H as a minor. More generally, our techniques apply to any graph with constant distance VC-dimension and *polynomial expansion* (or equivalently having strongly sublinear balanced separators). As a result for all such graphs one obtains a truly subquadratic-time randomized algorithm for computing their diameter.

We note that all our results also hold for *radius* computation. Our approach is based on the work of Chazelle and Welzl who proved the existence of spanning paths with strongly sublinear *stabbing number* for every hypergraph of constant VC-dimension. We show how to compute such paths efficiently by combining known algorithms for the stabbing number problem with a clever use of ε -nets, region decomposition and other partition techniques.

*This work was supported by project PN 19 37 04 01 “New solutions for complex problems in current ICT research fields based on modelling and optimization”, funded by the Romanian Core Program of the Ministry of Research and Innovation (MCI) 2019-2022. This work was also supported by a grant of Romanian Ministry of Research and Innovation CCCDI-UEFISCDI, project no. 17PCCDI/2018.

[†]Supported by Inria Gang project-team, and ANR project DISTANCIA (ANR-17-CE40-0015).

[‡]Supported by Irif laboratory from CNRS and Paris University, and ANR project Multimod (ANR-17-CE22-0016).

1 Introduction

In this paper we present new results on exact diameter computation within several classes of unweighted (undirected) graphs with a geometric flavor. We recall that the diameter of an unweighted graph is the maximum number of edges on a shortest path. Beyond its many practical applications, this fundamental problem in Graph Theory has attracted a lot of attention in the fine-grained complexity study of polynomial-time solvable problems [1, 4, 8, 16, 19, 24, 27, 35, 57]. More precisely, for every n -vertex m -edge unweighted graph the textbook algorithm for computing its diameter runs in time $\mathcal{O}(nm)$. In a seminal paper [57] this roughly quadratic running-time (in the size $n + m$ of the input) was matched by a quadratic *lower-bound*, assuming the Strong Exponential-Time Hypothesis (SETH). We stress that for graphs with millions of nodes and edges, quadratic time is already prohibitive.

The conditional lower-bound of [57] also holds for sparse graphs *i.e.*, with only $m = \mathcal{O}(n)$ edges [1]. However it does *not* hold for many well-structured graph classes [1, 11, 13, 22, 15, 23, 25, 27, 33, 36, 38, 54]. Our work proposes some new advances on the characterization of graph families for which we can compute the diameter in truly subquadratic time.

1.1 Related work

Before we detail our contributions, we wish to mention a few recent (and not so recent) results that are most related to our approach.

Interval graphs. An early example of linear-time solvable special case for diameter computation is the class of interval graphs [54]. For every interval graph G and for any integer k , if we first compute an interval representation for G in linear-time [39] then we can compute by dynamic programming, for every vertex v , the contiguous segment of all the vertices at a distance $\leq k$ from v in G . It takes almost linear-time and it implies a straightforward quasi linear-time algorithm for diameter computation. More efficient algorithms for diameter computation on interval graphs and related graph classes were proposed in [22]. Nevertheless we will show in what follows that interval orderings are a powerful tool for diameter computation on more general geometric graph classes.

Bounded-treewidth graphs. More recently, quasi linear-time algorithms for diameter computation on bounded-treewidth graphs were presented in [1, 13] with almost optimal dependency on the treewidth parameter. The cornerstone of these algorithms is the use of k -range trees in order to detect the furthest pairs that are disconnected by some small-cardinality separators. Since then a few other applications of k -range trees and, more generally, orthogonal range searching for diameter computation, have been presented in [27, 28]. In our work we uncover deeper connections between diameter computation and range searching techniques from computational geometry.

Planar graphs. Finally, in a recent breakthrough paper [15], Cabello presented the first truly subquadratic algorithm for diameter computation on planar graphs (see also [38] for improvements on his work). For that he combined r -divisions: a recursive decomposition technique for planar graphs and other hereditary graph classes with sublinear balanced separators, with a clever use of additively weighted Voronoi diagrams. Cabello conjectured that his algorithm could be generalized to bounded-genus graphs. The long version of [38] indicates that their techniques could allow such a generalization if computing the diameter of a graph embedded onto a surface of genus g reduces

to the planar case with $O(g)$ holes in the regions of some r -division. Although it is known that such a graph can be decomposed into planar subgraphs by removing $2g$ shortest paths [44, 34], such reduction is not clear, and we could not find references formally supporting this. More recently, Li and Parter proposed a distributed algorithm for planar diameter which is based on metric compression [48] and uses a VC-dimension argument to bound the number distance profiles with respect to a given subset of nodes. Following the basics of planar algorithms, we partly reuse r -divisions within our algorithms. However we replace the intricate use of Voronoi diagrams with a quite different approach that is based on some interval representations of the balls of a given radius in a graph. Our approach is also based on a VC-dimension argument but in a very different way than [48]. In doing so, we can obtain truly subquadratic-time algorithms for diameter computation on bounded genus graphs (and more generally, on any proper minor-closed graph family) while avoiding a great deal of topological complications. Note that our approach works similarly for computing the radius whereas it not clear whether the Voronoi diagram approach does.

We stress that for the three aforementioned graph classes, the techniques used for computing their diameter are quite different from each other. Our work is a first step toward unifying all these previous results for unweighted graphs in a single framework (note that some of the aforementioned results also hold in the directed weighted case).

1.2 Our contributions

We study the parameterization of graph diameter by the *VC-dimension* of various hypergraphs. More precisely, a set Y is *shattered* by a hypergraph \mathcal{H} if by intersecting Y with all hyperedges of \mathcal{H} one obtains the power-set of Y . The VC-dimension of \mathcal{H} is then defined as the largest cardinality of a subset shattered by \mathcal{H} . This powerful notion was first introduced by Vapnik and Chervonenkis in [60]. Since then it has found applications in sampling complexity and machine learning, among other domains. We refer to [47] for early work on VC-dimension in graphs. In particular, the VC-dimension of a graph G is defined as the VC-dimension of its closed neighbourhood hypergraph: whose hyperedges are the closed neighbourhoods of vertices in G . Graphs of bounded interval number and proper minor-closed graph classes are two examples of graph families with a *constant* upper-bound on their VC-dimension [47, 20].

First example. As an appetizer we first consider an n -vertex split graph with clique-number $\log^{\mathcal{O}(1)} n$, that is a notoriously hard case for diameter computation [8]. Given such a split graph G with stable set S and maximal clique K , we can pre-process G in linear-time so as to partition the vertices of S into *twin classes*: with two vertices in S being called twins if and only if they have the same neighbourhood in K (*e.g.*, see [23]). If the VC-dimension of G is at most d then, by the Sauer-Shelah-Perles Lemma [58, 59] the number of twin classes is in $\mathcal{O}(|K|^d) = \log^{\mathcal{O}(d)} n$. Therefore, after some linear-time preprocessing, we are left with computing the diameter on a graph of *polylogarithmic order*! Unfortunately, such simple brute-force arguments are no longer sufficient for split graphs of arbitrary clique-size.

Overview of our techniques. In order to generalize our approach to any graph of constant VC-dimension, we use the central notion of *spanning paths with low stabbing number*. Chazelle and Welzl [18] defined a spanning path for a hypergraph \mathcal{H} as a total ordering of its vertex-set. The

stabbing number of such a path is, up to 1, the maximum number of maximal intervals of which a hyperedge in \mathcal{H} can be the union (we refer to Sec. 2 for a formal definition).

Assume for now that we are given a spanning path with stabbing number t for the closed neighbourhood hypergraph of G . Then in linear time, we can compute for every vertex v the ends of the $\mathcal{O}(t)$ intervals of which $N_G[v]$ is the union. We denote this set of intervals by $I(v)$ in what follows. Then, in order to decide whether G has diameter at most two, it is sufficient to check whether for every vertex u we have $\bigcup_{v \in N_G[u]} I(v) = V$. Since we only need to consider the ends of such intervals, this verification phase takes time $\mathcal{O}(\deg_G(u) \cdot t)$ for a vertex of degree $\deg_G(u)$, and so, $\mathcal{O}(tm)$ total time. Note that such running-time is always subquadratic if t is sublinear in n . Overall, we reduced the diameter-two problem to the computation of a spanning path with low stabbing number for the closed neighbourhood hypergraph.

Motivated by range searching problems, Chazelle and Welzl proved the existence of spanning paths with *strongly sublinear* stabbing number for every hypergraph of constant VC-dimension [18]! Following this approach, we obtain our first main result in this paper:

Theorem 1. *For every $d > 0$, there exists a constant $\varepsilon_d \in (0; 1)$ such that in deterministic time $\tilde{\mathcal{O}}(mn^{1-\varepsilon_d})$ we can decide whether a graph of VC-dimension at most d has diameter two.*

We stress that in contrast to Theorem 1, under the Strong Exponential-Time we cannot decide whether a *general* graph has diameter at most two in truly subquadratic time [57].

On our way to prove Theorem 1 our main difficulty was to show how to compute for a hypergraph \mathcal{H} a spanning path of low stabbing number. Computing a spanning path of minimum stabbing number is NP-hard [6]. However, there exist approximation algorithms for this problem that run in polynomial time [6, 40]. Their approximation ratio is logarithmic, that is fine for our applications. Unfortunately, the fastest known algorithms require us to solve a linear program. So far, the best known algorithms for this intermediate problem run in superquadratic time [21]. We show how to decrease the running-time of this part, at the price of a slightly increased stabbing number. For that, we carefully apply the deterministic algorithm resulting from [18] to some arbitrary partition of \mathcal{H} in subhypergraphs of sublinear size. This nice trick might be of independent interest. We thus state the following theorem where the size of a hypergraph is defined as the sum of its hyperedge cardinalities.

Theorem 2. *For every $d > 0$, there exists a constant $\varepsilon_d \in (0; 1)$ such that in $\tilde{\mathcal{O}}(m + n^{2-\varepsilon_d})$ deterministic time, for every n -vertex hypergraph \mathcal{H} of VC-dimension at most d and size m , we can compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$.*

Moreover, $\varepsilon_d = \frac{1}{2^{d+1} \lfloor c(d+1) - 1 \rfloor + 1}$ for some universal constant $c > 2$.

From VC-dimension to *distance* VC-dimension. In order to go beyond Theorem 1, we need to consider a stronger notion of VC-dimension for graphs. The *distance VC-dimension*¹ of G is equal to the VC-dimension of its *ball hypergraph*: of which the hyperedges are all possible balls in G . Note that a bounded distance VC-dimension implies a bounded VC-dimension, but the converse a priori does not hold. Nevertheless, and perhaps surprisingly, there are still many classes of graphs with a constant distance VC-dimension. These classes include, among others: interval graphs, planar graphs [20] and, more generally, any proper minor-closed graph family (from Remark 3 in [20]), as well as graphs of bounded rank-width [10].

¹Our definition of distance VC-dimension is slightly weaker than the one proposed in [10].

Theorem 3. *There exists a Monte Carlo algorithm such that, for every positive integers d and k , we can decide whether a graph of distance VC-dimension at most d has diameter at most k . The running time is in $\tilde{O}(k \cdot mn^{1-\varepsilon_d})$, where $\varepsilon_d \in (0; 1)$ only depends on d .*

Eppstein proved in [33] that for any constant k , we can decide in linear time whether the diameter of a planar graph is at most k . Our result can be seen as a generalization of his to any graph class of constant distance VC-dimension – but at the price of a superlinear running-time. Furthermore, our techniques also apply to superconstant diameters, say polylogarithmic in n , or even polynomial in n provided the exponent is in $o(\varepsilon_d)$.

Our main technical contribution in this part is the efficient computation of spanning paths with strongly sublinear stabbing number for some *dense* hypergraphs of constant VC-dimension. More precisely, the ℓ -neighbourhood hypergraph of G has for hyperedges the balls of radius ℓ in G . For instance, the 1-neighbourhood hypergraph of G is exactly its closed neighbourhood hypergraph. In order to prove Theorem 3, we reduce the problem of deciding whether a graph has diameter at most k to the computation of a spanning path with low stabbing number for its $(k - 1)$ -neighbourhood hypergraph. In this sense, the proofs of Theorems 1 and 3 are very similar. However, an additional difficulty here is that we cannot have direct access to this $(k - 1)$ -neighbourhood hypergraph. Indeed, in the worst case all hyperedges of this hypergraph may have a cardinality in $\Omega(n)$, and then storing the hypergraph itself would already require quadratic space.

We overcome this issue by computing an ε -net [42, 60] in order to partition the vertices of the graph in a small number of groups, with every two vertices in the same group having almost the same ball of radius $k - 1$. By selecting only one vertex per group, we so reduce the number of hyperedges (*i.e.*, balls of radius $k - 1$) to be considered. Finally, once a spanning path was computed for this smaller hypergraph, for every unselected vertex we compute the symmetric difference between its ball of radius $k - 1$ and the one of the unique vertex taken in its group. Our solution in order to do that efficiently is to first compute a spanning path with low stabbing number for the $(k - 2)$ -neighbourhood hypergraph. This is where the dependency on k occurs, as overall we will need to compute a spanning path for $k - 1$ consecutive hypergraphs. Our algorithm is randomized and succeeds with high probability. The use of randomization comes from the ε -net construction. Although deterministic algorithms do exist for that [14], it is not clear whether they can be used as efficiently as the simple sampling technique of the randomized algorithm. We leave open the question of finding a deterministic variant of Theorem 3.

We note that this above technique can be applied under slightly weaker hypothesis than the one we state in Theorem 3. For instance, Nešetřil and Ossona de Mendez proved that for all *nowhere dense* graph classes (*i.e.*, a broad generalization of proper minor-closed graph classes and bounded-degree graphs), for any graph in the class and for any constant k , the VC-dimension of the k -neighbourhood hypergraph is constantly upper-bounded [53]. It allows us to derive the following weaker version of our Theorem 3:

Theorem 4. *Let \mathcal{G} be a class of nowhere dense graphs. There exists a Monte Carlo algorithm such that, for every constant $k = \mathcal{O}(1)$, for any graph in \mathcal{G} we can decide whether its diameter is at most k in time $\tilde{O}(mn^{1-\varepsilon_{\mathcal{G}}(k)})$, for some constant $\varepsilon_{\mathcal{G}}(k) \in (0; 1)$ that only depends on k .*

Let us mention that under SETH, Theorem 4 is the best result that we can hope for nowhere dense graph classes. Indeed, bounded-degree graphs are nowhere dense and, under SETH, we cannot compute their diameter in truly subquadratic time even if it is in $\omega(\log n)$ [35].

We conjecture that on every graph family of constant distance VC-dimension, we can compute the diameter in truly subquadratic time. Our next main result shows the conjecture to be true for any monotone graph family with strongly sublinear balanced separators, *a.k.a* the graphs of *polynomial expansion* [30].

Theorem 5. *Let \mathcal{G} be a monotone graph class with strongly sublinear balanced separators. Then there exists a Monte Carlo algorithm such that, for every $d > 0$, we can compute the diameter of any graph in \mathcal{G} of distance VC-dimension at most d in time $\tilde{O}(n^{2-\varepsilon_{\mathcal{G}}(d)})$, for some constant $\varepsilon_{\mathcal{G}}(d) \in (0; 1)$ that only depends on d .*

Let us recall that H -minor free graphs have a constant distance VC-dimension from Remark 3 in [20] (see also [10]), and that they all have strongly sublinear balanced separators [2, 45, 62]. Therefore, as an important consequence of Theorem 5, we get a truly subquadratic-time algorithm for computing the diameter on all the proper minor-closed graph classes.

It might be tempting, in the above Theorem 5, to drop the assumption that the distance VC-dimension must be bounded. Unfortunately, this cannot be done assuming SETH. Indeed, there is also an equivalence between the graphs of strongly sublinear treewidth and those monotone graph classes with strongly sublinear balanced separators [31]; however it follows from [1] that under SETH, we cannot compute the diameter in truly subquadratic time already for n -vertex graphs of treewidth $\omega(\log n)$. Conversely, not all graph classes with constant distance VC-dimension have strongly sublinear separators. This can be seen, *e.g.*, with interval graphs.

The speed-up of Theorem 5 follows from a faster computation of spanning paths for the neighbourhood hypergraphs. More precisely, we explain how to compute a spanning path for the $2k$ -neighbourhood hypergraph of G from a spanning path of its k -neighbourhood hypergraph. Note that in doing so, we only need to consider *logarithmically* many intermediate hypergraphs in order to compute such spanning path. Our approach for that consists in computing a first (suboptimal) representation of the $2k$ -neighbourhood of every vertex. Then, as for Theorem 3, we partition the vertices into a small number of groups and we select a unique vertex in each group. The suboptimal representations are used at the end of the algorithm in order to compute, for every unselected vertex, the symmetric difference between its ball of radius $2k$ and the one of the unique vertex taken in its group. So the problem becomes how to compute efficiently these suboptimal representations?

For that, we use a rather classical divide-and-conquer approach. Federickson [37] proved that a planar graph can be edge-covered with $\mathcal{O}(n/r)$ subgraphs of order at most r such that at most $\mathcal{O}(\sqrt{r})$ vertices of each subgraph can be contained in another subgraph of this decomposition. His construction directly follows from the planar separator theorem of Lipton and Tarjan [49], and as such it can be easily adapted for any monotone graph family with sublinear balanced separators [43]². For illustrating our method, we now focus in this introduction on the planar case. We can first compute, for some well-chosen $r = n^\gamma$, $\gamma \in (0; 1)$, a decomposition as described above. For every two vertices in a same subgraph, we can check whether they are at distance at most $2k$ by checking whether their balls of radius k intersect; assuming r is small enough, and we precomputed a spanning path with low stabbing number for the k -neighbourhood hypergraph, this phase can be implemented in order to run in truly subquadratic time. Then for every subgraph of

²Note that Federickson proposed several refinements of his construction in [37], some of which do use the fact that the input graph is planar. We will use in our proofs an even weaker version of his result than the one presented in this introduction.

the decomposition, we compute a breadth-first search from each of the $\mathcal{O}(\sqrt{r})$ boundary vertices that are also contained in another subgraph. Overall, there can only be $\mathcal{O}(n/\sqrt{r})$ such boundary vertices, and so, it takes truly subquadratic time. Furthermore in doing so, we computed for every subgraph of the decomposition the $\mathcal{O}(r\sqrt{r})$ distances between the boundary vertices and all the others. For any vertex v that is *not* on the boundary, we observe that a vertex in another subgraph can be at a distance $\leq 2k$ from v if and only if it is at distance $\leq 2k - \text{dist}_G(v, x)$ from some vertex x on the boundary ($\mathcal{O}(\sqrt{r})$ balls to be considered). Our strategy consists in computing a spanning path with low stabbing number for some “boundary hypergraph” whose hyperedges are the $\mathcal{O}(r\sqrt{r} \times (n/r)) = \mathcal{O}(n\sqrt{r})$ balls that we need to consider. We encounter a similar problem as for Theorem 3 because storing this hypergraph may require superquadratic space. Fortunately, we can encode this hypergraph in a much more compact way by taking advantage of (i) the fact that we can only have $\mathcal{O}(n/\sqrt{r})$ different centers for the balls, and (ii) that all the balls with a same center have a chain-like inclusion structure.

Although we keep the focus on computing the diameter, we shall stress in Sec. 2.4 that all our techniques can also be applied to *radius* computation (i.e., see Remark 1). Our algorithms almost need no particular information about the graph structure in order to be applied. In fact, we do not even need to compute the (distance) VC-dimension of the input graph! From the applicative point of view, this observation (further discussed in Sec. 2.4) is quite important. Indeed, computing the VC-dimension is W[1]-hard [26] and LogNP-hard [55].

1.3 Organization of the paper

In Sec. 2 we formally introduce the concepts of (distance) VC-dimension and stabbing number, along with some of their basic properties. Then, we explain in Sec. 3 how to compute a spanning path with strongly sublinear stabbing number for a hypergraph of constant VC-dimension (Theorem 2). As a direct application, we give a short proof of Theorem 1. Our techniques are generalized in Sec. 4 so as to prove Theorems 3 and 4. Finally, our main technical result (Theorem 5) is proved in Sec. 5. For that, we will need to recall some useful results on the graphs of *polynomial expansion* [30]. We discuss some possible future work in Sec. 6.

2 Preliminaries

After recalling a few basic definitions about graphs and hypergraphs (Sec. 2.1 and 2.2) we introduce our framework for computing the diameter of a graph in Sec. 2.3 and 2.4.

2.1 Graphs and Diameter

For any undefined graph terminology, see [7]. Throughout all this paper we only consider graphs that are undirected, unweighted and connected. For every graph $G = (V, E)$, let $n := |V|$ be its order and $m := |E|$ be its size. We denote by $N_G(v)$ and $N_G[v] := N_G(v) \cup \{v\}$ the open and closed neighbourhoods of vertex v , respectively. The degree of v is equal to $|N_G(v)|$ and is denoted by $\text{deg}_G(v)$ in what follows. The length of a path is its number of edges, and the distance $\text{dist}_G(u, v)$ between $u, v \in V$ is equal to the length of a shortest uv -path. For every $v \in V$ and $k \geq 0$, the k -neighbourhood of v , also known as the ball of center v and radius k , is defined as $N_G^k[v] = \{u \in V \mid \text{dist}_G(u, v) \leq k\}$. For instance, $N_G^1[v]$ is exactly the closed neighbourhood of v . The diameter of G is equal to $\text{diam}(G) = \max_{u, v \in V} \text{dist}_G(u, v)$.

Problem 1 (DIAMETER).

Input: A graph $G = (V, E)$.

Output: The diameter of G .

Theorem 6 ([57]). *Under the Strong Exponential-Time Hypothesis, we cannot decide whether a graph has diameter at most two in time $\mathcal{O}(mn^{1-\varepsilon})$, for any $\varepsilon > 0$.*

2.2 Hypergraphs

More generally, a hypergraph is a pair $\mathcal{H} = (X, R)$ with X being the set of vertices and $R \subseteq 2^X$ being the set of hyperedges. See also [5] for any undefined hypergraph terminology. Let $n := |X|$, $m := \sum_{q \in R} |q|$ and $r := |R|$ be the order, the size and the number of hyperedges of \mathcal{H} , respectively. For every vertex $x \in X$, let $R_x := \{q \in R \mid x \in q\}$. The *dual* of \mathcal{H} is the hypergraph $\mathcal{H}^* := (R, X^*)$, where $X^* := \{R_x \mid x \in X\}$. In particular, \mathcal{H} and \mathcal{H}^{**} are isomorphic.

Several hypergraphs can be related to a graph G :

- The closed neighbourhood hypergraph, denoted by $\mathcal{N}_1(G)$, has vertex-set $X = V$ and hyperedge-set $R = \{N_G[v] \mid v \in V\}$;
- More generally, for every fixed $\ell \geq 0$, the ℓ -neighbourhood hypergraph of G is defined as $\mathcal{N}_\ell(G) = (V, \{N_G^\ell[v] \mid v \in V\})$. We stress that $\mathcal{N}_\ell(G)$ and its dual $\mathcal{N}_\ell^*(G)$ are isomorphic [10].
- Finally, the ball hypergraph of G , simply denoted by $\mathcal{B}(G)$, has for hyperedges the balls of all possible centers and radii in G . Equivalently, $\mathcal{B}(G) = \bigcup_{\ell \geq 0} \mathcal{N}_\ell(G)$.

2.3 VC-dimension

Let $\mathcal{H} = (X, R)$ be a fixed hypergraph. A subset $Y \subseteq X$ is *shattered* by \mathcal{H} if, for every $Y' \subseteq Y$, there exists a hyperedge $q \in R$ such that $Y \cap q = Y'$. Then, the Vapnik-Chervonenkis dimension of \mathcal{H} (abbreviated in what follows to *VC-dimension*) is the largest cardinality of a shattered subset. Similarly, the *dual VC-dimension* of \mathcal{H} is the VC-dimension of its dual \mathcal{H}^* . We will often use the following (easy) properties in our analysis:

Lemma 1 (Sauer-Shelah-Perles, [58, 59]). *Every n -vertex hypergraph of VC-dimension at most d has $\mathcal{O}(n^d)$ hyperedges.*

Lemma 2 ([18]). *Every hypergraph of VC-dimension d has dual VC-dimension at most 2^{d+1} .*

Lemma 3 ([46]). *For every hypergraph $\mathcal{H} = (X, R)$ and $Y \subseteq X$, let $R[Y] = \{q \cap Y \mid q \in R\}$. Then, the VC-dimension of $\mathcal{H}[Y] := (Y, R[Y])$ is at most the VC-dimension of \mathcal{H} .*

VC-dimension for graphs. The *VC-dimension* of a graph G is defined as the VC-dimension of its closed neighbourhood hypergraph $\mathcal{N}_1(G)$. For instance, K_h -minor free graphs (and so, H -minor free graphs for any H of order at most h) have VC-dimension at most $h - 1$ [3]. Every k -interval graph has VC-dimension in $\mathcal{O}(k \log k)$ [28]. Other classes of constant VC-dimension – at most three – are unit disk graphs, chordal bipartite graphs, C_4 -free bipartite graphs, graphs of girth at least five and undirected path graphs [9].

The *distance VC-dimension* of a graph G is defined as the VC-dimension of its ball hypergraph $\mathcal{B}(G)$. Chepoi, Estellon and Vaxès proved in [20] that planar graphs have distance VC-dimension at most 4, and remarked that more generally every K_h -minor free graph has distance VC-dimension at most $h - 1$. Bousquet and Thomassé proved in [10] that graphs of bounded distance VC-dimension also generalize graphs of bounded rankwidth. Indeed, every graph of rankwidth k has distance VC-dimension at most $3 \cdot 2^{k+1} + 1$. For purpose of illustration, we next adapt a proof from [9] in order to show that interval graphs have distance VC-dimension at most two:

Lemma 4. *Every interval graph has distance VC-dimension at most 2.*

Proof. Let $G = (V, E)$ be an interval graph. We fix an interval model for G . For every $v \in V$, let $I(v) = [a_v, b_v]$ be the corresponding interval in the representation. Suppose now by contradiction that there is a set $S = \{v_1, v_2, v_3\}$ that is shattered by $\mathcal{B}(G)$. W.l.o.g., $a_{v_1} < a_{v_2} < a_{v_3}$. Since S is shattered, there exist some $u \in V$ and $k \geq 0$ such that $N_G^k[u] \cap S = \{v_1, v_3\}$. But then, let $I_{k-1}(u) := \bigcup_{w \in N_G^{k-1}[u]} I(w)$ be the contiguous segment of all the vertices at a distance $\leq k - 1$ from u . Note that $I_{k-1}(u) \cap I(v_2) = \emptyset$ because we assume that $v_2 \notin N_G^k[u]$. In this situation, either $I_{k-1}(u) \subseteq]-\infty, a_{v_2}[$ or $I_{k-1}(u) \subseteq]b_{v_2}, \infty[$ where $]x, y[= [x, y] \setminus \{x, y\}$ denotes the open interval between x and y . In fact we must have $I_{k-1}(u) \subseteq]b_{v_2}, \infty[$ because otherwise, $I_{k-1}(u) \cap I(v_3) = \emptyset$ and so, $v_3 \notin N_G^k[u]$, a contradiction. Since $I_{k-1}(u) \cap I(v_1) \neq \emptyset$, it implies that $b_{v_1} > b_{v_2}$, and so, $I(v_2) \subseteq I(v_1)$. As a result we have $N_G[v_2] \subseteq N_G[v_1]$. But then, for any $w \in V$ and $\ell \geq 1$, we have $v_2 \in N_G^\ell[w] \implies v_1 \in N_G^\ell[w]$. The latter contradicts our hypothesis that S is shattered. \square

2.4 Stabbing number and applications to DIAMETER

A spanning tree of $\mathcal{H} = (X, R)$ is a tree T whose node-set is exactly X . The *stabbing number* of such spanning tree T is the least k such that, for every hyperedge $q \in R$, there exist at most k edges $uv \in E(T)$ such that $|q \cap \{u, v\}| = 1$ (we also say that uv is *stabbed* by q). Given a set $q \subseteq X$, we let $E_T(q) = \{uv \in E(T) \mid u \in q, v \notin q\}$ of all edges stabbed by q . Finally, the stabbing number of \mathcal{H} is the minimum stabbing number over its spanning *paths*. Indeed, as noted in [18], every spanning tree T can be transformed into a spanning path of stabbing number at most twice bigger than for T . Therefore, there is essentially no loss of generality in restricting ourselves to spanning paths.

Lemma 5 ([18]). *Every n -vertex hypergraph of dual VC-dimension d has stabbing number $\tilde{O}(n^{1-\frac{1}{d}})$.*

Overall it follows from Lemmata 2 and 5 that any n -vertex hypergraph of VC-dimension at most d has strongly sublinear stabbing number in $\tilde{O}(n^{1-\frac{1}{2d+1}})$. We stress that the proof of Lemma 5 is constructive but that it cannot be transformed into a truly subquadratic-time algorithm. Efficient computations of spanning paths with sublinear stabbing number – or related data structures – were proposed for many special cases from computational geometry [17, 50, 61].

Problem 2 (*f*-APPROX STABBING NUMBER).

Input: A hypergraph $\mathcal{H} = (X, R)$ of VC-dimension at most d .

Output: A spanning path P of stabbing number at most $\tilde{O}(n^{1-\frac{1}{f(d)}})$ and, for every $q \in R$, the set $E_P(q) = \{uv \in E(P) \mid u \in q, v \notin q\}$ of all edges stabbed by q .

For simplicity of exposition, we will assume throughout the remainder of this paper that the VC-dimension of all the hypergraphs considered is part of the input. However in practice, we can easily weaken this assumption as follows. Given some “guess” d on the VC-dimension of the input, we can modify our proposed solutions so that they either output a spanning path whose stabbing number is at most $\tilde{O}(n^{1-\frac{1}{f(d)}})$, for some function f , or conclude that the VC-dimension of the input is larger than d . By dichotomic search, we so can compute some minimum d^* such that, for any $d \geq d^*$, our algorithms always output a spanning path of stabbing number $\tilde{O}(n^{1-\frac{1}{f(d)}})$. We stress that d^* is at most the VC-dimension of G , but that it can be much smaller in practice.

Reduction from diameter computation. We now recall the following simple but beautiful approach that we use in order to solve DIAMETER on graphs of constant VC-dimension.

Lemma 6. *Let G be a graph and $k \geq 2$. If the hypergraph $\mathcal{N}_{k-1}(G)$ has VC-dimension at most d , and we can solve f -APPROX STABBING NUMBER for $\mathcal{N}_{k-1}(G)$ in time $T(n, m)$, then we can decide whether G has diameter at most k in time $\tilde{O}(T(n, m) + mn^{1-\frac{1}{f(d)}})$.*

Proof. Let us first compute a spanning path P of stabbing number at most $\tilde{O}(n^{1-\frac{1}{f(d)}})$ for $\mathcal{N}_{k-1}(G)$. By the hypothesis, it takes $\mathcal{O}(T(n, m))$ time. For every $v \in V$ we can compute from $E_P(N_G^{k-1}[v])$ a set $I_{k-1}(v)$ of t_v intervals, where $|E_P(N_G^{k-1}[v])| - 1 \leq t_v \leq |E_P(N_G^{k-1}[v])| + 1$, such that $\bigcup I_{k-1}(v) = N_G^{k-1}[v]$. This preprocessing phase takes time $\mathcal{O}(|E_P(N_G^{k-1}[v])|) = \tilde{O}(n^{1-\frac{1}{f(d)}})$, and so, $\tilde{O}(n^{2-\frac{1}{f(d)}})$ total time. Then in order to decide whether $\text{diam}(G) \leq k$, we are left to decide whether for every $u \in V$ we have $\bigcup_{v \in N_G[u]} I_{k-1}(v) = V$. For that, it suffices to collect the $\tilde{O}(\text{deg}_G(u) \cdot n^{1-\frac{1}{f(d)}})$ ends of the intervals in $\bigcup_{v \in N_G[u]} I_{k-1}(v)$, and then to order them lexicographically. As a result, this last verification phase can be done in total time $\tilde{O}(mn^{1-\frac{1}{f(d)}})$. \square

Remark 1. The *radius* of a graph G is equal to $\text{rad}(G) = \min_{u \in V} \max_{v \in V} \text{dist}_G(u, v)$. Under the Hitting Set conjecture, we cannot compute the radius of a graph in truly subquadratic-time [1]. We here observe that we can easily modify the framework of Lemma 6 in order to decide whether a graph has radius at most k . Indeed, for that it suffices to check whether there exists at least one vertex u such that $\bigcup_{v \in N_G[u]} I_{k-1}(v) = V$.

Our main task in the remainder of this article will be to solve f -APPROX STABBING NUMBER efficiently on ℓ -neighbourhood hypergraphs, for some increasing function f . Then, we can apply Lemma 6 in order to efficiently solve DIAMETER.

3 Computation of Spanning paths with low Stabbing Number

We prove in this section our first main result in the paper, whose statement is reminded below:

Theorem 1. *For every $d > 0$, there exists a constant $\varepsilon_d \in (0; 1)$ such that in deterministic time $\tilde{O}(mn^{1-\varepsilon_d})$ we can decide whether a graph of VC-dimension at most d has diameter two.*

We will need the following result in our proofs:

Lemma 7 ([18]). *There is a deterministic polynomial-time algorithm that outputs, for every n -vertex hypergraph \mathcal{H} of VC-dimension at most d , a spanning path of stabbing number $\mathcal{O}(n^{1-1/2^{d+1}} \log n)$.*

This above lemma is a consequence of Theorem 4.3 in [18] and the discussion about the complexity of the algorithm resulting from their proof. We note that their result applies to infinite range spaces too, with the initial step in their proof reducing to the finite case. In order to derive Lemma 7 from [18], we use the bound on the dual VC-dimension resulting from Lemma 2 and the fact that no initial step is required as we start from a *finite* range space. Better randomized algorithms can be obtained through the approximation results in [6, 40]. They are expressed for spanning trees but easily convert to paths as previously noted. The algorithms in [6, 40] use LP relaxation and randomized rounding. It is not immediately clear if they can be derandomized using classical techniques. Indeed, the algorithm from [40] works by phases. During a phase, it needs to solve an ILP relaxation and then to apply some randomized rounding technique. In the worst case, this main phase is repeated $\mathcal{O}(\log n)$ times. We observe that even by using the best known upper-bounds on the time complexity of linear programming [21], this overall process takes super-quadratic time. In what follows, we use the Sauer-Shelah-Perles Lemma (Lemma 1) in order to obtain better trade-offs between the running-time and the quality of our solution.

Theorem 2. *For every $d > 0$, there exists a constant $\varepsilon_d \in (0; 1)$ such that in $\tilde{\mathcal{O}}(m + n^{2-\varepsilon_d})$ deterministic time, for every n -vertex hypergraph \mathcal{H} of VC-dimension at most d and size m , we can compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$.*

Moreover, $\varepsilon_d = \frac{1}{2^{d+1}[c(d+1)-1]+1}$ for some universal constant $c > 2$.

Proof. Let $\eta \in (0; 1)$ to be fixed later in the proof. We arbitrarily partition the vertex-set X into subsets X_1, X_2, \dots, X_p such that $p = \mathcal{O}(n^{1-\eta})$ and, for every $1 \leq i \leq p$, $|X_i| = \mathcal{O}(n^\eta)$. Our aim is to apply Lemma 7 to the subhypergraphs $\mathcal{H}[X_1], \mathcal{H}[X_2], \dots, \mathcal{H}[X_p]$. We stress that all these subhypergraphs can be constructed in total $\mathcal{O}(m)$ -time, as follows: we scan all the hyperedges q once in order to compute $(q \cap X_i)_{1 \leq i \leq p}$; then, for every i , we use a linear-time sorting algorithm in order to suppress duplicated values in $\{q \cap X_i \mid q \in R\}$.

Claim 1. *Given $\mathcal{H}[X_1], \mathcal{H}[X_2], \dots, \mathcal{H}[X_p]$, we can compute a spanning path for \mathcal{H} of stabbing number $\tilde{\mathcal{O}}(n^{1-\frac{\eta}{2^{d+1}}})$. Moreover, it takes $\mathcal{O}(n^{1+\eta[c(d+1)-1]})$ time for some universal constant $c > 2$.*

Proof. By Lemma 3, every $\mathcal{H}[X_i]$ has VC-dimension at most d . This implies that $\mathcal{H}[X_i]$ has $\mathcal{O}(n^{\eta d})$ hyperedges (Lemma 1), and so it has size $\mathcal{O}(n^{\eta(d+1)})$. Furthermore by Lemma 7 we can compute deterministically a spanning path of stabbing number $\tilde{\mathcal{O}}\left(n^{\eta\left(1-\frac{1}{2^{d+1}}\right)}\right)$, in time $\mathcal{O}(n^{c\eta(d+1)})$ for some universal constant c .

Let P_1, P_2, \dots, P_p be the spanning paths that we obtain. We obtain a spanning path P for \mathcal{H} by concatenating all the P_i 's. For every $1 \leq i \leq p$, we recall that the stabbing number of P_i is in $\tilde{\mathcal{O}}\left(n^{\eta\left(1-\frac{1}{2^{d+1}}\right)}\right)$. Therefore by construction, the stabbing number of P is in $\tilde{\mathcal{O}}\left(p \cdot n^{\eta\left(1-\frac{1}{2^{d+1}}\right)} + p - 1\right) = \tilde{\mathcal{O}}\left(n^{1-\frac{\eta}{2^{d+1}}}\right)$. \diamond

Let P be the spanning path obtained with Claim 1. Finally, for every $q \in R$ we compute the set $E_P(q)$ of all edges of P stabbed by q , in total $\mathcal{O}(m)$ -time, simply by scanning once all the hyperedges. The total running-time is in $\mathcal{O}(m + p \cdot n^{c\eta(d+1)}) = \mathcal{O}(m + n^{1+\eta[c(d+1)-1]})$. Overall, we achieve a good trade-off between running-time and approximation factor if we have $2 - \frac{\eta}{2^{d+1}} = 1 + \eta[c(d+1) - 1]$. Therefore we set $\eta = \frac{1}{c(d+1) + \frac{1}{2^{d+1}} - 1}$, and then $\varepsilon_d = \frac{\eta}{2^{d+1}} = \frac{1}{2^{d+1}[c(d+1)-1]+1}$. \square

We observe that our analysis could be easily improved in some particular cases, *e.g.*, for all hypergraphs that are isomorphic to their dual.

We are now ready to prove the main result in this section:

Proof of Theorem 1. We apply Theorem 2 to the closed neighbourhood hypergraph of G . Then, the result follows from Lemma 6 applied to the function $f : d \rightarrow 1/\varepsilon_d$. \square

4 Bounded DIAMETER with ε -nets

For graphs of bounded *distance* VC-dimension we now generalize Theorem 1 from the previous section to larger values for the diameter.

Theorem 3. *There exists a Monte Carlo algorithm such that, for every positive integers d and k , we can decide whether a graph of distance VC-dimension at most d has diameter at most k . The running time is in $\tilde{O}(k \cdot mn^{1-\varepsilon_d})$, where $\varepsilon_d \in (0; 1)$ only depends on d .*

Our proof crucially relies on the concept of ε -net. We recall that for a hypergraph $\mathcal{H} = (X, R)$, a subset $Y \subseteq X$ is called an ε -net if, for every $q \in R$, we have $|q| \geq \varepsilon n \implies Y \cap q \neq \emptyset$.

Lemma 8 ([42, 60]). *For every hypergraph of VC-dimension at most d , any random subset of size $\Omega(\frac{d}{\varepsilon} \log(\frac{1}{\varepsilon\delta}))$ is an ε -net with probability $1 - \delta$.*

We will also need the following result:

Lemma 9 ([18]). *For every hypergraph $\mathcal{H} = (X, R)$, let $\hat{R} := \{q\Delta q' \mid q, q' \in R\}$ be the set of symmetric differences between hyperedges. If \mathcal{H} has VC-dimension at most d then, $\hat{\mathcal{H}} := (X, \hat{R})$ has bounded VC-dimension.*

We observe that no explicit upper bound on the VC-dimension of $\hat{\mathcal{H}}$ was stated in [18]. Nevertheless it can be easily deduced from their proof that it is in $\mathcal{O}(d \log d)$ (see also [32]).

The following partition lemma which derives from the two above lemmata is the cornerstone of our algorithm.

Lemma 10. *Let $G = (V, E)$ be a graph of distance VC-dimension at most d , and let S be any random subset of size $\tilde{\Theta}(d/\varepsilon)$. Then w.h.p., for every $\ell \geq 0$ and for every $u, v \in V$ such that $N_G^\ell[u] \cap S = N_G^\ell[v] \cap S$, we have $|N_G^\ell[u] \Delta N_G^\ell[v]| = \tilde{O}(\varepsilon n)$.*

Proof. Let $\hat{R} = \{N_G^{\ell_1}[x] \Delta N_G^{\ell_2}[y] \mid x, y \in V \text{ and } \ell_1, \ell_2 \geq 0\}$ be the set of the symmetric differences between the balls of G . Since G has distance VC-dimension at most d then, by Lemma 9, the hypergraph $\hat{\mathcal{H}} = (V, \hat{R})$ has VC-dimension in $\mathcal{O}(d \log d)$. Then by Lemma 8, w.h.p. S is an ε -net for $\hat{\mathcal{H}}$. Therefore, for every $\ell \geq 0$ and for every $u, v \in V$, $|N_G^\ell[u] \Delta N_G^\ell[v]| > \varepsilon n \implies (N_G^\ell[u] \Delta N_G^\ell[v]) \cap S \neq \emptyset$. We stress that $(N_G^\ell[u] \Delta N_G^\ell[v]) \cap S \neq \emptyset \implies N_G^\ell[u] \cap S \neq N_G^\ell[v] \cap S$. \square

This above partition lemma will be useful in order to group the vertices in a small number of groups, with every two vertices in a group having almost the same ball of radius ℓ . Here there is a trade-off between the number of groups (that we upper-bound by using the Sauer-Shelah-Perles Lemma) and, for every two vertices in the same group, the maximum number of vertices in which their respective balls of radius ℓ can differ.

More precisely, our approach in the next two sections can be summarized as follows:

1. We compute a spanning path P'_k for $\mathcal{N}_k(G)$ of low *average* stabbing number, with the latter being equal to $\frac{1}{n} \cdot \sum_{v \in V} |E_{P'_k}(N_G^k[v])|$;
2. Then, we compute an ε -net, for some well-chosen ε , and in doing so we partition the vertex-set into $p(\varepsilon)$ disjoint groups $V_1, V_2, \dots, V_{p(\varepsilon)}$. For every j we select a unique $v_j \in V_j$. We restrict ourselves to $\mathcal{H}_k := (V, \{N_G^k[v_j] \mid 1 \leq j \leq p(\varepsilon)\})$. We compute a spanning path P_k of low stabbing number for this subhypergraph.
3. We observe that if P_k is a spanning path of stabbing number t for \mathcal{H}_k , then it is also a spanning path of stabbing number $t + \mathcal{O}(\varepsilon n)$ for $\mathcal{N}_k(G)$. Finally, for every $1 \leq j \leq p(\varepsilon)$, we consider the unselected vertices $u \in V_j \setminus \{v_j\}$ sequentially. We compute the set of all the edges in $E(P_k)$ that are stabbed by $N_G^k[u]$. For that, it suffices to compute the $\mathcal{O}(\varepsilon n)$ vertices of $N_G^k[u] \Delta N_G^k[v_j]$. We do so efficiently by using the auxiliary spanning path P'_k .

We next give a first application of our approach (we will give another such application in the proof of Theorem 5).

Proof of Theorem 3. Let ε_d be the constant of Theorem 2. We shall prove the following claim by finite induction:

Claim 2. *For every $1 \leq i \leq k - 1$, we can compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{N}_i(G)$. Moreover, it can be done in time $\tilde{\mathcal{O}}(i \cdot mn^{1-\varepsilon_d})$.*

The proof of Theorem 3 will follow from this claim and Lemma 6 by taking $i = k - 1$.

Proof. By Theorem 2, the claim is true for the base case $i = 1$. Assume by our induction hypothesis that the claim holds for $i - 1$. We divide the remainder of the proof into two subclaims.

Subclaim 1. *Let P_{i-1} be a spanning path of stabbing number t for $\mathcal{N}_{i-1}(G)$. We can transform P_{i-1} into a spanning path P'_i for $\mathcal{N}_i(G)$, such that $\sum_{v \in V} |E_{P'_i}(N_G^i[v])| = \mathcal{O}(tm)$. Moreover, the transformation takes time $\mathcal{O}(tm)$.*

Proof. Let $u \in V$. Then in time $\mathcal{O}(\deg_G(u) \cdot t)$, we can collect the edge-sets $E_{P_{i-1}}(N_G^{i-1}[w])$ of all the edges of P_{i-1} that are stabbed by w , for $w \in N_G^i[u]$. We compute from these edge-sets a (suboptimal) representation of $N_G^i[u]$ into $\mathcal{O}(\deg_G(u) \cdot t)$ intervals of P_{i-1} . \square

Subclaim 2. *Let P'_i be a spanning path for $\mathcal{N}_i(G)$, such that $\sum_{v \in V} |E_{P'_i}(N_G^i[v])| = \mathcal{O}(tm)$. Then, in time $\tilde{\mathcal{O}}((n^{1-\varepsilon_d} + t) \cdot m)$, we can compute a spanning path P_i of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$.*

Proof. Let $\varepsilon := \Theta(n^{-\varepsilon_d})$. We perform a breadth-first search from every vertex in some random subset S of cardinality $\tilde{\mathcal{O}}(d/\varepsilon) = \tilde{\mathcal{O}}(n^{\varepsilon_d})$. In doing so we define an equivalence relation \sim on V such that $u \sim v \iff^{def} N_G^i[u] \cap S = N_G^i[v] \cap S$. We so partition V into some groups V_1, V_2, \dots, V_p . Since by the hypothesis G has distance VC-dimension at most d then, by Lemma 1 we have $p = \mathcal{O}(|S|^d) = \tilde{\mathcal{O}}(n^{\varepsilon_d})$. Furthermore by Lemma 10, we have w.h.p. $u \sim v \implies |N_G^i[u] \Delta N_G^i[v]| = \tilde{\mathcal{O}}(\varepsilon n) = \tilde{\mathcal{O}}(n^{1-\varepsilon_d})$. The algorithm now proceeds as follows:

1. For every $1 \leq j \leq p$, we select a unique $v_j \in V_j$, and then we start a breadth-first search from this vertex. Since $p = \tilde{\mathcal{O}}(n^{\varepsilon_d})$ and we have $\varepsilon_d \ll 1/d$, this phase can be implemented in time $\tilde{\mathcal{O}}(mn^{\varepsilon_d}) = \tilde{\mathcal{O}}(mn^{1-\varepsilon_d})$, that is truly subquadratic.

2. Let $R_i := \{N_G^i[v_j] \mid 1 \leq j \leq p\}$, and let $\mathcal{H}_i := (V, R_i)$. Note that since $\mathcal{H}_i \subseteq \mathcal{B}(G)$, the VC-dimension of \mathcal{H}_i is at most d . Furthermore, the order and size of \mathcal{H}_i are, respectively, n and $m_i := \mathcal{O}(pn) = \tilde{\mathcal{O}}(n^{1+\varepsilon d})$. By Theorem 2, we can compute a spanning path P_i for \mathcal{H}_i of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon d})$ in time $\tilde{\mathcal{O}}(m_i + n^{2-\varepsilon d}) = \tilde{\mathcal{O}}(n^{1+\varepsilon d} + n^{2-\varepsilon d}) = \tilde{\mathcal{O}}(n^{1-\varepsilon d}m)$.

3. We observe that P_i is a spanning path of $\mathcal{N}_i(G)$ of stabbing number:

$$\tilde{\mathcal{O}}(n^{1-\varepsilon d}) + \max_{1 \leq j \leq p} \max_{u \in V_j \setminus \{v_j\}} |N_G^i[u] \Delta N_G^i[v_j]| = \tilde{\mathcal{O}}(n^{1-\varepsilon d}).$$

We are now left with computing, for every $1 \leq j \leq p$ and $u \in V_j \setminus \{v_j\}$, the set $E_{P_i}(N_G^i[u])$ of all the edges stabbed by the ball of radius i centered at u . For that, since we are already given $E_{P_i}(N_G^i[v_j])$, it suffices to compute $N_G^i[u] \Delta N_G^i[v_j]$. We proceed in three steps:

- By our hypothesis, we computed a spanning path P'_i for $\mathcal{N}_i(G)$, such that $\sum_{u \in V} |E(N_G^i[u])| = \mathcal{O}(tm)$. Then, we can compute from P'_i a (suboptimal) representation $I_i(u)$ of $N_G^i[u]$ into $\mathcal{O}(|E_{P'_i}(N_G^i[u])|)$ intervals. In doing so, we also compute within the same amount of time a representation $\overline{I_i(u)}$ of $V \setminus N_G^i[u]$ into $\mathcal{O}(|E_{P'_i}(N_G^i[u])|)$ intervals of P'_i . Overall this step takes total time $\tilde{\mathcal{O}}(tm)$.
- Let $\sigma_i : V \rightarrow V(P'_i)$ be the permutation that maps every vertex to its position in the spanning path P'_i . For every $1 \leq j \leq p$, we construct two balanced binary search trees whose items are, respectively, $\{\sigma_i(x) \mid x \in N_G^i[v_j]\}$ and $\{\sigma_i(y) \mid y \notin N_G^i[v_j]\}$. Overall, this takes total time $\tilde{\mathcal{O}}(np) = \tilde{\mathcal{O}}(n^{1+\varepsilon d}) = \tilde{\mathcal{O}}(mn^{1-\varepsilon d})$.
- Finally, let us again consider some $u \in V_j \setminus \{v_j\}$ for some j . For every interval from $I_i(u)$, we want to enumerate the vertices of $V \setminus N_G^i[v_j]$ that lie on this interval. Since we stored all of $V \setminus N_G^i[v_j]$ into a balanced binary search tree, this can be done in time $\mathcal{O}(\log n)$ plus $\mathcal{O}(1)$ extra time per solution. In the same way, for every interval from $\overline{I_i(u)}$, we enumerate the vertices of $N_G^i[v_j]$ that lie on this interval. For a fixed u , the total time for this step is in $\tilde{\mathcal{O}}(|I_i(u)| + |\overline{I_i(u)}| + |N_G^i[u] \Delta N_G^i[v_j]|) = \tilde{\mathcal{O}}(|E_{P'_i}(N_G^i[u])| + n^{1-\varepsilon d})$. Therefore, this last step takes total time $\tilde{\mathcal{O}}(tm + n^{2-\varepsilon d})$.

◦

Now, by the induction hypothesis we get a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon d})$ for $\mathcal{N}_{i-1}(G)$. By Subclaim 1 we transform such spanning path into a spanning path P'_i for $\mathcal{N}_i(G)$, where $\sum_{u \in V} |E_{P'_i}(N_G^i[u])| = \tilde{\mathcal{O}}(mn^{1-\varepsilon d})$. Finally, by Subclaim 2 we can use P'_i in order to compute, in time $\tilde{\mathcal{O}}(mn^{1-\varepsilon d})$, a spanning path P_i of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon d})$. The above algorithm achieves proving that our claim holds for i . ◊

Summarizing, by Claim 2 we can compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon d})$ for the hypergraph $\mathcal{N}_{k-1}(G)$, in time $\tilde{\mathcal{O}}(k \cdot mn^{1-\varepsilon d})$. By Lemma 6 it implies that we can also decide whether G has diameter at most k , and if so compute $\text{diam}(G)$ exactly, in time $\tilde{\mathcal{O}}(k \cdot mn^{1-\varepsilon d})$. ◻

4.1 Application to nowhere dense graph classes

A closer look at the proof of Theorem 3 shows that it also holds if, instead of having bounded distance VC-dimension, there rather exists some constant d such that, for every $1 \leq i \leq k-1$,

the VC-dimension of the i -neighbourhood hypergraph is at most d (the latter value is sometimes called the distance- i VC-dimension of the graph [53]). It has algorithmic implications for some special cases of sparse graphs. Namely, H is an r -shallow minor of a graph G if it can be obtained from some subgraph of G by the contraction of pairwise disjoint subgraphs of radius at most r [56]; a graph family \mathcal{G} is termed *nowhere dense* if, for any r , there exists a graph H_r which is not an r -shallow minor for any graph in \mathcal{G} [52]. Of interest here is that, for any graph class \mathcal{G} nowhere dense, and for any i , the distance- i VC-dimension of any graph in \mathcal{G} is upper-bounded by some constant d_i [53]. By choosing $d := \max_{1 \leq i \leq k-1} d_i$, we thus obtain the following weaker version of Theorem 3 for nowhere dense graphs:

Theorem 4. *Let \mathcal{G} be a class of nowhere dense graphs. There exists a Monte Carlo algorithm such that, for every constant $k = \mathcal{O}(1)$, for any graph in \mathcal{G} we can decide whether its diameter is at most k in time $\tilde{\mathcal{O}}(mn^{1-\varepsilon_{\mathcal{G}}(k)})$, for some constant $\varepsilon_{\mathcal{G}}(k) \in (0; 1)$ that only depends on k .*

We left open whether there exists a truly subquadratic-time FPT algorithm for diameter computation on nowhere dense graph classes (i.e., with no dependency on k in the exponent).

5 Diameter computation in truly Subquadratic time

We finally improve the results of Theorem 3 for a more restricted family of graphs of bounded distance VC-dimension. Before that, we need to introduce a bit more of graph terminology. A class of graphs is called *monotone* if it is closed by taking subgraphs. For a connected n -vertex graph G , a *separator* is a subset S such that $G \setminus S$ is disconnected. It is called *balanced* if every connected component of $G \setminus S$ has order at most $2n/3$. Finally, a class of graphs has *strongly sublinear balanced separator* if every connected n -vertex graph in the class has a balanced separator of cardinality at most $C \cdot n^\alpha$ for some constants C and $\alpha < 1$.

Theorem 5. *Let \mathcal{G} be a monotone graph class with strongly sublinear balanced separators. Then there exists a Monte Carlo algorithm such that, for every $d > 0$, we can compute the diameter of any graph in \mathcal{G} of distance VC-dimension at most d in time $\tilde{\mathcal{O}}(n^{2-\varepsilon_{\mathcal{G}}(d)})$, for some constant $\varepsilon_{\mathcal{G}}(d) \in (0; 1)$ that only depends on d .*

We postpone the technical proof of this result to Sec. 5.2. Let us emphasize that Theorem 5 cannot be applied to *all* graph classes of bounded distance VC-dimension. For instance, we proved in Lemma 4 that the intervals graphs have distance VC-dimension at most two. However, there exist intervals graphs with no balanced separators of sublinear size. We give some interesting cases where Theorem 5 *does* apply in Sec. 5.1.

Finally, we say that a class of graphs \mathcal{G} has *polynomial expansion* if there exists a polynomial p such that, for every r -shallow minor of a graph in \mathcal{G} (cf. Section 4.1), the average degree is at most $p(r)$. We want to stress that there is an equivalence between the monotone classes of graphs \mathcal{G} with strongly sublinear balanced separators and those of polynomial expansion [30]. In particular, the graphs in \mathcal{G} have bounded degeneracy, and so, they are sparse (i.e., with $m = \mathcal{O}(n)$ edges). We will often use this property in what follows.

5.1 Application to H -minor free graphs

Let us now review some interesting classes where Theorem 5 does apply. Since planar graphs have distance VC-dimension at most four [10] then, it follows from the planar separator theorem of

Lipton and Tarjan [49] that it is the case for planar graphs. Therefore, Theorem 5 gives us a new subquadratic-time algorithm for diameter computation on *unweighted* planar graphs, but with a slower running-time than for the algorithms presented in [15, 38]. More generally, the following separator theorem is from Alon et al.:

Lemma 11 ([2]). *Every K_h -minor free graph has a balanced separator of cardinality $\mathcal{O}(h^{3/2}\sqrt{n})$. Moreover, such a separator can be found in $\mathcal{O}(n^{3/2})$ time.*

See also [45, 62] for various trade-offs between the size of the separator and the time that is needed in order to find it. We recall that K_h -minor free graphs have distance VC-dimension at most $h - 1$ [10]. By combining this result with Lemma 11, we so prove the following meta-theorem:

Corollary 1. *The diameter of a H -minor free graph can be computed in time $\tilde{\mathcal{O}}(n^{2-\varepsilon_H})$, with a Monte Carlo algorithm, where $\varepsilon_H \in (0; 1)$ is a constant that only depends on H .*

For most values of H this is the first known subquadratic-time algorithm for diameter computation on H -minor free graphs. In particular, this is the first known subquadratic algorithm for diameter computation on (unweighted) bounded-genus graphs to the best of our knowledge (see the planar graphs paragraph in the introduction).

5.2 Proof of Theorem 5

The remainder of this section is devoted to the proof of Theorem 5. We start by presenting, in a separate subsection, all the required background on r -divisions.

Algorithmic aspects of r -divisions

Throughout all this section, let $\mathcal{G}_{\alpha, C}$ be the class of all the graphs G such that, for every connected h -vertex subgraph of G , there exists a balanced separator of order at most $C \cdot h^\alpha$. The following intermediate result is built upon a previous algorithm from Plotkin et al. [56].

Lemma 12 ([29]). *For every n -vertex m -edge graph $G \in \mathcal{G}_{\alpha, C}$, we can find a balanced separator of order $\mathcal{O}(n^{\frac{4+\alpha}{5}})$ in time $\mathcal{O}(mn^{\frac{4+\alpha}{5}}) = \mathcal{O}(n^{2-\frac{1-\alpha}{5}})$.*

We will also use the following simple result:

Lemma 13. *Let G be a graph and S a balanced separator. We can bipartition the connected components of $G \setminus S$ in two disjoint sets A and B such that $\min\{|A|, |B|\} \leq 2n/3$.*

Proof. Let C_1, C_2, \dots, C_k be the connected components of $G \setminus S$. We define $i_0 := \max\{i \mid |\bigcup_{j < i} C_j| \leq 2n/3\}$. Let $A' := \bigcup_{j < i_0} C_j$ and $B' := \bigcup_{j > i_0} C_j$. If $|B' \cup C_{i_0}| \leq 2n/3$ then we are done by setting $A := A'$, $B := B' \cup C_{i_0}$. Thus, from now on let us assume that $|B' \cup C_{i_0}| > 2n/3$. Note that since S is a balanced separator, it implies that $i_0 < k$. Then, by the very definition of i_0 we also have $|A' \cup C_{i_0}| > 2n/3$. Overall, $|A'| + 2|C_{i_0}| + |B'| > 4n/3$. Since $|A'| + |B'| + |C_{i_0}| < n$, we obtain $|C_{i_0}| > n/3$. We are done by setting $A := A' \cup B'$ and $B := C_{i_0}$. \square

Now, set $\beta := \frac{4+\alpha}{5} < 1^3$. By Lemma 12, for every n -vertex m -edge graph in $\mathcal{G}_{\alpha, C}$ we can compute a balanced separator of order $\mathcal{O}(n^\beta)$ in time $\mathcal{O}(n^{1+\beta})$. Following Federickson [37], we define an r -division for an n -vertex graph $G \in \mathcal{G}_{\alpha, C}$ as follows:

³More generally, let $\mathcal{G} \subseteq \mathcal{G}_{\alpha, C}$. We may choose any parameter $\beta \in [\alpha; 1)$ such that for all the graphs in \mathcal{G} we can compute a balanced separator of size $\mathcal{O}(n^\beta)$ in truly subquadratic-time. For instance by Lemma 11, if \mathcal{G} is proper minor-closed then we can set $\beta = \alpha = 1/2$.

- If $n \leq r$ then, we output G ;
- Otherwise, let S be a balanced separator of cardinality $\mathcal{O}(n^\beta)$. Since S is balanced then, by Lemma 13 we can partition the connected components of $G \setminus S$ in two disjoint sets A and B of cardinality $\leq 2n/3$. We end up computing an r -division for the induced subgraphs $G[A \cup S]$ and $G[B \cup S]$ separately. Note that since S is a separator, all edges of G are covered by these two subgraphs.

Therefore by construction, an r -division of a connected graph G is a collection of connected induced subgraphs of order at most r that cover all edges of G . We will use the terminology from [41]. In particular, the subgraphs in an r -division are termed *clusters*. A vertex is *interior* if it is contained in a unique cluster, otherwise it is a *boundary* vertex. Finally, if the sum of the orders of all the clusters is $n + q$ then, we call q the *excess*.

The following result is essentially a reformulation of [41, Lemma 2.2].

Lemma 14 ([41]). *Set $\beta := \frac{4+\alpha}{5}$. There exists a constant r_0 such that, for any n -vertex graph $G \in \mathcal{G}_{\alpha,C}$ and $r \geq r_0$, any r -division of G has an excess in $\mathcal{O}(n/r^{1-\beta})$.*

Note that in our applications, we will choose $r = n^\gamma$ for some $\gamma \in (0; 1)$ that only depends on β and on the distance VC-dimension.

It is easy to prove that an r -division can be computed in polynomial time [41]. Next we use the known connections between strongly sublinear separators and *polynomial expansion* [29] in order to bound the running-time by some truly subquadratic function.

Corollary 2. *Set $\beta := \frac{4+\alpha}{5}$. Then, for any n -vertex m -edge graph $G \in \mathcal{G}_{\alpha,C}$, we can compute an r -division in time $\tilde{\mathcal{O}}(n^{1+\beta})$.*

Proof. Let us assume that at the initialization step, $n > r$ (otherwise, we are done). We claim that it is sufficient to prove that the total number of edges in the final clusters is in $\mathcal{O}(n)$. Indeed, if this is true for the final clusters then, this is also true for the intermediate clusters at any given step of the decomposition. In particular, every step runs in time $\mathcal{O}(n^{1+\beta})$. Furthermore, since we only consider balanced separators of *sublinear* cardinality, for every n above some constant the two induced subgraphs constructed have truly sublinear order (say, $\leq 3n/4$). Therefore it takes $\mathcal{O}(\log n)$ steps to decrease the order of all the subgraphs in this collection to less than r . This upper-bound on the number of steps proves, as claimed, that the total running time is in $\tilde{\mathcal{O}}(n^{1+\beta})$.

We are left proving that the total number of edges in the final clusters is indeed in $\mathcal{O}(n)$. For that, let us consider any of the clusters W_i . Since $\mathcal{G}_{\alpha,C}$ is monotone, we have $W_i \in \mathcal{G}_{\alpha,C}$. Furthermore, every graph in $\mathcal{G}_{\alpha,C}$ must be $\mathcal{O}(1)$ -degenerate (e.g., see [29, Lemma 2 (b)] where the author proved a stronger result, namely that $\mathcal{G}_{\alpha,C}$ has polynomial expansion). It implies that W_i has size $\mathcal{O}(|V(W_i)|)$. Overall, if the total excess is q then, the total number of edges in the clusters is in $\mathcal{O}(n + q)$. By Lemma 14 we have $q = \mathcal{O}(n)$, and so the total number of edges is also in $\mathcal{O}(n)$. \square

Boundary Hypergraphs

Let G be a graph equipped with some r -division, and let ℓ be a positive integer. Roughly, our objective is to use the r -division in order to compute, for every vertex, a compact interval representation of its balls of radius ℓ . This leads us to the following natural object:

Definition 7. Let Λ_r be an r -division of a graph G , and let ℓ be a positive integer. The ℓ -boundary hypergraph $\mathcal{H}_{\ell,G}(\Lambda_r)$ has for vertex-set V . Moreover, for every cluster $W_i \in \Lambda_r$ and for every $u, v \in V(W_i)$, if v is a boundary vertex and $\text{dist}_G(u, v) < \ell$, then the ball $N_G^{\ell - \text{dist}_G(u, v)}[v]$ is a hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$.

To better understand this above construction, let W_i be a cluster, let $u \in V(W_i)$ be internal and let $z \notin V(W_i)$. Then, since an r -division is also an edge-covering, we have $\text{dist}_G(u, z) \leq \ell$ if and only if there exists a boundary vertex $v \in V(W_i)$ such that $\text{dist}_G(u, v) + \text{dist}_G(v, z) \leq \ell$. Equivalently, we have $z \in N_G^{\ell - \text{dist}_G(u, v)}[v]$.

Lemma 15. Set $\beta := \frac{4+\alpha}{5}$. Then, for any n -vertex graph $G \in \mathcal{G}_{\alpha,C}$, and for any r -division Λ_r , the ℓ -boundary hypergraph $\mathcal{H}_{\ell,G}(\Lambda_r)$ has $\mathcal{O}(nr^\beta)$ hyperedges.

Proof. For every $W_i \in \Lambda_r$, we create $\mathcal{O}(r \cdot b_i)$ hyperedges, where b_i denotes the number of boundary vertices in the cluster. We observe that $\sum_{W_i \in \Lambda_r} b_i$ is at most twice the excess. Then, by Lemma 14 we have $\mathcal{O}(r) \times \mathcal{O}(n/r^{1-\beta}) = \mathcal{O}(nr^\beta)$ hyperedges. \square

We stress that by Lemma 15, a boundary hypergraph may have a superlinear number of edges. Therefore, if we restrict ourselves to subquadratic-time computation, we cannot compute this hypergraph explicitly. Fortunately, we show next that this is not needed if one just wants to compute for this hypergraph a spanning path of low stabbing number.

Lemma 16. Set $\beta := \frac{4+\alpha}{5}$, and let $G \in \mathcal{G}_{\alpha,C}$ have distance VC-dimension at most d . Then, there exists a constant $\varepsilon_d \in (0, 1)$ that only depends on d and such that, for any r -division Λ_r , the stabbing number of $\mathcal{H}_{\ell,G}(\Lambda_r)$ is in $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$. Moreover, we can compute a spanning path reaching this upper bound in deterministic time $\tilde{\mathcal{O}}(n^2/r^{1-\beta} + n^{2-\varepsilon_d}r^\beta)$.

Proof. By construction, $\mathcal{H}_{\ell,G}(\Lambda_r)$ is a subhypergraph of $\mathcal{B}(G)$, the ball hypergraph of G . Therefore, the VC-dimension of $\mathcal{H}_{\ell,G}(\Lambda_r)$ is at most d . Let ε_d be the constant of Theorem 2. In order to prove the result, we are left proving that we can adapt the algorithm of Theorem 2 so that it runs in time $\tilde{\mathcal{O}}(nm/r^{1-\beta} + n^{2-\varepsilon_d}r^\beta)$ when it is given $\mathcal{H}_{\ell,G}(\Lambda_r)$ as input. For that, let F be the set of the boundary vertices. We have that $|F|$ is at most twice the excess, and so, by Lemma 14 we get $|F| = \mathcal{O}(n/r^{1-\beta})$.

1. We start with a breadth-first search from every vertex of F . This pre-processing phase takes time $\mathcal{O}(|F|m) = \mathcal{O}(n^2/r^{1-\beta})$. Furthermore, note that in doing so we computed all the pairs $(v, t) \in F \times [m]$ such that $N_G^t[v]$ is a hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$.
2. Let $\eta = 2^{d+1}\varepsilon_d$. We arbitrarily partition the vertex-set V into subsets V_1, V_2, \dots, V_p such that $p = \mathcal{O}(n^{1-\eta})$ and, for every $1 \leq i \leq p$, $|V_i| = \mathcal{O}(n^\eta)$. Furthermore, as explained in the proof of Theorem 2 (i.e., Claim 1), we can compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{H}_{\ell,G}(\Lambda_r)$ if we are given the subhypergraphs $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_p$ that are induced by V_1, V_2, \dots, V_p respectively. It takes time $\tilde{\mathcal{O}}(n^{1+\eta(c(d+1)-1)})$ for some constant c , that is in $\tilde{\mathcal{O}}(n^{2-\varepsilon_d})$.

In order to compute all the subhypergraphs \mathcal{H}_i , we could proceed by brute-force, as follows. For every i and for any boundary vertex v , we read the vertices of V_i by non-decreasing distance to v . Furthermore, if $N_G^t[v]$ is a hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$, then as soon as we exceed

distance t all the vertices read so far are exactly $N_G^t[v] \cap V_i$. Overall, for a fixed boundary vertex v we could obtain this way up to $\mathcal{O}(|V_i|)$ different subsets of order $\mathcal{O}(|V_i|)$ each. But unfortunately, that would give us a time complexity in $\mathcal{O}(|F||V_i|^2) = \mathcal{O}(n^{1+2\eta}/r^{1-\beta})$ for a given i , and so a total running time in $\mathcal{O}(n^{2+\eta}/r^{1-\beta})$. In order to lower this running-time, we proceed as follows.

- (a) For every $v \in F$, we group all the vertices in V_i at equal distance to v . We totally order this partition by increasing distance of its vertices to v . Doing so we get exactly $n_i := |V_i|$ ordered groups (possibly, by adding some empty groups in the sequence), denoted $V_i^1(v), V_i^2(v), \dots, V_i^{n_i}(v)$. Overall, this phase takes time $\mathcal{O}(|F||V_i|) = \mathcal{O}(n^{1+\eta}/r^{1-\beta})$.
- (b) Then, we introduce a complex subprocedure in order to gradually remove the duplicates from the sets $N_G^t[v] \cap V_i$, for $v \in F$ and $t \geq 0$. For every $j = 0 \dots n_i$, we map every boundary vertex v to $\bigcup_{j' \leq j} V_i^{j'}(v)$. More precisely, we maintain some collection of different subsets of V_i , denoted $\mathcal{P}_j = (V_i^{j,1}, V_i^{j,2}, \dots, V_i^{j,s_i(j)})$ (note that \mathcal{P}_j is a list of lists). For every $v \in F$ we ensure that there is a unique t such that $V_i^{j,t} = \bigcup_{j' \leq j} V_i^{j'}(v)$. Then, there is a pointer from vertex v to this t^{th} subset (equivalently, for every list in \mathcal{P}_j , we store an auxiliary list of all the corresponding vertices of F).

We will show next that it is easy to construct \mathcal{P}_{j+1} from \mathcal{P}_j , but that the natural method for doing so might generate some duplicates. Roughly, by using in our analysis the Sauer-Shelah-Perles lemma, we prove that it is more efficient to remove duplicates at every single step rather than doing it only once at the end of the subprocedure.

We observe that initially for $j = 0$, there is a unique subset $V_i^{0,1} = \emptyset$. Furthermore if all the subsets $V_i^{j,t}$ have been computed at step j , then we can compute those at step $j + 1$, as follows:

- For every $v \in F$, if we have $V_i^{j,t} = \bigcup_{j' \leq j} V_i^{j'}(v)$, then we add a copy of $V_i^{j+1}(v)$ into some buffer $b'_{j+1}(t)$ and a pointer from v to this copy. It takes time $\mathcal{O}(\sum_{v \in F} |V_i^{j+1}(v)|)$.
- Then, for every $1 \leq t \leq s_i(j)$, we remove all the duplicated subsets in the buffer $b'_{j+1}(t)$. The new buffer that we get is denoted $b_{j+1}(t)$. We can compute it by using partition refinement (*e.g.*, see [39]), that takes time $\mathcal{O}(\sum_{W \in b'_{j+1}(t)} |W|)$ up to some $\mathcal{O}(|V_i|)$ -time pre-processing. Overall the removal of all the duplicates, for all t , takes total time $\mathcal{O}(n^\eta + \sum_{v \in F} |V_i^{j+1}(v)|)$. Furthermore on our way to remove the duplicates, we also need to actualize the pointers between the boundary vertices and the buffer contents, that takes additional time $\mathcal{O}(|F|) = \mathcal{O}(n/r^{1-\beta})$.
- For every $1 \leq t \leq s_i(j)$, we can now refine $V_i^{j,t}$ in $|b_{j+1}(t)|$ new subsets. Every such subset has order $\mathcal{O}(n^\eta)$, and so this operation takes total time $\mathcal{O}(n^\eta |b_{j+1}(t)|)$. Overall, we obtain a new collection of $\mathcal{O}(\sum_t |b_{j+1}(t)|)$ subsets. Furthermore, on our way to construct this collection, we can add a pointer from every boundary vertex v to *one* subset equal to $\bigcup_{j' \leq j+1} V_i^{j'}(v)$ (there may be duplicated subsets). By carefully using the pointers added between the boundary vertices and the buffer contents during the previous phases, this operation takes additional time $\mathcal{O}(|F|) = \mathcal{O}(n/r^{1-\beta})$.
- Finally, since all the subsets in the new collection have order $\mathcal{O}(n^\eta)$, by using again partition refinement we can merge all the duplicated subsets in time $\mathcal{O}(|V_i| + n^\eta)$.

$\sum_t |b_{j+1}(t)| = \mathcal{O}(n^\eta \cdot \sum_t |b_{j+1}(t)|)$. We also need to actualize the pointers between the boundary vertices and the subsets, that takes total time $\mathcal{O}(|F|) = \mathcal{O}(n/r^{1-\beta})$.

Let us upper bound $s_i(j)$. For that we stress that every subset $V_i^{j,t}$ represents a different intersection of V_i with a ball of G , hence of a hyperedge of $\mathcal{B}(G)$. Since $\mathcal{B}(G)$ has VC-dimension at most d , by Lemma 3 so does its subhypergraph \mathcal{H}'_i induced by V_i . In particular, every $V_i^{j,t}$ is a hyperedge of \mathcal{H}'_i . By Lemma 1 we get that $s_i(j) = \mathcal{O}(n^{\eta d})$. In the same way, since for a fixed t the $|b_{j+1}(t)|$ new subsets that are obtained by refinement of $V_i^{j,t}$ are pairwise different, we have $|b_{j+1}(t)| \leq s_i(j+1) = \mathcal{O}(n^{\eta d})$. As a result, the passing from step j to step $j+1$ takes time:

$$\mathcal{O} \left(\left[\sum_{v \in F} |V_i^{j+1}(t)| \right] + n/r^{1-\beta} + n^\eta \cdot n^{\eta d} \cdot n^{\eta d} \right) = \mathcal{O} \left(\left[\sum_{v \in F} |V_i^{j+1}(t)| \right] + n/r^{1-\beta} + n^{(2d+1)\eta} \right).$$

There are $\mathcal{O}(n^\eta)$ loops, that gives us a total running time of:

$$\begin{aligned} \mathcal{O} \left(\left[\sum_{v \in F} \sum_{j=1}^{n_i} |V_i^{j+1}(t)| \right] + n^{1+\eta}/r^{1-\beta} + n^{(2d+2)\eta} \right) &= \mathcal{O} \left(\left[\sum_{v \in F} n^\eta \right] + n^{1+\eta}/r^{1-\beta} + n^{2(d+1)\eta} \right) \\ &= \mathcal{O} \left(n^{1+\eta}/r^{1-\beta} + n^{2(d+1)\eta} \right). \end{aligned}$$

- (c) Here the key observation is that $\bigcup_j \mathcal{P}_j$ contains the intersection with V_i of *all* the balls whose center is in F . We so computed a *superset* of order $\mathcal{O}(n^{(d+1)\eta})$ (i.e., $\mathcal{O}(n^{\eta d})$ per loop) that contains all possible intersections between a hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$ and V_i . Since every subset in $\bigcup_j \mathcal{P}_j$ represents the intersection of a hyperedge of $\mathcal{B}(G)$ with V_i , and furthermore $\mathcal{B}(G)$ has VC-dimension at most d , then for simplicity we may replace \mathcal{H}_i by the slightly larger hypergraph \mathcal{H}'_i of which these are the hyperedges (i.e., the hyperedges of \mathcal{H}'_i are the intersections of V_i with all the balls whose center is in F). Note that in order to compute \mathcal{H}'_i , it is sufficient to eliminate all the duplicated elements in this collection $\bigcup_j \mathcal{P}_j$, that takes total time $\mathcal{O}(n^{(d+2)\eta})$.

The running-time is in $\tilde{\mathcal{O}}(n^{1+\eta}/r^{1-\beta} + n^{2(d+1)\eta})$ for any fixed i . Therefore, the total running-time is in $\tilde{\mathcal{O}}(n^2/r^{1-\beta} + n^{1+[2(d+1)-1]\eta})$. Recall (see Theorem 2 and its proof) that we have $\tilde{\mathcal{O}}(n^{1+\eta[c(d+1)-1]}) = \tilde{\mathcal{O}}(n^{2-\varepsilon_d})$ for some constant $c > 2$. As a result, the running-time of this part is also in $\tilde{\mathcal{O}}(n^2/r^{1-\beta} + n^{2-\varepsilon_d})$.

3. By continuing the algorithm of Theorem 2 with the hypergraphs $\mathcal{H}'_1, \mathcal{H}'_2, \dots, \mathcal{H}'_p$, we get a spanning path of $\mathcal{H}_{\ell,G}(\Lambda_r)$ whose stabbing number is in $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$. It remains to compute, for every hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$, the set of the stabbed edges. For that, let $v \in F$ be fixed. We add all the radii t such that $N_G^t[v]$ is a hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$ in a balanced binary search tree T_v . Then, we scan all the edges xy of the spanning path. By symmetry let us assume that $\text{dist}_G(v, x) \leq \text{dist}_G(v, y)$. The edge xy is stabbed by all the hyperedges $N_G^t[v]$ such that $\text{dist}_G(v, x) \leq t < \text{dist}_G(v, y)$. Then by using T_v , after some pre-computation in time $\mathcal{O}(\log n)$ every value t in the range $[\text{dist}_G(v, x); \text{dist}_G(v, y))$ can be enumerated in constant-time.

Overall, by Lemma 15 there are $\mathcal{O}(nr^\beta)$ edges, and so the construction of all the balanced binary search trees takes time $\tilde{\mathcal{O}}(nr^\beta)$. Scanning all the edges, for every boundary vertex,

takes total time $\tilde{\mathcal{O}}(n^2/r^{1-\beta})$. Any other operation corresponds to an edge of the spanning path that is stabbed by a hyperedge of $\mathcal{H}_{\ell,G}(\Lambda_r)$, and as a result there can only be $\tilde{\mathcal{O}}(n^{1-\varepsilon_d}) \times \mathcal{O}(nr^\beta) = \tilde{\mathcal{O}}(n^{2-\varepsilon_d}r^\beta)$ such operations.

Altogether combined, the running time of the algorithm is in $\tilde{\mathcal{O}}(n^2/r^{1-\beta} + n^{2-\varepsilon_d}r^\beta)$. \square

The algorithm

We are now ready to prove the main result of this section.

Proof of Theorem 5. By a classical dichotomic argument it is sufficient to prove that for any k , we can decide whether $\text{diam}(G) \leq k$ in truly subquadratic time. Furthermore in order to solve this decision problem, by Lemma 6 we are left with computing a spanning path of strongly sublinear stabbing number for the $(k-1)$ -neighbourhood hypergraph. For that, let $b_0b_1 \dots b_{s-1}b_s$ be the binary decomposition of $k-1$ (from the most significant to the less significant bit). Furthermore, for every $0 \leq i \leq s$, let k_i be the positive integer of binary decomposition $b_0b_1 \dots b_{i-1}b_i$. We will prove by finite induction that one can compute a spanning path of strongly sublinear stabbing number for the k_i -neighbourhood hypergraph of G . More precisely, let ε_d be the constant of Theorem 2. We will compute such spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ in $\tilde{\mathcal{O}}(n^{2-\varepsilon'})$ time for some $\varepsilon' > 0$. Note that since $s = \mathcal{O}(\log k)$, that will indeed give us a truly subquadratic algorithm for deciding whether $\text{diam}(G) \leq k$.

If $i = 0$, then $b_0 = 1$ and the result follows from Theorems 1 and 2. Thus from now on assume $i > 0$. We observe that $k_i = 2k_{i-1} + b_i$. Furthermore if $b_i = 1$, then we explained in the proof of Theorem 3 (Subclaims 1 and 2) how to compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{N}_{k_i}(G)$ from such spanning path for $\mathcal{N}_{k_{i-1}}(G) = \mathcal{N}_{2k_{i-1}}(G)$; it takes $\tilde{\mathcal{O}}(mn^{1-\varepsilon_d}) = \tilde{\mathcal{O}}(n^{2-\varepsilon_d})$ time, that is truly subquadratic. In order to complete the proof, we now adapt this algorithm of Theorem 3 so as to compute a spanning path with stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{N}_{2k_{i-1}}(G)$ from such spanning path for $\mathcal{N}_{k_{i-1}}(G)$. Specifically, let C and $\alpha < 1$ be such that $\mathcal{G} \subseteq \mathcal{G}_{\alpha,C}$ and set $\beta := \frac{4+\alpha}{5} < 1$. The running time for computing the adequate spanning path will be $\tilde{\mathcal{O}}\left(n^{2-\left(1-\frac{1}{2-\beta}\right)\varepsilon_d}\right)$. We first prove the following intermediate result for any value $r > 0$ (we will use $r = \Theta(n^{\frac{\varepsilon_d}{2-\beta}})$):

Claim 3. *Let P_{i-1} be a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{N}_{k_{i-1}}(G)$. In $\tilde{\mathcal{O}}(n^{2-\varepsilon_d}r + n^2/r^{1-\beta})$ time, we can compute from P_{i-1} a spanning path P'_i for $\mathcal{N}_{2k_{i-1}}(G)$, such that $\sum_{v \in V} |E_{P'_i}(N_G^{2k_{i-1}}[v])| = \tilde{\mathcal{O}}(n(r + n^{1-\varepsilon_d}r^\beta))$.*

Proof. By Lemma 2 we can compute an r -division, denoted Λ_r , in time $\tilde{\mathcal{O}}(n^{1+\beta}) = \tilde{\mathcal{O}}(n^2/r^{1-\beta})$. Then, we proceed as follows.

1. We first consider all the clusters $W \in \Lambda_r$ sequentially. For every $x, y \in W$ we have $\text{dist}_G(x, y) \leq 2k_{i-1}$ if and only if we have $N_G^{k_{i-1}}[x] \cap N_G^{k_{i-1}}[y] \neq \emptyset$. Given a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{N}_{k_{i-1}}(G)$, this test can be easily done in time $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ (i.e., by sorting the ends of the $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ intervals on this spanning path that cover the balls $N_G^{k_{i-1}}[x]$ and $N_G^{k_{i-1}}[y]$). Furthermore by Lemma 14 we have $\sum_{W \in \Lambda_r} |W| = \Theta(n)$, and so this step takes time $\tilde{\mathcal{O}}(n^{1-\varepsilon_d}) \times \sum_{W \in \Lambda_r} \mathcal{O}(|W|^2) = \tilde{\mathcal{O}}(rn^{1-\varepsilon_d}) \times \sum_{W \in \Lambda_r} |W| = \tilde{\mathcal{O}}(rn^{2-\varepsilon_d})$. Overall for every $u \in V$, we computed all the vertices of $N_G^{2k_{i-1}}[u]$ that belong to a common cluster with u .

2. Let us now consider the $2k_{i-1}$ -boundary hypergraph $\mathcal{H}_{2k_{i-1},G}(\Lambda_r)$. By Lemma 16 we can compute a spanning path P'_i of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for this hypergraph, in time $\tilde{\mathcal{O}}(nm/r^{1-\beta} + n^{2-\varepsilon_d}r^\beta) = \tilde{\mathcal{O}}(n^2/r^{1-\beta} + n^{2-\varepsilon_d}r^\beta)$. Let $u \in V$. There are two cases:

- *Case u is a boundary vertex.* Since $N_G^{2k_{i-1}}[u]$ is a hyperedge of the boundary hypergraph, we have $|E_{P'_i}(N_G^{2k_{i-1}}[u])| = \tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ (already computed).
- *Case u is an internal vertex.* Let $W \in \Lambda_r$ be the unique cluster containing u , and set initially $E_{P'_i}(N_G^{2k_{i-1}}[u]) := \emptyset$. For every boundary vertex $v \in V(W)$, if $\text{dist}_G(u, v) < 2k_{i-1}$ then, we add all of $E_{P'_i}(N_G^{2k_{i-1}-\text{dist}_G(u,v)}[v])$ to $E_{P'_i}(N_G^{2k_{i-1}}[u])$. Assuming there are b_W boundary vertices in W , we obtain that $|E_{P'_i}(N_G^{2k_{i-1}}[u])| = \tilde{\mathcal{O}}(b_W \cdot n^{1-\varepsilon_d})$. Furthermore, this above set of stabbed edges defines a collection of intervals that covers exactly the balls $N_G^{2k_{i-1}-\text{dist}_G(u,v)}[v]$, for the boundary vertices $v \in V(W)$. Denote this collection of intervals by $I_i(u)$, and its complementary by $\overline{I_i(u)}$. By construction, every vertex on an interval of $I_i(u)$ is at a distance $\leq 2k_{i-1}$ to u ; conversely, since Λ_r is also an edge-covering, every vertex of $N_G^{2k_{i-1}}[u] \setminus W$ must be on one of these intervals. As a result, in order to construct $E_{P'_i}(N_G^{2k_{i-1}}[u])$, it suffices to update this set using the vertices of $N_G^{2k_{i-1}}[u] \cap W$ that lie on some interval of $\overline{I_i(u)}$. Note that in doing so, we can only modify the cardinality of $E_{P'_i}(N_G^{2k_{i-1}}[u])$ by an $\mathcal{O}(|W|) = \mathcal{O}(r)$.

Overall, we obtain that $\sum_{u \in V} |E_{P'_i}(N_G^{2k_{i-1}}[u])| = \tilde{\mathcal{O}}(nr + n^{1-\varepsilon_d} \cdot \sum_W (b_W \cdot |V(W)|)) = \tilde{\mathcal{O}}(nr + rn^{1-\varepsilon_d} \cdot \sum_W b_W)$. Again we observe that $\sum_W b_W$ is at most twice the excess, and so by Lemma 14 $\sum_W b_W = \mathcal{O}(n/r^{1-\beta})$. Therefore, $\sum_{u \in V} |E_{P'_i}(N_G^{2k_{i-1}}[u])| = \tilde{\mathcal{O}}(n(r + n^{1-\varepsilon_d}r^\beta))$.

◇

By combining Claim 3 with Subclaim 2, in $\tilde{\mathcal{O}}(n^{2-\varepsilon_d}r + n^2/r^{1-\beta} + n^{1-\varepsilon_d}m)$ time we compute a spanning path of stabbing number $\tilde{\mathcal{O}}(n^{1-\varepsilon_d})$ for $\mathcal{N}_{2k_{i-1}}(G)$, thereby completing the description of our algorithm. Overall, the running-time of our algorithm is optimized when we have $n^2/r^{1-\beta} = n^{2-\varepsilon_d}r$. As a result, a good choice is $r = \Theta(n^{\frac{\varepsilon_d}{2-\beta}})$. Finally, we stress that in this case, $n^{2-\varepsilon_d}r = n^{2-(1-\frac{1}{2-\beta})\varepsilon_d}$ is truly subquadratic, because $\beta < 1$ and so $1 - \frac{1}{2-\beta} > 0$. □

6 Open problems

We left open whether we can compute the diameter of all the graphs of constant distance VC-dimension and unbounded diameter in truly subquadratic time. In order to prove that it is the case, we stress that by our Theorem 3 we only need to consider the graphs of large diameter, i.e., above some polynomial.

Another interesting direction could be to extend the constant distance VC-dimension framework to weighted graphs.

Finally, we observe that there exist graph families of *unbounded* (distance) VC-dimension for which we can compute the diameter very efficiently. For instance, the class of all the graphs with a universal vertex has unbounded VC-dimension. Such graphs are a particular case of *dually chordal* graphs, for which we know how to compute the diameter in linear time [12]. We observe that the

ball hypergraphs of dually chordal graphs also admit some nice characterizations. Thus, it would be very interesting to study whether a truly subquadratic algorithm for computing the diameter could be derived from some common property of dually chordal graphs and graphs of constant distance VC-dimension (say, a bounded fractional Helly number [51]).

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.
- [2] N. Alon, P. Seymour, and R. Thomas. A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society*, 3(4):801–808, 1990.
- [3] M. Anthony, G. Brightwell, and C. Cooper. The Vapnik-Chervonenkis dimension of a random graph. *Discrete Mathematics*, 138(1-3):43–56, 1995.
- [4] A. Backurs, L. Roditty, G. Segal, V. Vassilevska Williams, and N. Wein. Towards tight approximation bounds for graph diameter and eccentricities. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*, pages 267–280. ACM, 2018.
- [5] C. Berge. *Graphs and hypergraphs*. North-Holland Pub. Co., 1973.
- [6] V. Bilo, V. Goyal, R. Ravi, and M. Singh. On the crossing spanning tree problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 51–60. Springer, 2004.
- [7] J. A. Bondy and U. S. R. Murty. *Graph theory*. Springer, 2008.
- [8] 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.
- [9] N. Bousquet, A. Lagoutte, Z. Li, A. Parreau, and S. Thomassé. Identifying codes in hereditary classes of graphs and VC-dimension. *SIAM Journal on Discrete Mathematics*, 29(4):2047–2064, 2015.
- [10] N. Bousquet and S. Thomassé. VC-dimension and Erdős–Pósa property. *Discrete Mathematics*, 338(12):2302–2317, 2015.
- [11] A. Brandstädt, V. Chepoi, and F. Dragan. The algorithmic use of hypertree structure and maximum neighbourhood orderings. *Discrete Applied Mathematics*, 82(1-3):43–77, 1998.
- [12] A. Brandstädt, F. Dragan, V. Chepoi, and V. Voloshin. Dually chordal graphs. *SIAM Journal on Discrete Mathematics*, 11(3):437–455, 1998.
- [13] K. Bringmann, T. Husfeldt, and M. Magnusson. Multivariate analysis of orthogonal range searching and graph distances parameterized by treewidth. In *IPEC*, 2018.
- [14] H. Brönnimann, B. Chazelle, and J. Matoušek. Product range spaces, sensitive sampling, and derandomization. *SIAM J. Comput.*, 28(5):1552–1575, 1999.

- [15] S. Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Transactions on Algorithms (TALG)*, 15(2):21, 2018.
- [16] M. Cairo, R. Grossi, and R. Rizzi. New bounds for approximating extremal distances in undirected graphs. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms (SODA 2016)*, pages 363–376. SIAM, 2016.
- [17] T. Chan. Optimal partition trees. *Discrete & Computational Geometry*, 47(4):661–690, 2012.
- [18] B. Chazelle and E. Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete & Computational Geometry*, 4(5):467–489, 1989.
- [19] S. Chechik, D. Larkin, L. Roditty, G. Schoenebeck, R. Tarjan, and V. Vassilevska Williams. Better approximation algorithms for the graph diameter. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms (SODA 2014)*, pages 1041–1052. SIAM, 2014.
- [20] V. Chepoi, B. Estellon, and Y. Vaxès. Covering planar graphs with a fixed number of balls. *Discrete & Computational Geometry*, 37(2):237–244, 2007.
- [21] M. Cohen, Y. Lee, and Z. Song. Solving linear programs in the current matrix multiplication time. In *STOC*, pages 938–942. ACM, 2019.
- [22] D. Corneil, F. Dragan, M. Habib, and C. Paul. Diameter determination on restricted graph families. *Discrete Applied Mathematics*, 113(2-3):143–166, 2001.
- [23] D. Coudert, G. Ducoffe, and A. Popa. Fully polynomial FPT algorithms for some classes of bounded clique-width graphs. In *SODA '18*, pages 2765–2784. SIAM, 2018.
- [24] S. Dahlgaard. On the hardness of partially dynamic graph problems and connections to diameter. In *43rd International Colloquium on Automata, Languages, and Programming, (ICALP 2016)*, pages 48:1–48:14, 2016.
- [25] P. Damaschke. Computing giant graph diameters. In *IWOCA*, pages 373–384. Springer, 2016.
- [26] R. Downey, P. Evans, and M. Fellows. Parameterized learning complexity. In *COLT*, pages 51–57. ACM, 1993.
- [27] G. Ducoffe. A New Application of Orthogonal Range Searching for Computing Giant Graph Diameters. In *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*, 2019.
- [28] G. Ducoffe, M. Habib, and L. Viennot. Fast diameter computation within split graphs. To appear in COCOA'19.
- [29] Z. Dvořák. On classes of graphs with strongly sublinear separators. *European Journal of Combinatorics*, 71:1–11, 2018.
- [30] Z. Dvorak and S. Norin. Strongly sublinear separators and polynomial expansion. *SIAM Journal on Discrete Mathematics*, 30(2):1095–1101, 2016.
- [31] Z. Dvořák and S. Norin. Treewidth of graphs with balanced separations. *Journal of Combinatorial Theory, Series B*, 137:137–144, 2019.

- [32] D. Eisenstat and D. Angluin. The VC dimension of k -fold union. *Information Processing Letters*, 101(5):181–184, 2007.
- [33] D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3-4):275–291, 2000.
- [34] D. Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 599–608. ACM/SIAM, 2003.
- [35] J. Evald and S. Dahlgaard. Tight Hardness Results for Distance and Centrality Problems in Constant Degree Graphs. Technical Report arXiv:1609.08403, ArXiv, 2016.
- [36] A. Farley and A. Proskurowski. Computation of the center and diameter of outerplanar graphs. *Discrete Applied Mathematics*, 2(3):185–191, 1980.
- [37] G. Federickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing*, 16(6):1004–1022, 1987.
- [38] P. Gawrychowski, H. Kaplan, S. Mozes, M. Sharir, and O. Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\tilde{O}(n^{5/3})$ time. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018)*, pages 495–514. SIAM, 2018.
- [39] M. Habib, R. McConnell, C. Paul, and L. Viennot. Lex-BFS and partition refinement, with applications to transitive orientation, interval graph recognition and consecutive ones testing. *Theoretical Computer Science*, 234(1-2):59–84, 2000.
- [40] S. Har-Peled. Approximating Spanning Trees with Low Crossing Number. Technical Report abs/0907.1131, arXiv, 2009.
- [41] S. Har-Peled and K. Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. *SIAM Journal on Computing*, 46(6):1712–1744, 2017.
- [42] D. Haussler and E. Welzl. ε -nets and simplex range queries. *Discrete & Computational Geometry*, 2(2):127–151, 1987.
- [43] M. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of computer and system sciences*, 55(1):3–23, 1997.
- [44] K. Kawarabayashi, P. N. Klein, and C. Sommer. Linear-space approximate distance oracles for planar, bounded-genus and minor-free graphs. In L. Aceto, M. Henzinger, and J. Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, volume 6755 of *Lecture Notes in Computer Science*, pages 135–146. Springer, 2011.
- [45] K. Kawarabayashi and B. Reed. A separator theorem in minor-closed classes. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 153–162. IEEE, 2010.
- [46] J. Kleinberg. Detecting a network failure. *Internet Mathematics*, 1(1):37–55, 2004.

- [47] E. Kranakis, D. Krizanc, B. Ruf, J. Urrutia, and G. Woeginger. The VC-dimension of set systems defined by graphs. *Discrete Applied Mathematics*, 77(3):237–257, 1997.
- [48] J. Li and M. Parter. Planar diameter via metric compression. In *STOC*, pages 152–163. ACM, 2019.
- [49] R. Lipton and R. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- [50] J. Matoušek. Spanning trees with low crossing number. *RAIRO-Theoretical Informatics and Applications-Informatique Théorique et Applications*, 25(2):103–123, 1991.
- [51] J. Matousek. Bounded VC-dimension implies a fractional Helly theorem. *Discrete & Computational Geometry*, 31(2):251–255, 2004.
- [52] J. Nešetřil and P. Ossona de Mendez. *Sparsity: graphs, structures, and algorithms*, volume 28. Springer Science & Business Media, 2012.
- [53] J. Nešetřil and P. Ossona de Mendez. Structural sparsity. *Russian Mathematical Surveys*, 71(1):79, 2016.
- [54] S. Olariu. A simple linear-time algorithm for computing the center of an interval graph. *International Journal of Computer Mathematics*, 34(3-4):121–128, 1990.
- [55] C. Papadimitriou and M. Yannakakis. On limited nondeterminism and the complexity of the VC dimension. *Journal of Computer and System Sciences*, 53(2):161–170, 1996.
- [56] S. Plotkin, S. Rao, and W. Smith. Shallow Excluded Minors and Improved Graph Decompositions. In *SODA*, volume 90, pages 462–470, 1994.
- [57] L. Roditty and V. Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *STOC*, pages 515–524. ACM, 2013.
- [58] N. Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972.
- [59] S. Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247–261, 1972.
- [60] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pages 11–30. Springer, 2015.
- [61] E. Welzl. On spanning trees with low crossing numbers. In *Data structures and efficient algorithms*, pages 233–249. Springer, 1992.
- [62] C. Wulff-Nilsen. Separator theorems for minor-free and shallow minor-free graphs with applications. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS 2011)*, pages 37–46. IEEE, 2011.