



The diameter of AT-free graphs

Guillaume Ducoffe

► To cite this version:

Guillaume Ducoffe. The diameter of AT-free graphs. Journal of Graph Theory, 2022, 99 (4), pp.594-614. [⟨10.1002/jgt.22754⟩](#). [⟨hal-03370937⟩](#)

HAL Id: hal-03370937

<https://hal.science/hal-03370937v1>

Submitted on 8 Oct 2021

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.



HAL Authorization

The diameter of AT-free graphs

Guillaume Ducoffe^{1,2}

¹National Institute for Research and Development in Informatics, Romania

²University of Bucharest, Romania

Abstract

A graph algorithm is *truly subquadratic* if it runs in $\mathcal{O}(m^b)$ time on connected m -edge graphs, for some positive $b < 2$. Roditty and Vassilevska Williams (STOC'13) proved that under plausible complexity assumptions, there is no truly subquadratic algorithm for computing the diameter of general graphs. In this work, we present positive and negative results on the existence of such algorithms for computing the diameter on some special graph classes. Specifically, three vertices in a graph form an *asteroidal triple* (AT) if between any two of them there exists a path that avoids the closed neighbourhood of the third one. We call a graph *AT-free* if it does not contain an AT. We first prove that for all m -edge AT-free graphs, one can compute all the eccentricities in truly subquadratic $\mathcal{O}(m^{3/2})$ time. For the AT-free bipartite graphs, it can be improved to linear time. Then, we extend our study to several subclasses of chordal graphs – all of them generalizing interval graphs in various ways –, as an attempt to understand which of the properties of AT-free graphs, or natural generalizations of the latter, can help in the design of fast algorithms for the diameter problem on broader graph classes. For instance, for all chordal graphs with a dominating shortest path, there is a linear-time algorithm for computing a diametral pair if the diameter is at least four. However, already for split graphs with a dominating edge, under plausible complexity assumptions, there is no truly subquadratic algorithm for deciding whether the diameter is either 2 or 3.

Keywords: AT-free graphs, Chordal graphs, Dominating pairs, Fine-grained complexity in P.

1 Introduction

For any undefined graph terminology, see [3]. All graphs studied in this paper are finite, simple and connected. Given a graph $G = (V, E)$, let $n = |V|$ be its order and $m = |E|$ be its size. Note that, since we assume G to be connected, $m \geq n - 1$. For a vertex $u \in V$, let $N_G(u) = \{v \in V \mid uv \in E\}$ and $N_G[u] = \{u\} \cup N_G(u)$ be, respectively, the open and closed neighbourhoods of u . The distance between two vertices $u, v \in V$ is equal to the minimum number of edges on a uv -path, and it is denoted by $\text{dist}_G(u, v)$. The maximum such distance between a fixed vertex u and all other vertices is called its eccentricity, denoted by $e_G(u) = \max_{v \in V} \text{dist}_G(u, v)$. We sometimes omit the subscript if the graph G is clear from the context. Finally, the *diameter* of a graph G is equal to $\text{diam}(G) = \max_{u, v \in V} \text{dist}(u, v) = \max_{u \in V} e(u)$. The purpose of this note is to study the complexity of diameter computation within some graph classes. On general graphs, the best known algorithms for this problem run in $\mathcal{O}(nm)$ time and in $\mathcal{O}(n^{\omega+o(1)})$ time, respectively, where $\omega < 2.3729$ is the exponent of square matrix multiplication [33]. In particular, both algorithms run in $\Omega(m^2)$ time on sparse graphs, where by sparse we mean that $m \leq c \cdot n$ for some universal constant c .

Improving the quadratic running time for diameter computation – in the number m of edges – is an important research problem, both in theory and in practice. Since any algorithm for this problem must run in $\Omega(m)$ time, some authors have considered whether we can compute the diameter in linear time or in quasi linear time on certain graph classes. Notably, this is the case for the *interval graphs* – i.e., the intersection graphs of intervals on the real line – [30]. All graph classes that are considered in this work are superclasses of the interval graphs. For general graphs, less than a decade ago [31], Roditty and Vassilevska Williams gave convincing evidence that the complexity of diameter computation *cannot* be improved by much. Specifically, the Strong Exponential-Time Hypothesis (SETH) says that for any $\varepsilon > 0$, there exists a k such that k -SAT on n variables cannot be solved in $\mathcal{O}((2 - \varepsilon)^n)$ time [23]. The Orthogonal-Vector problem (OV) takes as input two families A and B of n sets over some universe C , and it asks whether there exist $a \in A, b \in B$ s.t. $a \cap b = \emptyset$. Some older works have named this problem DISJOINT SETS [8]. The following result is due to Williams:

Theorem 1 ([36]). *Under SETH, for any $\varepsilon > 0$, there exists a constant $c > 0$ such that we cannot solve OV in $\mathcal{O}(n^{2-\varepsilon})$ time, even if $|C| \leq c \cdot \log n$.*

In what follows, by *truly subquadratic* we mean a running time in $\mathcal{O}(N^{2-\varepsilon})$, for some $\varepsilon > 0$, where N denotes the size of the input (for connected graphs, $N \approx m$). Theorem 1 was used in order to prove that many classic problems that can be solved in polynomial time do not admit a truly subquadratic algorithm (e.g., see [34] for a survey). In particular, as far as we are concerned in this note, under SETH there is no truly subquadratic algorithm for computing the diameter, even on sparse graphs [1]. This negative result has motivated a long line of papers, with some trying to characterize the graph classes for which there *does* exist a truly subquadratic algorithm for the diameter problem. We refer to [1, 5, 6, 11] for recent relevant work in this area. In particular, several authors have studied whether some important properties of interval graphs could imply on their own fast diameter computation algorithms. This is the case for Helly graphs [16], and graphs of bounded distance VC-dimension [18], but *not* for chordal graphs (a.k.a., graphs with no induced cycle of length larger than three) [4].

We recall that a graph is called AT-free if there does not exist a triple x, y, z such that, for any two of them, there exists a path that avoids the closed neighbourhood of the third one. The interval graphs are exactly the chordal AT-free graphs [26]. The complexity of (exact and approximate) diameter computation within AT-free graphs was studied in [9] (see also [14, 22]), where the authors emphasize a kind of duality between AT-free graphs and chordal graphs. For instance, on both graph classes, two consecutive executions of LexBFS always yield a vertex whose eccentricity is within one of the diameter – this is the so-called 2-sweep LexBFS algorithm, see Fig. 2. However, there is no constant c such that c consecutive executions of LexBFS on these graph classes always output the exact diameter. The authors from [9] further sketch a reduction from OV to diameter computation within AT-free graphs and chordal graphs, as evidence that the diameter problem on these graph classes cannot be solved in linear time. This same reduction was revisited in [4] in order to prove that, indeed, under SETH there is no truly subquadratic algorithm for computing the diameter on chordal graphs. But the same *cannot* be done for AT-free graphs, because for the latter, the reduction in [9] from OV to diameter computation already runs in $\Omega(n^2)$ time.

- **Our first main result** is an $\mathcal{O}(m^{3/2})$ -time algorithm in order to compute all the eccentricities (and so, the diameter) in an AT-free graph (Theorem 2).

The proof of this above result stays simple. Nevertheless, it comes to us as surprise given the evidence for SETH-hardness in [9]. On dense graphs – with $m \geq c \cdot n^2$ edges, for some constant c – our algorithm does no better than the brute-force $\mathcal{O}(nm)$ -time algorithm for All-Pairs Shortest-Paths.

- We show that it is unlikely the running time can be improved, in the following sense: computing the diameter, resp. all eccentricities, in an AT-free graph, is at least as hard as computing a simplicial vertex in an arbitrary graph, resp. as finding a triangle (Theorem 3). The best known algorithms for both problems run in $\mathcal{O}(nm)$ time and $\mathcal{O}(n^\omega)$ time.

Our reductions are inspired from the one in [9] and to reductions from these problems to many more graph problems in [25]. In particular, due to the large web of equivalences in [35] between triangle and matrix problems, we see our reduction from the triangle detection problem as strong evidence of impossibility for (combinatorial) subcubic computation of the eccentricities in dense AT-free graphs. Our reduction from the problem of computing a simplicial vertex to the diameter problem on dense AT-free graphs is also interesting because of similar reductions to a more diverse set of graph problems in [25]. By comparison, the complexity status of OV for dense instances *i.e.*, for which we have $C = \Theta(n)$, remains unclear.

For claw-free AT-free graphs, there is a linear-time algorithm in order to compute the diameter and all the vertices of minimum eccentricity [22]. As a special case of Theorem 2, we prove that the same holds for AT-free bipartite graphs: for these graphs, one can compute all the eccentricities in linear time (Corollary 1). Note that the AT-free bipartite graphs have attracted some attention on their own in the literature [20].

The *ball hypergraph* of G has for hyperedges the balls of all possible centers and radii in G . On our way, we observe that for the family of ball hypergraphs of AT-free graphs, classic geometric parameters such as the Helly number and the VC-dimension are unbounded (Proposition 1). It sets AT-free graphs apart from most known graph classes with a truly subquadratic algorithm for the diameter problem [18, 16]. We initiate the investigation of the complexity of diameter computation within graph classes sharing a common property with AT-free graphs. In this paper, we only study such classes which are subclasses of chordal graphs. This is for the following two main reasons. On one hand, under SETH, there is no truly subquadratic algorithm for computing the diameter, already for chordal graphs. On the other hand, chordal graphs are more structured than general graphs, thereby making easier the design and the analysis of our algorithms. We summarise our results for subclasses of chordal graphs (we postpone their technical definitions to appropriate places throughout the paper):

- For every chordal graph with *asteroidal number* at most k , there is a randomized $\mathcal{O}(km \log^2 n)$ -time algorithm in order to compute the diameter with high probability (Theorem 4). The same holds for chordal graphs with *leafage* at most k . In contrast to this positive result, it is easy to prove that under SETH, there is no truly subquadratic algorithm for diameter computation within k -AT-free chordal graphs, for every $k \geq 2$ (Proposition 3);
- For every chordal *dominating pair* graph, there is a truly subquadratic algorithm for computing the diameter (Theorem 7). For the larger class of chordal graphs with a dominating shortest path, there is a linear-time algorithm for computing a diametral pair if the diameter is at least four; however, already for split graphs with a dominating edge, under SETH there is no truly subquadratic algorithm for computing the diameter (Theorem 6);

- For every chordal graph with a dominating triple, if the diameter is at least 10, then there is a linear-time algorithm in order to compute a diametral pair (Theorem 9).

We stress that all the aforementioned graph classes generalize interval graphs in various ways, that are incomparable one with another. For chordal Helly graphs and chordal graphs of bounded VC-dimension: two other generalizations of interval graphs that are incomparable with each other and with the other subclasses presented above, there also exist truly subquadratic algorithms for the diameter problem [16]. We left open whether Theorem 4 can be derandomized. It is also open whether the lower bound of 10 on the diameter for Theorem 9 is tight. Finally, we point out the dichotomy of Theorem 6, where we show that the *only* difficulty for a chordal graph with a dominating shortest-path is to decide whether the diameter is either two or three. By contrast, for any $d \geq 1$, under SETH there is no truly subquadratic algorithm for distinguishing the chordal graphs of diameter $\leq 2d$ from those of diameter $\geq 2d + 1$ (to see this, just start from any split graph, and replace every vertex of its stable set by a path of length $d - 1$).

2 All eccentricities for AT-free graphs

2.1 Preliminaries

We start recalling a few results from prior work, that we will use in this paper.

Graph searches. The Lexicographic Breadth-First Search (LexBFS) is a standard algorithmic procedure, that runs in linear time [32]. We give a pseudo-code in Fig. 1. Note that we can always enforce a start vertex u by assigning to it an initial non empty label. Then, for a given graph $G = (V, E)$ and a start vertex u , $LexBFS(u)$ denotes the corresponding execution of LexBFS. Its output is a numbering σ over the vertex-set (namely, the reverse of the ordering in which vertices are visited during the search). In particular, if $\sigma(i) = x$, then $\sigma^{-1}(x) = i$.

```

Input: A graph  $G = (V, E)$ 
Output: An ordering  $\sigma$  of the vertices of  $V$ 
begin
  assign the label  $\emptyset$  to each vertex ;
  for  $i = n$  to 1 do
    pick an unnumbered vertex  $x$  with the largest label in the lexicographic order ;
    for each unnumbered neighbour  $y$  of  $x$  do
      add  $i$  to  $label(y)$  ;
     $\sigma(i) \leftarrow x$  /* number  $x$  by  $i$  */ ;
  end

```

Figure 1: Algorithm LexBFS [32].

The 2-sweep LexBFS consists of two consecutive execution of LexBFS, with the start vertex of the second execution being the last one visited during the first LexBFS. See also Fig. 2. More generally, for any positive integer c , the algorithm c -sweep LexBFS consists of c consecutive applications of LexBFS, with for any $i \geq 2$, the start vertex of the i^{th} execution being the last vertex visited during the $(i - 1)^{th}$. We prove most of our results using 3-sweep LexBFS, but often use the known properties of 2-sweep LexBFS in our proofs.

```

Input: A graph  $G$ 
Output: A vertex  $v$ 
begin
  Let  $w$  be an arbitrary vertex;
   $u \leftarrow$  the last vertex numbered by LexBFS( $w$ );
   $v \leftarrow$  the last vertex numbered by LexBFS( $u$ );
  return  $v$ ;
end

```

Figure 2: Algorithm 2-sweep [9].

AT-free graphs. We now recall several useful properties of LexBFS orderings within AT-free graphs. A *dominating pair* in a graph is a pair (u, v) of two vertices such that, for every vertex x , every uv -path intersects its closed neighbourhood $N[x]$.

Lemma 1 ([10]). *For an AT-free graph $G = (V, E)$, let u be the last vertex visited during a LexBFS, and let $\sigma = \text{LexBFS}(u)$. Then, for every vertex y , the pair (u, y) is dominating for the subgraph induced by $\{z \in V \mid \sigma^{-1}(z) \geq \sigma^{-1}(y)\}$. In particular, if $\sigma(1) = v$, then (u, v) is a dominating pair for G .*

Lemma 2 ([15]). *For an AT-free graph $G = (V, E)$, let u be the last vertex visited during a LexBFS, and let $\sigma = \text{LexBFS}(u)$. If $x, y \in V$ are such that: $xy \notin E$, $\text{dist}(u, x) = \text{dist}(u, y) = i$, and $\sigma^{-1}(x) < \sigma^{-1}(y)$, then we have $N(x) \cap \{z \in V \mid \text{dist}(u, z) = i - 1\} \subseteq N(y) \cap \{z \in V \mid \text{dist}(u, z) = i - 1\}$. In particular, $\text{dist}(x, y) = 2$.*

2.2 The algorithm

In what follows, we prove that all eccentricities in an AT-free graph can be computed in truly subquadratic time (Theorem 2). Before that, we need to prove a few intermediate lemmas.

Lemma 3. *Let (u, v) be a dominating pair in a graph $G = (V, E)$ (not necessarily AT-free). Every vertex $x \in V$ s.t. $e(x) \geq 3$ is at distance $e(x)$ from a vertex in $N[u] \cup N[v]$.*

Proof. Let $x \in V$ be s.t. $e(x) \geq 3$, and let $y \in V$ s.t. $\text{dist}(x, y) = e(x)$. In what follows, we assume $y \notin N[u] \cup N[v]$ (otherwise, we are done). Name P a shortest uv -path in G s.t. $\text{dist}(y, P)$ is minimized. In particular, we have $y \in V(P)$ if and only if $\text{dist}(u, v) = \text{dist}(u, y) + \text{dist}(y, v)$. We pick any two vertices $x^* \in N[x] \cap V(P)$ and $y^* \in N[y] \cap V(P)$, that always exist because, by the hypothesis, P is dominating. By symmetry, we may assume that y^* is on the subpath of P between u and x^* . We claim that $e(x) = \text{dist}(u, x)$, that will prove the lemma. Suppose for the sake of contradiction that it is not the case. Consider the (not necessarily simple) uv -path P' that is obtained from the concatenation of $P[v, x^*]$: the subpath of P between v and x^* , with an arbitrary shortest xu -path Q . Since (u, v) is a dominating pair, $N[y] \cap V(P') \neq \emptyset$. However, we prove in what follows that y has no neighbour on $P[v, x^*]$. Suppose by contradiction that there exists a vertex $z \in N[y]$ on the subpath $P[v, x^*]$. Then, $\text{dist}(y^*, z) \leq 2$. Observe that we have $\text{dist}(y^*, z) = \text{dist}(y^*, x^*) + \text{dist}(x^*, z)$. Furthermore, $x^* \notin N[y]$ because we assume $\text{dist}(x, y) = e(x) \geq 3$. Since P is a shortest uv -path, it follows that $y^*, z \in N(y) \cap N(x^*)$ and $\text{dist}(y^*, z) = 2$. But then, we could replace x^* by y on our shortest uv -path, thus contradicting the minimality of $\text{dist}(y, P)$. As a result, and since $N[y] \cap V(P') \neq \emptyset$, we must have $N[y] \cap V(Q) \neq \emptyset$. Since Q is a

shortest ux -path and we also have $u \notin N[y]$, we derive from the above that $\text{dist}(x, y) \leq \text{dist}(x, u)$, thus contradicting our assumption that $e(x) > \text{dist}(x, u)$. \square

By using Lemma 3, we will prove that in order to compute all eccentricities, it suffices to perform a BFS from all vertices in the first and two last distance layers of some shortest-path tree, whose root is the vertex output by the algorithm 3-sweep LexBFS. However, there does not seem to be a simple way in order to upper bound the number of vertices in these layers. We complete our result, as follows:

Lemma 4. *For an AT-free graph $G = (V, E)$, let u be the last vertex visited during a LexBFS, and let $\sigma = \text{LexBFS}(u)$. Let $x, y \in V$ be s.t.: $xy \notin E$, $\text{dist}(u, x) = \text{dist}(u, y) = i$, and $\sigma^{-1}(x) < \sigma^{-1}(y)$. For any vertex $z \in V$, if $\text{dist}(u, z) < i$ then we always have $\text{dist}(y, z) \leq \text{dist}(x, z)$.*

In particular, if G has no universal vertex and $\text{dist}(u, x) = \text{dist}(u, y) = e(u)$, then $e(x) \geq e(y)$.

Proof. Let $z \in V$ be such that $\text{dist}(u, z) < i$. In particular, $\sigma^{-1}(z) > \sigma^{-1}(y) > \sigma^{-1}(x)$. For any shortest xz -path P , we claim that $N[y] \cap V(P) \neq \emptyset$. Note that, since we assume $xy \notin E$, this will prove that $\text{dist}(y, z) \leq \text{dist}(x, z)$. In order to prove the claim, let $x' \in V(P)$ minimize $\sigma^{-1}(x')$. By construction, $\sigma^{-1}(x') \leq \sigma^{-1}(x) < \sigma^{-1}(y)$. Let $P_{x'}$ be the subpath of P between x' and z . We complete $P_{x'}$ into an ux' -path P' (not necessarily simple) by adding to it a shortest uz -path Q . In particular, all the vertices $w \in V(P') = V(P_{x'}) \cup V(Q)$ satisfy $\sigma^{-1}(w) \geq \sigma^{-1}(x')$. Since we also have $\sigma^{-1}(y) > \sigma^{-1}(x')$, it follows by Lemma 1, $N[y] \cap V(P') \neq \emptyset$. Furthermore, since any vertex $w \in V(Q) \setminus \{z\}$ satisfies $\text{dist}(u, w) \leq \text{dist}(u, z) - 1 \leq \text{dist}(u, y) - 2$, we get $N[y] \cap (V(Q) \setminus \{z\}) = \emptyset$. The latter proves, as claimed, $N[y] \cap V(P) \supseteq N[y] \cap V(P_{x'}) \neq \emptyset$.

Finally, let us assume that $\text{dist}(u, x) = \text{dist}(u, y) = e(u)$. If $z \in V$ is such that $\text{dist}(u, z) < e(u)$, then we proved above that we have $\text{dist}(y, z) \leq \text{dist}(x, z)$. Otherwise, $\text{dist}(u, z) = e(u)$, and by Lemma 2 we get that $\text{dist}(y, z) \leq 2$. Altogether combined, $e(y) \leq \max\{e(x), 2\}$. If moreover, G has no universal vertex, then $e(x) \geq 2$, and so, we conclude that $e(y) \leq e(x)$. \square

If we execute three consecutive LexBFS instead of two, then a weaker converse of Lemma 4 for *lower* distance layers (instead of upper ones) also holds. Namely:

Lemma 5. *For an AT-free graph $G = (V, E)$, let u be the last vertex visited during a LexBFS, let $\sigma = \text{LexBFS}(u)$, and let $v = \sigma(1)$ be the last vertex visited during LexBFS(u). Let $x, y \in V$ be s.t.: $xy \notin E$, $\text{dist}(v, x) = \text{dist}(v, y) = i$, and $\sigma^{-1}(x) < \sigma^{-1}(y)$. For any vertex $z \in V$, if $\text{dist}(v, z) > i$ then we always have $\text{dist}(y, z) \leq \max\{\text{dist}(x, z), 2\}$.*

Proof. Since we have $\sigma^{-1}(x) < \sigma^{-1}(y)$, $\text{dist}(u, x) \geq \text{dist}(u, y) = j$. Now, let $z \in V$ be s.t. $\text{dist}(v, z) > i$. If furthermore, $\text{dist}(u, z) < j$, then we claim that $\text{dist}(y, z) \leq \text{dist}(x, z)$. Indeed, if $\text{dist}(u, x) = \text{dist}(u, y)$, then the claim follows from Lemma 4. From now on, we assume $\text{dist}(u, x) > \text{dist}(u, y)$. Any shortest xz -path must contain a vertex x_z s.t. $\text{dist}(u, x_z) = \text{dist}(u, y)$. We may further assume $x_z \in N(x)$ (otherwise, since by Lemma 2 we have $\text{dist}(y, x_z) \leq 2$, $\text{dist}(y, z) \leq 2 + \text{dist}(x_z, z) \leq \text{dist}(x, z)$, and we are done). In particular, $\text{dist}(u, x) = j + 1$. In the same way, we may assume that y, x_z are nonadjacent, and that any neighbour $x' \in N(x_z)$ on a shortest $x_z z$ -path satisfies $\text{dist}(u, x') = j - 1$. Then, consider an xu -path P starting with $[x, x_z, x']$ and continuing with any shortest $x' u$ -path. By Lemma 1, P dominates all vertices w s.t. $\sigma^{-1}(w) \geq \sigma^{-1}(x)$. In particular, $N[y] \cap V(P) \neq \emptyset$. Since we assume $x, x_z \notin N[y]$, it implies $x' \in N(y)$. In such case, $\text{dist}(y, z) \leq 1 + \text{dist}(x', z) < \text{dist}(x, z)$, thereby proving our claim.

We are left with the case $\text{dist}(u, z) \geq j$. In order to prove the lemma, it suffices to prove that $\text{dist}(y, z) \leq 2$. If $\text{dist}(u, z) = j$, then this follows from Lemma 2. Thus, from now on we assume $\text{dist}(u, z) > j$. In order to conclude here, we need the following observation: for every vertex w , we always have $\text{dist}(u, v) \leq \text{dist}(u, w) + \text{dist}(w, v) \leq \text{dist}(u, v) + 2$. Indeed, the first inequality follows from the triangular inequality. The second inequality follows from Lemma 1 saying that any shortest uv -path is dominating (to see this, consider any $w^* \in N[w]$ on such a shortest-path, and observe that $\text{dist}(u, w) + \text{dist}(w, v) \leq \text{dist}(u, w^*) + \text{dist}(w^*, v) + 2 = \text{dist}(u, v) + 2$). In our case, $\text{dist}(u, z) + \text{dist}(z, v) \geq j + 1 + i + 1 = \text{dist}(u, y) + \text{dist}(y, v) + 2$. Therefore, y is on a shortest uv -path Q , and in addition $\text{dist}(u, z) = j + 1$, $\text{dist}(v, z) = i + 1$. By Lemma 1, there exists a $z^* \in N[z] \cap V(Q)$. The only possibility w.r.t. $\text{dist}(u, z)$, $\text{dist}(v, z)$ is to have $z^* = y$, and so, $\text{dist}(y, z) = 1$. \square

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

Theorem 2. *For every m -edge AT-free graph $G = (V, E)$, we can compute the eccentricities in $\mathcal{O}(m^{\frac{3}{2}})$ time.*

Proof. Let u be the last vertex visited during a LexBFS, and let $\sigma = \text{LexBFS}(u)$. Similarly, let $v = \sigma(1)$ be the last vertex visited during $\text{LexBFS}(u)$, and let $\tau = \text{LexBFS}(v)$. Set $d = e(v)$. For any $0 \leq i \leq d$, we define $L_i = \{w \in V \mid \text{dist}(v, w) = i\}$. We compute the set A from $L_1 \cup \{v\}$ and σ as follows. We scan all the vertices $v' \in L_1 \cup \{v\}$ by increasing value of $\sigma^{-1}(v')$, removing from this set the non-neighbours of v' . Similarly, we compute the set B (resp., C) from L_{d-1} (resp., L_d) and τ as follows. We scan all the vertices $u' \in L_{d-1}$ (resp., $u' \in L_d$) by increasing value of $\tau^{-1}(u')$, removing from this set the non-neighbours of u' . The three of A, B, C can be computed in linear time. Furthermore, by construction every vertex added in A (resp., in B or in C) is adjacent to all the vertices previously added into this set (otherwise, this vertex should have been removed). Therefore, the sets A, B, C are cliques, and so, their cardinality is in $\mathcal{O}(\sqrt{m})$.

For every vertex $w \in V$, we claim that we have:

$$e(w) = \begin{cases} 1 & \text{if } w \text{ is a universal vertex} \\ \max\{2\} \cup \{\text{dist}(w, x) \mid x \in A \cup B \cup C\} & \text{otherwise.} \end{cases}$$

Indeed, w is universal if and only if $e(w) = 1$. The above formula is also trivially true if $e(w) = 2$. Therefore, let us assume in what follows $e(w) \geq 3$. By Lemma 1 $(\tau(1), v)$ is a dominating pair, and therefore by Lemma 3, w is at distance $e(w)$ from some vertex in $N[\tau(1)] \cup N[v]$. Note that $N[\tau(1)] \cup N[v] \subseteq \{v\} \cup L_1 \cup L_{d-1} \cup L_d$. Let $y \in \{v\} \cup L_1 \cup L_{d-1} \cup L_d$ s.t. $e(w) = \text{dist}(y, w)$. We consider in what follows four different cases:

- Case $y = v$. By construction, $v \in A$. Therefore, the above formula for $e(w)$ holds in this case.
- Case $y \in L_1$. We assume w.l.o.g. $\sigma^{-1}(y)$ is minimized for this property. Suppose for the sake of contradiction $y \notin A$ (otherwise, we are done). By the construction of A , there exists a $x \in L_1$ s.t. $\sigma^{-1}(x) < \sigma^{-1}(y)$ and x, y are nonadjacent. Furthermore, $\text{dist}(v, w) > 1$ because otherwise, we would get $e(w) = \text{dist}(y, w) \leq 2$. By Lemma 5, $\text{dist}(y, w) \leq \max\{2, \text{dist}(x, w)\}$. Since in addition we have $\text{dist}(y, w) = e(w) \geq 3$, we conclude in this case that we have $\text{dist}(x, w) = e(w)$, that contradicts the minimality of $\sigma^{-1}(y)$.

- Case $y \in L_d$. The proof is quite similar as for the previous case, but slightly simpler. We assume w.l.o.g. $\tau^{-1}(y)$ is minimized for this property. Suppose for the sake of contradiction $y \notin C$. By the construction of C , there exists a $x \in L_d$ s.t. $\tau^{-1}(x) < \tau^{-1}(y)$ and x, y are nonadjacent. Furthermore, $\text{dist}(v, w) < d$ because otherwise, we would get by Lemma 2 $e(w) = \text{dist}(y, w) \leq 2$. By Lemma 5, $\text{dist}(y, w) \leq \text{dist}(x, w)$. Since in addition we have $\text{dist}(y, w) = e(w)$, we conclude in this case that we have $\text{dist}(x, w) = e(w)$, that contradicts the minimality of $\tau^{-1}(y)$.
- Case $y \in L_{d-1}$. Here also, the proof is essentially the same as for the previous case. In fact, we only need to detail the special case $w \in L_d$. We claim that we have $e(w) = 3$. Indeed, let $x \in N(w) \cap L_{d-1}$. By Lemma 2 we have $\text{dist}(w, y) \leq 1 + \text{dist}(x, y) \leq 3$. We may further assume w.l.o.g. $d - 1 > 1$ (otherwise, $y \in \{v\} \cup L_1$). But then, $\text{dist}(v, w) = d \geq 3 = e(w)$.

Finally, we claim that the above formula can be computed, for all vertices, in total $\mathcal{O}(m^{3/2})$ time. Indeed, it suffices to perform a BFS from every vertex of $A \cup B \cup C$, and we observed above that there are only $\mathcal{O}(\sqrt{m})$ many such vertices. \square

We observe that the actual running time of our algorithm is in $\mathcal{O}(km)$ on m -edge AT-free graphs with *clique-number* at most k . In particular, we get:

Corollary 1. *For every AT-free bipartite graph, we can compute the eccentricities in linear time.*

2.3 Digression: properties of the neighbourhood hypergraph

For a graph $G = (V, E)$, its neighbourhood hypergraph is $\mathcal{N}(G) = (V, \{N[v] \mid v \in V\})$. Note that it is a subhypergraph of the ball hypergraph (as defined in the introduction). The *Helly number* of $\mathcal{N}(G)$ is the smallest k s.t. every family of k -wise intersecting neighbourhoods of G (i.e., hyperedges) have a nonempty common intersection. Its *VC-dimension* is the largest d s.t., for some vertex-subset X of cardinality d , for every subset $Y \subseteq X$, there exists a $v \in V$ s.t. $N[v] \cap X = Y$ (we say that X is shattered by $\mathcal{N}(G)$). It was proved recently that a bounded Helly number or bounded VC-dimension for the ball hypergraph implies fast diameter and radius computation algorithms [18, 16]. Actually, for most graph classes for which we know how to compute the diameter in truly subquadratic time, one of these two parameters above is always bounded. For instance, on interval graphs, both parameters are at most two. We prove that such a property does not hold for the AT-free graphs:

Proposition 1. *There are AT-graphs whose neighbourhood hypergraph has a Helly number (resp., VC-dimension) that is arbitrarily large.*

Proof. We take the opportunity to recall the reduction in [9] from OV to AT-free graphs. Let A, B be two families of sets over a universe C . The graph $H_{A,B,C}$ has vertex-set $A \cup B \cup C$. Its edge-set is as follows. The sets A, B, C are cliques. For every $a \in A$ and $c \in C$, a and c are adjacent if and only if $c \in a$. In the same way, for every $b \in B$ and $c \in C$, b and c are adjacent if and only if $c \in b$. In [9], the authors observed that $H_{A,B,C}$ is a cocomparability graph (and so, it is AT-free). We prove that for some suitable A, B and C , the neighbourhood hypergraph of $H_{A,B,C}$ has arbitrarily large Helly number (resp., VC-dimension).

First, for any fixed d , we consider a family A over some universe C and VC-dimension $\geq d$. Let X be of cardinality d and shattered by A . For every $Y \subseteq X$, there exists a set $a \in A$ s.t. $a \cap X = Y$.

In particular, identifying $X \subseteq C$ with a vertex-subset of $H_{A,\emptyset,C}$, we also get $N[a] \cap X = Y$. As a result, the VC-dimension of $\mathcal{N}(H_{A,\emptyset,C})$ is at least d (since X is shattered by this hypergraph).

In the same way, for any fixed k , let F be a minimal family of Helly number $\geq k+1$ over some universe C . By minimality of the family, all sets in F k -wise intersect, but they have an empty common intersection. We arbitrarily bipartition F into nonempty subfamilies A and B . Then, we claim that the neighbourhood hypergraph of $H_{A,B,C}$ has Helly number at least $k+1$. Indeed, by construction the neighbour sets $N[a], a \in A$ and $N[b], b \in B$ k -wise intersect. However, since their common intersection must be in C , then by the choice of F the latter is empty. \square

2.4 Hardness results

Our algorithm for computing the eccentricities in an AT-free graph, presented in Theorem 2, is combinatorial, and it runs in $\mathcal{O}(n^3)$ time on n -vertex (dense) graphs. We show, via the following reduction, it is unlikely that such a running time can be improved without the use of algebraic techniques. If we do allow such algebraic manipulations, then we can use the well-known Seidel's algorithm [33] in order to compute all the distances in a general n -vertex graph in $\mathcal{O}(n^\omega \log n) = \tilde{\mathcal{O}}(n^{2.3729})$ time. Combined with Theorem 2, it leads to an $\tilde{\mathcal{O}}(n^{\omega/3}m) = \tilde{\mathcal{O}}(n^{0.8}m)$ -time algorithm for computing the eccentricities of n -vertex m -edge AT-free graphs.

Theorem 3. *The problem of deciding whether an n -vertex graph contains a simplicial vertex can be reduced in $\mathcal{O}(n^2)$ time to the diameter problem on $\mathcal{O}(n)$ -vertex AT-free graphs.*

Furthermore, the problem of finding a triangle in an n -vertex graph can be reduced in $\mathcal{O}(n^2)$ time to the problem of computing all the diametral vertices (and so, all eccentricities) in an $\mathcal{O}(n)$ -vertex AT-free graph.

Proof. Consider an n -vertex graph $G = (V, E)$. We construct the following graph H :

- $V(H) = V_1 \cup V_2 \cup V'_2 \cup V_3 \cup V'_3 \cup V_4$ where each set V_i or V'_i is a copy of V . For every $v \in V$ we denote by v_i (resp., v'_i) its copy in V_i (resp., in V'_i);
- The sets $V_1, V_2 \cup V'_2, V_3 \cup V'_3, V_4$ are cliques;
- We add edges $v_1 v'_2, v'_2 v'_3$ for every $v \in V$;
- For every $v \in V$, we add the edges $v_1 u_2$ for every $u \in N_G(v)$. Similarly, we add the edges $v_3 u_4$ for every $u \in N_G(v)$. However, we add the edges $v_2 w_3$ for every $w \in V \setminus N_G[v]$;
- Finally, for every $v \in V$, we add the edges $v'_3 w_4$ for every $w \in V \setminus v$.

This graph H has $6n$ vertices, and it can be constructed in $\mathcal{O}(n^2)$ time. We claim that H is AT-free. Indeed, suppose by contradiction that there exists an AT x, y, z in H . Since the three of x, y, z are pairwise non adjacent, each clique $V_1, V_2 \cup V'_2, V_3 \cup V'_3, V_4$ contains at most one such vertex. Then, let us totally order these four cliques as $V_1 \prec V_2 \cup V'_2 \prec V_3 \cup V'_3 \prec V_4$, and let us define the following partial ordering over $V(H)$: $u \prec v$ if and only if, for $A, B \in \{V_1, V_2 \cup V'_2, V_3 \cup V'_3, V_4\}$ such that $u \in A, v \in B$, we have $A \prec B$. W.l.o.g., $x \prec y \prec z$. But then, there is a clique containing y that disconnects x from z . The latter contradicts that x, z are in the same connected component of $H \setminus N_H[y]$. Therefore, we proved as claimed that H is AT-free.

By construction, all vertices in H are pairwise at distance at most 3, except maybe some pairs of vertices $v_1 \in V_1, w_4 \in V_4$. Furthermore, if $v \neq w$ then, due to the existence of the path

$[v_1, v'_2, v'_3, w_4]$, we have $\text{dist}_H(v_1, w_4) = 3$. As a result, for every $v \in V$, $e_H(v_1) = \text{dist}_H(v_1, v_4) \in \{3, 4\}$, and $\text{diam}(H) = \max_{v \in V} e_H(v_1)$. It now remains to characterize the vertices $v \in V$ such that $e_H(v_1) = 4$. We claim these vertices to be exactly the simplicial vertices of G . To see this, assume $\text{dist}_H(v_1, v_4) = 3$, and let $[v_1, x, y, v_4]$ be a corresponding shortest path. By construction, $x \in V_2 \cup V'_2$, $y \in V_3 \cup V'_3$. Since v'_2 is the only neighbour of v_1 in V'_2 , v'_3 is the only neighbour of v'_2 in $V_3 \cup V'_3$ and v'_3 and v_4 are non adjacent, $x \notin V'_2$. In particular, since no vertex of V_2 has a neighbour in V'_3 , $y \notin V'_3$. Overall, $x = u_2$ and $y = w_3$ for some $u, w \in V$. By construction of H , $u, w \in N_G(v)$. Furthermore, since $u_2 w_3 \in E(H)$, we get $w \in V \setminus N_G[u]$. Conversely, if $u \neq w$ are non adjacent neighbours of v then, due to the existence of the path $[v_1, u_2, w_3, v_4]$, we get $\text{dist}_H(v_1, v_4) = 3$. Summarizing $\text{dist}_H(v_1, v_4) = 3$ if and only if vertex v has two non adjacent neighbours in G . Equivalently, $e_H(v_1) = \text{dist}_H(v_1, v_4) = 4$ if and only if v is simplicial in G .

Our claim implies that $\text{diam}(H) = 4$ if and only if G contains a simplicial vertex. Furthermore, if $\text{diam}(H) = 4$, then the diametral vertices are exactly those v_1, v_4 such that v is simplicial in G . It is known [25] that the problem of finding a triangle in an n -vertex graph can be reduced in $\mathcal{O}(n^2)$ time to the problem of counting the number of simplicial vertices in an $\mathcal{O}(n)$ -vertex graph. \square

3 Chordal graphs with bounded asteroidal number

Let us recall that an asteroidal set in a graph $G = (V, E)$ is an independent set $A \subseteq V$ s.t., for every $v \in A$, all vertices in $A \setminus \{v\}$ are in the same connected component of $G \setminus N[v]$. In particular, an AT is exactly an asteroidal set of cardinality three. The *asteroidal number* of a graph G is the largest cardinality of its asteroidal sets. Since with this terminology, the interval graphs are exactly the chordal graphs of asteroidal number two, the following result generalizes the celebrated linear-time algorithm [30] for computing the diameter in this graph class:

Theorem 4. *There is a randomized $\mathcal{O}(km \log^2 n)$ -time algorithm for computing w.h.p. the diameter of chordal graphs with asteroidal number at most k .*

Our results are based on a general framework for computing the diameter of chordal graphs [16]. We recall that a split graph G is a graph whose vertex-set can be bipartitioned into a clique K and a stable set S . We may assume G to be given under its *sparse representation*, defined in [17] as being the hypergraph $(K \cup S, \{N_G[s] \mid s \in S\})$. The SPLIT-OV problem is a special case of OV where $A = \{N_G[a] \mid a \in S_A\}$ and $B = \{N_G[b] \mid b \in S_B\}$, for some split graph G and for some partition S_A, S_B of its stable set. An instance of SPLIT-OV can be encoded as a triple (G, S_A, S_B) , with G being given under its sparse representation. Note that the size of such an instance is dominated by $\ell = \sum_{s \in S_A \cup S_B} |N_G[s]|$, that is $\leq m + n$.

Theorem 5 (Theorem 8' from [16]). *For a subclass \mathcal{C} of chordal graphs, let \mathcal{S} be the subclass of all split graphs that are induced subgraphs of a chordal graph in \mathcal{C} . If for every (G, S_A, S_B) , with $G \in \mathcal{S}$ connected, we can solve SPLIT-OV in $\mathcal{O}(\ell^b)$ time, for some $b \geq 1$, then there is a randomized $\mathcal{O}(m^b \log^2 n)$ -time algorithm for computing w.h.p. the diameter of chordal graphs in \mathcal{C} .*

Recall that having asteroidal number at most k is a hereditary property. Hence, by Theorem 5, in order to prove Theorem 4, it suffices to solve SPLIT-OV in $\mathcal{O}(k\ell)$ time for connected split graphs of asteroidal number at most k . Before presenting such an algorithm (Proposition 2), we need a few preparatory lemmas.

Lemma 6. *For a split graph G with stable set S , and $A \subseteq S$ of cardinality $|A| \geq 3$, A is an asteroidal set if and only if the neighbour sets $N(a)$, $a \in A$ are pairwise incomparable w.r.t. inclusion.*

Proof. In one direction, if $N(a') \subseteq N(a)$ for some distinct $a, a' \in A$, then we claim that A is not an asteroidal set. Indeed, a' is isolated in $G \setminus N[a]$, and so disconnected from $A \setminus \{a, a'\}$, that is nonempty because $|A| \geq 3$. In the other direction, let us assume the neighbour sets $N(a)$, $a \in A$ to be pairwise incomparable w.r.t. inclusion. Let $a \in A$ be arbitrary. By the hypothesis all the subsets $N(a') \setminus N(a)$, $a' \in A \setminus \{a\}$ are nonempty. Since G is a split graph, $N(A \setminus \{a\}) \setminus N(a) = \bigcup_{a' \in A \setminus \{a\}} N(a') \setminus N(a)$ is a clique. Therefore, the vertices of $A \setminus \{a\}$ are in the same connected component of $G \setminus N[a]$. \square

Lemma 7. *Let \mathcal{F} be a family of pairwise different subsets, and $\ell = \sum_{S \in \mathcal{F}} |S|$. If there are at most k inclusionwise minimal subsets in \mathcal{F} , then all these subsets can be computed in total $\mathcal{O}(k\ell)$ time.*

Proof. It suffices to prove that in $\mathcal{O}(\ell)$ time, we can compute an inclusionwise minimal subset and remove all its supersets from \mathcal{F} . Specifically, let $S \in \mathcal{F}$ be of minimum cardinality. Since all sets in \mathcal{F} are pairwise different, S is inclusion minimal. Furthermore, if we first mark all elements in S , and we store the cardinality $|S|$ of this set, then we can remove all its supersets from \mathcal{F} by scanning the family in total $\mathcal{O}(\ell)$ time. \square

Combining Lemmas 6 and 7, we get the following algorithm for SPLIT-OV on split graphs of asteroidal number at most k :

Proposition 2. *For every (G, S_A, S_B) where G has asteroidal number at most $k \geq 2$, we can solve SPLIT-OV in $\mathcal{O}(k\ell)$ time.*

Proof. Two vertices $a, a' \in S_A$ are *twins* if $N(a) = N(a')$. It is now folklore that by using partition refinement techniques, we can compute all twin classes of S_A in $\mathcal{O}(\ell)$ time [21]. Note also that in order to solve SPLIT-OV, it is sufficient to keep only one vertex per twin class of S_A . Thus, from now on we assume that the neighbour sets $N(a)$, $a \in S_A$ are pairwise different. Observe that if $a, a' \in S_A$ and $b \in S_B$ satisfy $N(a') \subset N(a)$ and $N(a) \cap N(b) = \emptyset$, then $N(a') \cap N(b) = \emptyset$. In particular, we may further restrict S_A to the vertices a s.t. $N(a)$ is inclusion wise minimal. By Lemma 6, there are at most k such vertices. Therefore, by Lemma 7, we can compute all these vertices in total $\mathcal{O}(k\ell)$ time. Finally, for each $a \in S_A$ s.t. $N(a)$ is inclusion wise minimal, we can compute a vertex $b \in S_B$ s.t. $N(a) \cap N(b) = \emptyset$ (if any) simply by scanning $N(a)$ and all the neighbour sets $N(b)$, $b \in S_B$. It takes total $\mathcal{O}(\ell)$ time. \square

Theorem 4 now follows from Theorem 5 and Proposition 2. \square

Chordal graphs are exactly the intersection graphs of a collection of subtrees of a host tree [19]. We call such a representation a tree model. The leafage of a chordal graph is the smallest number of leaves amongst its tree models. Since the leafage is an upper bound on the asteroidal number [27], we obtain:

Corollary 2. *There is a randomized $\mathcal{O}(km \log^2 n)$ -time algorithm for computing w.h.p. the diameter of chordal graphs with leafage at most k .*

Finally, a k -AT in a graph is a triple x, y, z of vertices such that, for any two of them, there exists a path that avoids the ball of radius k of the third one, where the latter is defined, for any vertex

v , as $N^k[v] = \{u \mid \text{dist}(u, v) \leq k\}$. A graph is k -AT-free if it does not contain any k -AT. Note that in particular, the AT-free graphs are exactly the 1-AT-free graphs. This generalization of AT-free graphs was first proposed in [28], where the performances of BFS for the latter were studied. Before concluding this section, we make the following simple observation:

Proposition 3. *Under SETH, there is no truly subquadratic algorithm for computing the diameter of k -AT-free chordal graphs, for every $k \geq 2$.*

Proof. Under SETH, there is no truly subquadratic algorithm for computing the diameter of split graphs [4]. Therefore, it suffices to prove that split graphs are k -AT-free, for every $k \geq 2$. Indeed, for a split graph with clique K , the ball of radius 2 for *any* vertex contains K . In particular, its removal either leaves an empty graph, a singleton, or an independent set. As a result, there can be no k -AT in a split graph, for every $k \geq 2$. \square

4 Chordal graphs with a dominating shortest-path

Recall (see Lemma 1) that every AT-free graph has a dominating pair. A *dominating pair* graph is one s.t. every induced subgraph has a dominating pair. We study diameter computation within chordal dominating pair graphs in Sec. 4.2. A weaker property for a graph is to have a dominating shortest-path. We first study chordal graphs with a dominating shortest-path, in Sec. 4.1, for which we derive an interesting dichotomy result.

4.1 Dichotomy theorem

The purpose of this section is to prove the following result:

Theorem 6. *For every chordal graph with a dominating shortest-path, if the algorithm 3-sweep LexBFS outputs a vertex of eccentricity $d \geq 3$, then it is the diameter. In particular, there is a linear-time algorithm for deciding whether the diameter is at least four on this graph class.*

However, already for split graphs with a dominating edge, under SETH there is no truly subquadratic algorithm for deciding whether the diameter is either two or three.

We start with the following easy lemma:

Lemma 8 ([12]). *If a graph has a dominating shortest-path, then it has a dominating diametral path.*

Roughly, the first part of Theorem 6 follows from Lemma 8 combined with the following properties of LexBFS on chordal graphs:

Lemma 9 ([9]). *Let u be the vertex of a chordal graph G last visited by a LexBFS, and let x, y be a pair of vertices such that $\text{dist}(x, y) = \text{diam}(G)$. If $e(u) < \text{diam}(G)$ then $e(u)$ is even, $\text{dist}(u, x) = \text{dist}(u, y) = e(u)$ and $e(u) = \text{diam}(G) - 1$.*

Corollary 3 ([9]). *If the vertex u of a chordal graph G last visited by a LexBFS has odd eccentricity, then $e(u) = \text{diam}(G)$.*

Proof of Theorem 6. Let u be last visited by a LexBFS, and let v be last visited by $\text{LexBFS}(u)$. Assume toward a contradiction $e(v) < \text{diam}(G)$. Since $e(u) = \text{dist}(u, v) \leq e(v)$, and by Lemma 9, $e(u) \geq \text{diam}(G) - 1$, we have $e(u) = e(v) = \text{dist}(u, v) = \text{diam}(G) - 1$. Let x, y be the two ends of a dominating diametral path P , that exists by Lemma 8. We pick $u^* \in V(P) \cap N[u]$ and $v^* \in V(P) \cap N[v]$, that exist because P is dominating. W.l.o.g., $\text{dist}(x, u^*) \leq \text{dist}(x, v^*)$. Then:

$$\begin{aligned} \text{diam}(G) &= \text{dist}(x, y) = \text{dist}(x, u^*) + \text{dist}(u^*, v^*) + \text{dist}(v^*, y) \\ &\geq (\text{dist}(x, u) - 1) + (\text{dist}(u, v) - 2) + (\text{dist}(v, y) - 1) \\ &= (\text{diam}(G) - 2) + (\text{diam}(G) - 3) + (\text{diam}(G) - 2) = 3 \cdot \text{diam}(G) - 7 \end{aligned}$$

where the equalities of the last line follow from Lemma 9. It implies $\text{diam}(G) \leq 3$. Therefore, if $e(v) \geq 3$, we cannot have $e(v) < \text{diam}(G)$.

For the second part of the theorem, we essentially rely on a previous observation from [17]. The bichromatic diameter problem consists of, given a graph and two vertex-subsets A and B , to computing the maximum distance between any vertex of A and any vertex of B . Under SETH, we cannot compute the bichromatic diameter of split graphs [4]. Now, given a split graph G with clique K along with two subsets A, B in its stable set, we can add in linear time fresh new vertices $a, b \notin V(G)$ s.t. $N[a] = A \cup K \cup \{a, b\}$ and $N[b] = B \cup K \cup \{a, b\}$. In doing so, we get a new split graph G' , for which solving the diameter is equivalent to computing the bichromatic diameter of G . Note that by construction, ab is a dominating edge of G' . As a result, already for split graphs with a dominating edge, we cannot decide whether the diameter is either two or three. \square

4.2 Chordal dominating pair graphs

We complete the results of Sec 4.1 by showing that the stronger property of having a dominating pair *for each induced subgraph* implies a truly subquadratic algorithm for diameter computation.

Theorem 7. *There is a truly subquadratic algorithm for computing the diameter of chordal dominating pair graphs.*

Note that in contrast to Theorem 7, the dichotomy result of Theorem 6 also applies to the chordal graphs with a dominating pair (*a.k.a.*, chordal weakly dominating pair graphs). Indeed, having a dominating shortest-path and a dominating edge are a weaker and a stronger property than having a dominating pair, respectively.

The remainder of this subsection is devoted to the proof of Theorem 7. Note that according to Theorem 6, the only difficulty is in order to decide whether the diameter is either two or three. If we further use Theorem 5 (at the price of having a randomized algorithm), then we are left solving SPLIT-OV for the split dominating pair graphs. The following result is an easy corollary of the characterization proven in [13]:

Lemma 10 ([13]). *A split graph is a dominating pair graph if and only if it is B_1 -free, where B_1 is the graph of Fig 3.*

Our main technical contribution in this subsection is as follows:

Lemma 11. *Every B_1 -free split graph has VC-dimension at most 3.*




Figure 3: Forbidden induced subgraph B_1 .

Proof. Let $G = (K \cup S, E)$ be a split graph with clique K and stable set S . Suppose for the sake of contradiction G is B_1 -free and there exists a $X \subseteq K \cup S$, $|X| \geq 4$, that is shattered.

We first prove that either $X \subseteq K$ or $X \subseteq S$. Indeed, by contradiction let $u \in X \cap K$, $v \in X \cap S$. Since X is shattered, there exists a z s.t. $N[z] \cap X = \{u, v\}$. Furthermore, since G is a split graph and $v \in S$, either $z = v$ or $z \in K$. But if $z = v$, then, for any $z' \in N[z]$ we have $u \in N[z']$, and so there can be no such z' s.t. $N[z'] \cap X = \{v\}$. Thus, necessarily, $z \in K$. It implies $X \setminus \{u, v\} \subseteq S$ (otherwise, since $K \cap X \subseteq N[z]$, we could not have $N[z] \cap X = \{u, v\}$, a contradiction). However, let $w \in X \setminus \{u, v\}$. Again, since X is shattered, there exists a z' s.t. $N[z'] \cap X = \{v, w\}$. But necessarily, $z' \in K$, and therefore we also have $u \in N[z']$, a contradiction. The latter proves, as claimed, either $X \subseteq K$ or $X \subseteq S$. Note that the above still applies if $|X| \geq 3$, and that we did not use in our proof the fact that G is B_1 -free.

If $X \subseteq K$ then, let $u, v, w \in X$. There exist z_u, z_v, z_w s.t. $N[z_u] \cap X = \{u\}$, $N[z_v] \cap X = \{v\}$, $N[z_w] \cap X = \{w\}$. Necessarily, $z_u, z_v, z_w \in S$. But then, u, v, w, z_u, z_v, z_w induce a copy of B_1 , that is a contradiction. Conversely, if $X \subseteq S$ then, let $u, v, w, r \in X$. There exist z_u, z_v, z_w s.t. $N[z_u] \cap X = \{u, r\}$, $N[z_v] \cap X = \{v, r\}$, $N[z_w] \cap X = \{w, r\}$. Necessarily, $z_u, z_v, z_w \in K$. But then, u, v, w, z_u, z_v, z_w induce a copy of B_1 , that is a contradiction. \square

Lemma 11 gives us the opportunity to use powerful techniques from previous work in order to solve SPLIT-OV.

Theorem 8 (special case of Theorem 1 in [18]). *For every $d > 0$, there exists a constant $\varepsilon_d \in (0, 1)$ s.t. for any (G, S_A, S_B) where G has VC-dimension at most d , we can solve SPLIT-OV in deterministic $\tilde{O}(\ell n^{1-\varepsilon_d})$ time.*

Some additional remarks are needed. First of all, Theorem 1 in [18] addresses a monochromatic variant of DISJOINT SETS, that is slightly different than OV. A simple trick (presented in the proof of Theorem 10 in [16]) allows us to reduce SPLIT-OV to this monochromatic variant, up to increasing the VC-dimension from d to some value in $\mathcal{O}(d \log d)$. Second, the time complexity for Theorem 8 is in $\tilde{O}(\ell n^{1-\varepsilon_d})$, and not in $\mathcal{O}(\ell^b)$ for some $b \geq 1$ as it was stated in Theorem 5. However, since $\ell n^{1-\varepsilon_d} = \mathcal{O}(\ell^{2-\varepsilon_d})$, Theorem 5 can still be applied in this case. \square

5 Dominating triples

A *dominating target* for a graph G is a subset of vertices D s.t. every connected graph containing all of D is a dominating set. In particular, dominating pairs are exactly dominating targets of cardinality two. An interesting generalization of Lemma 1 is that every graph of asteroidal number at most k contains a dominating target of cardinality at most k [24]. In this last section, we study chordal graphs with a dominating triple (dominating set of cardinality three).

Theorem 9. *For every chordal graph with a dominating triple, if the algorithm 3-sweep LexBFS outputs a vertex of eccentricity $d \geq 10$, then it is the diameter. In particular, there is a linear-time algorithm for deciding whether the diameter is at least 10 on this graph class.*

We use in our proof a few results from Metric Graph Theory, that we now introduce. A *geodesic triangle* with corners $x, y, z \in V$, denoted in what follows by $\Delta(x, y, z)$, is the union $P(x, y) \cup P(y, z) \cup P(z, x)$ of three shortest-paths connecting its corners. The three of $P(x, y), P(y, z), P(z, x)$ are called the sides of the triangle. We say that $\Delta(x, y, z)$ is δ -slim if the maximum distance between any vertex on one side $P(x, y)$ and the other two sides $P(y, z) \cup P(z, x)$ is at most δ . A graph is called δ -slim if all its geodesic triangles are.

Lemma 12 ([29]). *Every chordal graph is 1-slim.*

Three vertices x, y, z form a *metric triangle* if, for any choice of shortest-paths $P(x, y), P(y, z), P(z, x)$ connecting them, the latter can only pairwise intersect at their endpoints. If furthermore, $\text{dist}(x, y) = \text{dist}(y, z) = \text{dist}(z, x) = k$, then we say of this metric triangle that it is equilateral of size k . The following result was proved in [2] for meshed graphs, that are a superclass of chordal graphs.

Lemma 13 ([2]). *Every metric triangle in a chordal graph G is equilateral.*

We next refine Lemma 13 by also bounding the size of metric triangles for chordal graphs. The following result is known by some researchers in Metric Graph Theory, but we were unable to find a reference.

Corollary 4. *Every metric triangle in a chordal graph G is equilateral of size at most two.*




Figure 4: Situation in the proof of the first case of Corollary 4.

Proof. Suppose for the sake of contradiction that there exists a metric triangle x, y, z of size at least three. By Lemma 13, this triangle is equilateral of size $k \geq 3$. Fix three shortest-paths $P(x, y), P(y, z), P(z, x)$ connecting these vertices, thus obtaining a geodesic triangle $\Delta(x, y, z)$. Let $u \in P(x, y)$ s.t. $\text{dist}(x, u) = 2$. By Lemma 12, there exists a neighbour $v \in N(u) \cap (P(y, z) \cup P(z, x))$. We need to consider two cases.

- *Case $v \in P(z, x)$.* Then, $\text{dist}(x, v) \geq 2$ because otherwise, there would be a shortest xy -path going through the edge uv , thus contradicting that x, y, z is a metric triangle. In the same way, $\text{dist}(z, v) \geq k - 2$, and therefore, we have: $\text{dist}(x, v) = 2$, $\text{dist}(z, v) = k - 2$ (see Fig. 4). Let $[x, s, u]$ and $[x, t, v]$ be shortest subpaths of $P(x, y)$ and $P(z, x)$, respectively. Since $[x, s, u, v, t, x]$ is a cycle of length five, there exists a chord. However, the only possible such chord is st (i.e., because $ut \in E$ and $vs \in E$ imply t is on a shortest xy -path and s is on a shortest xz -path, respectively). In this situation, u, v, s, t induces a cycle of length four, a contradiction.

- *Case $v \in P(y, z)$.* We prove next that $k \leq 4$. Indeed, this is the case if $v \in \{y, z\}$. Otherwise, since x, y, z is a metric triangle, we cannot have v on a shortest xy -path nor on a shortest xz -path. Then, $\text{dist}(y, v) \geq k - 2$, $\text{dist}(z, v) \geq k - 2$, and therefore, $k = \text{dist}(y, z) = \text{dist}(y, v) + \text{dist}(v, z) \geq 2k - 4$. It implies $k \leq 4$. If $k = 4$ then, since we also have $\text{dist}(y, u) = 2$, we are back to the previous case up to replacing x by y . From now on, let us assume $k = 3$. Since the geodesic triangle considered is a cycle of length nine, there exists a chord. However, consider any such chord st . Vertices s and t are adjacent to some vertices amongst $\{x, y, z\}$, and so, the edge st is on a path of length three between two corners of the triangle. Since x, y, z is a metric triangle, the only possible chords are, for every of x, y, z , between their two neighbours on different sides of the triangle. However, in this situation we are left with an induced cycle of length at least six, thus contradicting that G is chordal.

□

Note that the bound of Corollary 4 is sharp, as it is shown, *e.g.*, by the 3-sun.

Finally, a *quasi-median* for $x, y, z \in V$ is a triple $x^*, y^*, z^* \in V$ s.t.:

$$\begin{cases} \text{dist}(x, y) = \text{dist}(x, x^*) + \text{dist}(x^*, y^*) + \text{dist}(y^*, y) \\ \text{dist}(y, z) = \text{dist}(y, y^*) + \text{dist}(y^*, z^*) + \text{dist}(z^*, z) \\ \text{dist}(z, x) = \text{dist}(z, z^*) + \text{dist}(z^*, x^*) + \text{dist}(x^*, x). \end{cases}$$

Chalopin et al. [7] observed that every triple of vertices has a quasi median which is a metric triangle.

Proof of Theorem 9. Let x, y, z be a dominating triple, and let x^*, y^*, z^* be a corresponding pseudo median that is also a metric triangle. We fix shortest-paths $P(x, x^*), P(y, y^*), P(z, z^*)$ along with a geodesic triangle $\Delta(x^*, y^*, z^*)$ with sides $P(x^*, y^*), P(y^*, z^*), P(z^*, x^*)$. In doing so, since x^*, y^*, z^* is a pseudo median, we also get shortest-paths $P(x, y), P(y, z), P(z, x)$. For instance, $P(x, y) = P(x, x^*) \cup P(x^*, y^*) \cup P(y^*, y)$. Furthermore, since x, y, z is a dominating triple, the union H of these above shortest-paths is a dominating set of G . Let u be last visited by a LexBFS, and let v be last visited by $\text{LexBFS}(u)$. Finally, let (s, t) be a diametral pair of G . Assume toward a contradiction that $\text{diam}(G) \geq 11$ and $e(v) < \text{diam}(G)$ (we will explain at the end of the proof how to lower the bound on the diameter to 10).

We first prove as an intermediate claim that for some two vertices $\alpha, \beta \in \{x, y, z\}$, we have that $\text{dist}(u, P(\alpha, \beta)) + \text{dist}(v, P(\alpha, \beta)) \leq 3$. Indeed, let $u^* \in N[u] \cap H$ and $v^* \in N[v] \cap H$. If $u^*, v^* \in P(\alpha, \beta)$ for some $\alpha, \beta \in \{x, y, z\}$, then we get $\text{dist}(u, P(\alpha, \beta)) + \text{dist}(v, P(\alpha, \beta)) \leq 2$. From now on, we assume that it is not the case. If $u^*, v^* \in \Delta(x^*, y^*, z^*)$, then since by Corollary 4 this metric triangle is equilateral of size ≤ 2 , $\text{dist}(u^*, v^*) \leq 3$. However, it implies by Lemma 9 that $\text{diam}(G) \leq \text{dist}(u, v) + 1 \leq \text{dist}(u^*, v^*) + 3 \leq 6$, a contradiction. Hence, let us assume for instance that $u^* \in P(x, x^*)$ but $v^* \in P(y^*, z^*)$ is on the only side of $\Delta(x^*, y^*, z^*)$ that is not on a shortest-path between x and one of $\{y, z\}$ (all other cases are symmetrical to this one). Since x^*, y^*, z^* is equilateral of size at most two, for any $\beta \in \{y, z\}$ we get $\text{dist}(v^*, P(x^*, \beta^*)) = 1$. As a result, $\text{dist}(u, P(x, \beta)) + \text{dist}(v, P(x, \beta)) \leq 1 + 2 = 3$, thus proving the claim. Furthermore, as a by-product of our proof, we also have in this case $\max\{\text{dist}(u, P(\alpha, \beta)), \text{dist}(v, P(\alpha, \beta))\} \leq 2$.

We prove as another intermediate claim that for some $w \in \{s, t\}$ we have $\text{dist}(w, P(\alpha, \beta)) \leq 2$. Indeed, let $s^* \in N[s] \cap H$ and $t^* \in N[t] \cap H$. If one amongst s^* and t^* is a vertex of

$\Delta(x^*, y^*, z^*)$, say it is the case of s^* , then since the latter triangle is equilateral of size at most two, $\text{dist}(s, P(\alpha, \beta)) \leq 1 + \text{dist}(s^*, P(\alpha^*, \beta^*)) \leq 2$. From now on, we assume that s^* and t^* are not vertices of $\Delta(x^*, y^*, z^*)$. If we write $\{\alpha, \beta, \gamma\} = \{x, y, z\}$, then we may assume w.l.o.g. $s^*, t^* \in P(\gamma, \gamma^*)$ (otherwise, one amongst s^*, t^* is a vertex of $P(\alpha, \beta)$, and so we are done). We prove as a subclaim that $\text{dist}(u^*, P(\alpha, \gamma) \cup P(\beta, \gamma)) \leq 1$. Indeed, to see this, it suffices to recall that either $u^* \in P(\alpha, \alpha^*) \cup P(\beta, \beta^*)$ or $u^* \in P(\alpha^*, \beta^*)$. In the former case, $\text{dist}(u^*, P(\alpha, \gamma) \cup P(\beta, \gamma)) = 0$, while in the latter case, since we have a metric triangle that is equilateral of size two, $\text{dist}(u^*, P(\alpha^*, \gamma^*)) = \text{dist}(u^*, P(\beta^*, \gamma^*)) = 1$. Thus, choosing $u' \in P(\alpha, \gamma) \cup P(\beta, \gamma)$ at minimum distance from u^* :

$$\begin{aligned} \text{dist}(s^*, u^*) &\leq \text{dist}(u^*, u') + (\text{dist}(\gamma, u') - \text{dist}(\gamma, s^*)) \\ &\leq 1 + \text{diam}(G) - \text{dist}(\gamma, s^*). \end{aligned}$$

In particular, $\text{dist}(s, u) \leq \text{dist}(s^*, u^*) + 2 \leq \text{diam}(G) + 3 - \text{dist}(\gamma, s^*)$. Since we have $\text{dist}(s, u) = \text{diam}(G) - 1$ by Lemma 9, it follows from the above inequalities that we have $\text{dist}(\gamma, s^*) \leq 4$. We get in the exact same way $\text{dist}(\gamma, t^*) \leq 4$. However, it implies $\text{diam}(G) = \text{dist}(s, t) \leq 2 + \text{dist}(s^*, t^*) \leq 2 + \text{dist}(s^*, \gamma) + \text{dist}(\gamma, t^*) \leq 10$, a contradiction.

Overall, let $u', v', w' \in P(\alpha, \beta)$ at minimum distance from u, v, w respectively. By Lemma 9, u, v, w are pairwise at distance $\text{diam}(G) - 1$. Let $\{r_1, r_2, r_3\} = \{u, v, w\}$ s.t. r'_2 is (metrically) between r'_1 and r'_3 onto the shortest-path $P(\alpha, \beta)$. Then:

$$\begin{aligned} \text{diam}(G) &\geq \text{dist}(\alpha, \beta) \\ &\geq \text{dist}(r'_1, r'_2) + \text{dist}(r'_2, r'_3) \\ &\geq \text{dist}(r_1, r_2) + \text{dist}(r_2, r_3) - \text{dist}(r_1, r'_1) - 2 \cdot \text{dist}(r_2, r'_2) - \text{dist}(r_3, r'_3) \\ &\geq 2(\text{diam}(G) - 1) - \text{dist}(u, u') - \text{dist}(v, v') - \text{dist}(w, w') \\ &\quad - \max\{\text{dist}(u, u'), \text{dist}(v, v'), \text{dist}(w, w')\} \\ &\geq 2 \cdot \text{diam}(G) - 2 - 3 - 2 - 2 = 2 \cdot \text{diam}(G) - 9. \end{aligned}$$

But then, $\text{diam}(G) \leq 9$, a contradiction.

Finally, let us sketch how we can decide in linear time whether the diameter is at least equal to d , for any $d \geq 10$. For $d = 10$, we claim that it suffices to execute a 2-sweep LexBFS. Indeed, by Lemma 9, the output vertex has eccentricity either 9 or 10, and if it is 9 then, by Corollary 3, we have $\text{diam}(G) = 9 < 10$. Otherwise, $d \geq 11$ and we apply the algorithm 3-sweep LexBFS, whose correctness follows from the above analysis. \square

We left open the following intriguing question. For any $k \geq 2$, does there exist a $d_k > 0$ s.t., for any (chordal) graph with a dominating target of cardinality at most k , we can decide in truly subquadratic time whether the diameter is at least d_k , and if so compute the diameter exactly?

Acknowledgements

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

References

- [1] A. Abboud, V. Vassilevska Williams, and J. Wang, *Approximation and fixed parameter sub-quadratic algorithms for radius and diameter in sparse graphs*, Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete Algorithms (SODA), SIAM, 2016, pp. 377–391.
- [2] H.-J. Bandelt and V. Chepoi, *Graphs with connected medians*, SIAM Journal on Discrete Mathematics **15** (2002), no. 2, 268–282.
- [3] J. A. Bondy and U. S. R. Murty, *Graph theory*, 2008.
- [4] 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** (2016), 51–67.
- [5] M. Borassi, P. Crescenzi, and L. Trevisan, *An axiomatic and an average-case analysis of algorithms and heuristics for metric properties of graphs*, Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2017, pp. 920–939.
- [6] S. Cabello, *Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs*, ACM Transactions on Algorithms (TALG) **15** (2018), no. 2, 1–38.
- [7] J. Chalopin, V. Chepoi, P. Papasoglu, and T. Pecatte, *Cop and robber game and hyperbolicity*, SIAM Journal on Discrete Mathematics **28** (2014), no. 4, 1987–2007.
- [8] V. Chepoi and F. Dragan, *Disjoint set problem*, 1992.
- [9] D. Corneil, F. Dragan, M. Habib, and C. Paul, *Diameter determination on restricted graph families*, Discrete Applied Mathematics **113** (2001), no. 2-3, 143–166.
- [10] D. Corneil, S. Olariu, and L. Stewart, *Linear time algorithms for dominating pairs in asteroidal triple-free graphs*, SIAM Journal on Computing **28** (1999), no. 4, 1284–1297.
- [11] D. Coudert, G. Ducoffe, and A. Popa, *Fully polynomial FPT algorithms for some classes of bounded clique-width graphs*, ACM Transactions on Algorithms (TALG) **15** (2019), no. 3, 1–57.
- [12] J. Deogun and D. Kratsch, *Diametral path graphs*, International Workshop on Graph-Theoretic Concepts in Computer Science, Springer, 1995, pp. 344–357.
- [13] ———, *Dominating pair graphs*, SIAM Journal on Discrete Mathematics **15** (2002), no. 3, 353–366.
- [14] F. Dragan, *Almost diameter of a house-hole-free graph in linear time via LexBFS*, Discrete applied mathematics **95** (1999), no. 1-3, 223–239.
- [15] F. Dragan, E. Köhler, and A. Leitert, *Line-distortion, bandwidth and path-length of a graph*, Algorithmica **77** (2017), no. 3, 686–713.
- [16] G. Ducoffe and F. Dragan, *A story of diameter, radius and Helly property*, Networks **77** (2021), no. 3, 435–453.

- [17] G. Ducoffe, M. Habib, and L. Viennot, *Fast diameter computation within split graphs*, International Conference on Combinatorial Optimization and Applications, Springer, 2019, pp. 155–167.
- [18] ———, *Diameter computation on H -minor free graphs and graphs of bounded (distance) VC-dimension*, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2020, pp. 1905–1922.
- [19] F. Gavril, *The intersection graphs of subtrees in trees are exactly the chordal graphs*, Journal of Combinatorial Theory, Series B **16** (1974), no. 1, 47–56.
- [20] J. Gorzny and J. Huang, *End-vertices of LBFS of (AT-free) bigraphs*, Discrete Applied Mathematics **225** (2017), 87–94.
- [21] 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** (2000), no. 1-2, 59–84.
- [22] H. Hempel and D. Kratsch, *On claw-free asteroidal triple-free graphs*, Discrete Applied Mathematics **121** (2002), no. 1-3, 155–180.
- [23] R. Impagliazzo and R. Paturi, *On the complexity of k -SAT*, Journal of Computer and System Sciences **62** (2001), no. 2, 367–375.
- [24] T. Kloks, D. Kratsch, and H. Müller, *On the structure of graphs with bounded asteroidal number*, Graphs and Combinatorics **17** (2001), no. 2, 295–306.
- [25] D. Kratsch and J. Spinrad, *Between $O(nm)$ and $O(n^\alpha)$* , SIAM Journal on Computing **36** (2006), no. 2, 310–325.
- [26] C. Lekkerkerker and J. Boland, *Representation of a finite graph by a set of intervals on the real line*, Fundamenta Mathematicae **51** (1962), no. 1, 45–64.
- [27] I.-J. Lin, T.A. McKee, and D.B. West, *The leafage of a chordal graph*, Discussiones Mathematicae Graph Theory **18** (1998), no. 1, 23–48.
- [28] R. Machado and C. de Figueiredo, *Linear-time graph distance and diameter approximation*, International Transactions in Operational Research **23** (2016), no. 5, 843–851.
- [29] A. Mohammed and F. Dragan, *Slimness of graphs*, Discrete Mathematics & Theoretical Computer Science **21** (2019).
- [30] S. Olariu, *A simple linear-time algorithm for computing the center of an interval graph*, International Journal of Computer Mathematics **34** (1990), no. 3-4, 121–128.
- [31] L. Roditty and V. Vassilevska Williams, *Fast approximation algorithms for the diameter and radius of sparse graphs*, Proceedings of the forty-fifth annual ACM symposium on Theory of computing (STOC), 2013, pp. 515–524.
- [32] D. Rose, R. Tarjan, and G. Lueker, *Algorithmic aspects of vertex elimination on graphs*, SIAM Journal on computing **5** (1976), no. 2, 266–283.

- [33] R. Seidel, *On the all-pairs-shortest-path problem in unweighted undirected graphs*, J. Comput. Syst. Sci. **51** (1995), no. 3, 400–403.
- [34] V. Vassilevska Williams, *On some fine-grained questions in algorithms and complexity*, Proceedings of the ICM, vol. 3, World Scientific, 2018, pp. 3431–3472.
- [35] V. Vassilevska Williams and R.R. Williams, *Subcubic equivalences between path, matrix, and triangle problems*, Journal of the ACM (JACM) **65** (2018), no. 5, 1–38.
- [36] R. Williams, *A new algorithm for optimal 2-constraint satisfaction and its implications*, Theoretical Computer Science **348** (2005), no. 2-3, 357–365.