Optimization and approximation on systems of geometric objects

van Leeuwen, E.J.

Chapter 6

Density and Unit Disk Graphs

The thickness of a unit disk graph is a good parameter by which to investigate the complexity of various graph optimization problems. We even obtain polynomial-time algorithms if the thickness is small. However, a given unit disk graph might not have small thickness. To alleviate this, we introduce a new notion for unit disk graphs, called *density*. Intuitively, the density of a set of disks is the number of disk centers in any 1×1 box. Using this notion, we are able to give a tight upper bound on the thickness of a unit disk graph.

Moreover, the density is instrumental in the design of a set of new approximation schemes for unit disk graphs. Using a uniform approach, we are able to obtain an eptas on unit disk graphs of bounded density and a ptas on general unit disk graphs for all studied graph optimization problems. These schemes both generalize and improve on previous work on approximation algorithms for unit disk graphs.

6.1 The Density of Unit Disk Graphs

The density of a unit disk graph is defined analogously to the thickness. Assume that we are given an n-vertex unit disk graph G with a representation $\mathcal{D} = \{\mathcal{D}(v) = (c_v, r_v) \mid v \in V(G)\}$, where $c_v \in \mathbb{R}^2$ is the center of the disk corresponding to vertex $v \in V(G)$ and $r_v = 1/2$ is its radius.

The density of a unit disk graph is determined by a grid decomposition of a representation of that graph. Given an angle α ($0 \leq \alpha < \pi/2$) and a point $p \in \mathbb{R}^2$, partition the plane using an infinite grid, such that each grid square has width and height 1, the grid is rotated (clockwise) by α with respect to the x-axis, and the corner of some grid square coincides with p. The horizontal and vertical lines defining the grid are the *horizontal and vertical grid boundaries*. Observe that the partitioning of the plane imposed by the grid remains the same after a rotation of $\pi/2$ around p. Hence it is valid to restrict α to $0 \leq \alpha < \pi/2$.

A disk is said to be *in* a grid square if its center is contained in the interior of the square or the center lies on the left vertical or top horizontal grid boundary determining the square. Given (α, p), this induces a *grid decomposition* of $\mathcal{D}$.

Definition 6.1.1 *Given (α, p), the* density *of a set of disks $\mathcal{D}$ is the maximum number of disks in any grid square induced by the grid decomposition of $\mathcal{D}$ determined by (α, p).*

For any (fixed) angle $0 \leq \alpha < \pi/2$, the *density* $d^*_\alpha(\mathcal{D})$ is the minimum density of any grid decomposition (α, p) over all $p \in \mathbb{R}^2$. The *max-density* $\bar{d}_\alpha(\mathcal{D})$ is the maximum density of any grid decomposition (α, p) over all $p \in \mathbb{R}^2$.

Definition 6.1.2 *The* density $d^*(\mathcal{D})$ *of a set of unit disks $\mathcal{D}$ is the minimum density $d^*_\alpha(\mathcal{D})$ over all angles $0 \leq \alpha < \pi/2$. The* max-density $\bar{d}(\mathcal{D})$ *is the maximum max-density $\bar{d}_\alpha(\mathcal{D})$ over all angles $0 \leq \alpha < \pi/2$.*

The density and max-density of a given set of unit disks can be computed in polynomial time by enumerating all relevant angles and points [258].

Observe that the notion of density is more general than the notion of λ-*precision* unit disk graphs [154], in which the disk centers are at least λ apart.

The studied optimization problems are all NP-hard when restricted to unit disk graphs of bounded density. Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set are NP-hard on arbitrary unit disk graphs [194, 267, 17, 67], even if the degree is at most 3 and (except for Maximum Independent Set and Minimum Vertex Cover) the graph is bipartite [65]. To show NP-hardness in case of bounded density, we can adapt a reduction by Clark, Colbourn, and Johnson [67] from Maximum Independent Set and Minimum Vertex Cover on planar graphs of degree 3 and 4 to the same problems on unit disk graphs, giving the following theorem [258].

Theorem 6.1.3 *Maximum Independent Set and Minimum Vertex Cover are NP-hard on unit disk graphs of density* 1.

Minimum Connected Dominating Set was proved NP-hard on unit disk graphs by Lichtenstein [194]. The instances of Connected Dominating Set constructed in this proof have density 3. The NP-hardness gadget given by Clark, Colbourn, and Johnson [67] however has density 1. This is also true for their gadget for Minimum Dominating Set. These results imply that Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set are NP-hard on unit disk graphs of any (fixed) density.

Because Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set on unit disk graphs of density 1 are NP-hard and the values of their optima are bounded by a polynomial in the instance size, they cannot have an fptas, unless P=NP. In Section 6.4, we exhibit further inapproximability results for these problems.

6.2 Relation to Thickness

A first strategy to deal with the NP-hardness of the studied optimization problems on unit disk graphs of bounded density is to consider fast exact algorithms (with exponential running time). We can do this by bounding the thickness of a unit disk graph in terms of its density.

Theorem 6.2.1 *For any n-vertex unit disk graph with representation $\mathcal{D}$ of max-density $d = \bar{d}(\mathcal{D})$,*

$$t^*(\mathcal{D}) \leq \bar{t}(\mathcal{D}) \leq 5.7 \cdot \sqrt{nd \log n}.$$

Moreover, this bound is tight (up to constants).

Proof: The theorem essentially follows from a result by Alon, Katchalski, and Pulleyblank [10]. We follow their proof.

If $d > n/(16 \log n)$, the theorem is trivial, so assume that $d \leq n/(16 \log n)$. Let $k = \lfloor \sqrt{n}/\sqrt{d \log n} \rfloor$ and note that $k \geq 4$. For each integer $0 \leq i < k/2$, the thickness $\bar{t}_{\alpha_i}(\mathcal{D})$ with $\alpha_i = \pi \cdot i/k$ is equivalent to the maximum number of disks intersecting any line at angle α_i with respect to the x-axis. Let l_i be a line with angle α_i intersecting the largest number of disks of $\mathcal{D}$. Then for any $0 \leq i < k/2$, $\bar{t}_{\alpha_i}(\mathcal{D})$ equals the number of disks intersecting l_i and none of the other l_j plus the number of disks intersecting l_i and at least one other l_j.

We first bound the second quantity. Consider $i \neq j$ with $0 \leq i, j < k/2$. Then using that $\sin \alpha \geq 2\alpha/\pi$ for any $0 \leq \alpha \leq \pi/2$, the disk centers of all disks intersecting both l_i and l_j can be contained in a 2 by $2 + \lceil k/|i-j| \rceil$ rectangle, such that two sides of this rectangle are parallel to l_i. Hence the number of disks intersecting both l_i and l_j is at most $(4 + 2\lceil k/|i-j|\rceil)d$. For any fixed i, the number of disks intersecting l_i and at least one of the other l_j's is at most

$$\begin{aligned} d \cdot 2 \sum_{h=1}^{\lfloor k/2 \rfloor} \left(4 + 2\lceil k/h \rceil\right) &\leq 6kd + 4kd \cdot \sum_{h=1}^{\lfloor k/2 \rfloor} 1/h \\ &\leq 6kd + 4kd \cdot (0.58 + 1/4 + \ln k - \ln 2) \\ &\leq 3kd \log k + 0.28kd \log k + 2.78kd \log k \\ &= 6.06 \cdot kd \log k \\ &\leq 3.03 \cdot \sqrt{dn \log n}. \end{aligned}$$

The first quantity can only be bounded existentially. By the pigeonhole principle, there is a value of i such that the number of disks intersected by l_i is at most $n/(k/2) \leq (8/3) \cdot \sqrt{dn \log n}$. Hence for this value of i, $\bar{t}_{\alpha_i}(\mathcal{D}) \leq 5.7 \cdot \sqrt{dn \log n}$. Therefore $\bar{t}(\mathcal{D}) \leq 5.7 \cdot \sqrt{dn \log n}$.

Adapting a construction by Alon, Katchalski, and Pulleyblank [10], we can give for any $d \leq h$ a set of $\mathrm{O}(dh^2/\log h)$ unit disks of max-density d and of thickness $\Omega(hd)$. In other words, for any n, a set of n unit disks with thickness $\Omega(\sqrt{dn \log n})$ and max-density d exists. $\square$

Using Theorem 5.2.5, we immediately obtain the following corollary.

Corollary 6.2.2 *The strong pathwidth of any n-vertex unit disk graph with representation $\mathcal{D}$ is at most $5.7 \cdot \sqrt{dn \log n}$, where d is the max-density of $\mathcal{D}$.*

This naturally implies a bound of $5.7 \cdot \sqrt{dn \log n}$ on the relaxed pathwidth of a unit disk graph. It is possible to improve considerably on this bound though. Van Leeuwen [258] shows that for any set of unit disks $\mathcal{D}$ there exists a slab of width 1 containing at most $(2 + 4/\pi)\sqrt{dn} + \mathrm{o}(\sqrt{dn})$ disks such that the disks outside the slab are partitioned into two pieces of at most $2n/3$ pieces each. In other words, the unit disk graph has a $\sqrt{dn}$-separator theorem. Smith and Wormald [244] show that the constant in this bound can be further improved to $2\sqrt{dn}$ using a circular separator. Using a result of Bodlaender [36], this implies the following bound on the relaxed pathwidth.

Theorem 6.2.3 *The (relaxed) pathwidth of any n-vertex unit disk graph with representation $\mathcal{D}$ of max-density d is at most $6\sqrt{dn}$.*

One can use these bounds on the strong and relaxed pathwidth to analyze the worst-case running times of the algorithms given in the previous chapter.

Theorem 6.2.4 *For a n-vertex unit disk graph with representation $\mathcal{D}$ of max-density d, Maximum Independent Set and Minimum Vertex Cover can be solved in $\mathrm{O}(n\sqrt{dn}\, 2^{6\sqrt{dn}})$ time, Minimum Dominating Set in $\mathrm{O}(n\sqrt{dn}\, 3^{6\sqrt{dn}})$ time, and Minimum Connected Dominating Set in $\mathrm{O}(dn^2\, 2^{22.8\sqrt{nd\log n}})$ time.*

This follows from Theorem 5.3.3, 5.3.4, 5.3.9, and 5.4.5.

Further improvement follows from work by Fu [112]. He showed that if $d = 1$, a $1.2126\sqrt{n}$-separator exists. By mapping disk centers to a grid and then using this separator, Fu shows the following.

Theorem 6.2.5 (Fu [112]) *Maximum Independent Set and Minimum Vertex Cover can be solved in $O^*(2^{\mathrm{O}(\sqrt{n})})$ time.*

The $O^*(\cdot)$ means that we omit polynomially bounded terms. The technique used by Fu is believed to extend to Minimum Dominating Set as well. We conjecture that using the techniques developed in Section 5.4, one can obtain an $O^*(2^{\mathrm{O}(\sqrt{n})})$ time algorithm for Minimum Connected Dominating Set.

In this context of exact algorithms, we should also mention results on the parameterized complexity of these problems. Alber and Fiala [7] gave an $n^{\mathrm{O}(\sqrt{k})}$-time algorithm to determine whether a unit disk graph has an independent set of cardinality at least k. If the unit disk graph has constant precision, this improves to $2^{\mathrm{O}(\sqrt{k})}$ time. Marx [202, 203] showed however that Maximum Independent Set and Minimum Dominating Set are W[1]-hard on arbitrary unit disk graphs. Hence it is unlikely that these problems are fixed-parameter tractable, unless FPT=W[1].

6.3 Approximation Schemes

Another way to get around the NP-hardness of the graph optimization problems on unit disk graphs of bounded density is to restrict to a polynomial running time, but allow the algorithm to return an approximation to the optimum. In particular, we are interested in approximation schemes, giving a $(1 + \epsilon)$-approximation for any $\epsilon > 0$. We present a unified approach that yields optimal approximation schemes for Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set on unit disk graphs with a known representation. The density of this representation is crucial to the analysis. For each of the aforementioned problems, we give an approximation scheme that is both an eptas if the density is bounded and a ptas in the general case. The running times of these schemes improves on the running times achieved by previous schemes for these problems.

The approximation schemes use the shifting technique, originally proposed by Baker [22] and Hochbaum and Maass [150]. Here we use a decomposition of the disks similar to one proposed by Hunt et al. [154].

Assume that we are given a unit disk graph G with representation $\mathcal{D}$, such that each disk in $\mathcal{D}$ has radius $1/2$. First, we find a grid decomposition (α, p) of $\mathcal{D}$ of minimum density $d = d^*(\mathcal{D})$. We may assume that $\alpha = 0$ and $p = (0, 0)$. Now we can speak of the columns and the rows of the grid decomposition, i.e. row r_i for some $i \in \mathbb{Z}$ contains all grid squares between the lines $y = i$ and $y = i + 1$. The idea of the proposed schemes is to group (the disks in) several consecutive rows together in a *strip*. Decomposing the plane in this way, we obtain a *strip decomposition*. The strips will each have bounded thickness, making it easier to solve the problems we consider. We then combine the solutions of these subproblems to a solution for the global problem. By repeating this for several appropriately constructed strip decompositions, we show that for (at least) one strip decomposition, the solution we obtain is the required approximation to the optimum.

6.3.1 Maximum Independent Set

Let $k \geq 2$ be an integer (whose precise value we determine later). Decompose the rows of the grid such that the b-th strip consists of rows r_i with $bk + 1 \leq i \leq (b + 1)k - 1$ for any $b \in \mathbb{Z}$. Observe that rows where $i \equiv 0 \pmod{k}$ are not in any strip. Hence the strips can be thought of as being independent. Let $\mathcal{D}^b \subseteq \mathcal{D}$ denote the set of disks contained in the b-th strip and $S^b \subseteq \mathcal{D}$ the set of disks contained in row r_{bk} or $r_{(b+1)k}$.

Lemma 6.3.1 *For any $b \in \mathbb{Z}$, the thickness of $\mathcal{D}^b$ is at most $(k - 1)d$.*

Proof: The columns of the grid decomposition induce a slab decomposition. Any column contains $k - 1$ grid squares of the b-th strip and thus the centers of at most $(k - 1)d$ disks. The lemma follows. □

Using this lemma, we can already conclude from Theorem 5.3.3 that for any $b \in \mathbb{Z}$, one can compute the cardinality of a maximum independent set in $O(n\, 2^{2kd})$ time. We can improve on this by more refined analysis.

We require the following auxiliary results. Let $e = 2.718...$ be the base of the natural logarithm.

Lemma 6.3.2 *Let c_1, s be positive integers and $c_2 \geq 1$ a number. Then a set of cardinality $c_1 s$ has at most $c_2 s \cdot (c_1 e)^{c_2 s}$ distinct subsets of cardinality at most $\lfloor c_2 s \rfloor$.*

Proof: Suppose that $\lfloor c_2 s \rfloor < c_1 s/2$. Using Åslund's [18] upper bound on the binomial coefficient and that the function x^x is convex,

$$\begin{aligned}
\binom{c_1 s}{\lfloor c_2 s \rfloor} &\leq \frac{(c_1 s)^{c_1 s}}{(\lfloor c_2 s \rfloor)^{\lfloor c_2 s \rfloor}\,(c_1 s - \lfloor c_2 s \rfloor)^{c_1 s - \lfloor c_2 s \rfloor}} \\
&\leq \frac{(c_1 s)^{c_1 s}}{(c_2 s)^{c_2 s}\,(c_1 s - c_2 s)^{c_1 s - c_2 s}} \\
&= \left(\frac{(c_1 s)^{c_1}}{(c_2 s)^{c_2}\,(c_1 s - c_2 s)^{c_1 - c_2}}\right)^s \\
&= \left(\frac{c_1^{c_1}}{c_2^{c_2}\,(c_1 - c_2)^{c_1 - c_2}}\right)^s \\
&\leq \left(\frac{c_1^{c_2} \cdot c_1^{c_1 - c_2}}{(c_1 - c_2)^{c_1 - c_2}}\right)^s \\
&= \left(c_1^{c_2} \cdot \left(\frac{c_1}{c_1 - c_2}\right)^{c_1 - c_2}\right)^s \\
&= \left(c_1^{c_2} \cdot \left(1 + \frac{c_2}{c_1 - c_2}\right)^{c_1 - c_2}\right)^s \\
&\leq (c_1 e)^{c_2 s}.
\end{aligned}$$

Hence the number of subsets is at most $c_2 s \cdot (c_1 e)^{c_2 s}$.

If $\lfloor c_2 s \rfloor \geq c_1 s/2$, the number of distinct subsets is at most $2^{c_1 s}$. If $c_1 \geq 2$, then $2^{c_1 s} \leq 2^{2 c_2 s} = 4^{c_2 s} \leq (c_1 e)^{c_2 s}$. If $c_1 = 1$, then $2^{c_1 s} \leq (c_1 e)^{c_1 s} \leq (c_1 e)^{c_2 s}$. The lemma follows. □

Lemma 6.3.3 *Consider the slab decomposition induced by Lemma 6.3.1. The maximum cardinality of any independent set of the disks in any $c \geq 1$ consecutive slabs is at most $4(c+1)k/\pi$.*

Proof: All disks in these c slabs are contained in an appropriately placed $c+1$ by k rectangle. A simple area bound gives the lemma. □

Lemma 6.3.4 *For any $b \in \mathbb{Z}$, one can compute a maximum independent set I^b of $\mathcal{D}^b$ in $\mathrm{O}(k^2 dn\,(ed)^{12k/\pi})$ time.*

Proof: Consider the algorithm for computing a maximum independent set as described in the proof of Lemma 5.3.1 in the case where we have a strong path decomposition. For any i and any independent set $A_i \subseteq X_i$,

$$\mathsf{size}_i(A_i) = \max_{A_{i-1}}\{|A_{i-1}| + \mathsf{size}_{i-1}(A_{i-1})\},$$

where the maximum is over all independent sets $A_{i-1} \subseteq X_{i-1} - N(A_i)$. Furthermore, $X_0 = \emptyset$ and $\mathsf{size}_0(\emptyset) = 0$.

Assume we are given a strong path decomposition induced by Lemma 6.3.1. It suffices to enumerate those sets A_i and A_{i-1} for which $A_i \cup A_{i-1}$ is an independent set. Following Lemma 6.3.3, no independent set of two consecutive slabs has cardinality more than $12k/\pi$. By Lemma 6.3.1, $|X_i| + |X_{i-1}| \leq 2(k-1)d < 2kd$. Then Lemma 6.3.2 gives that all of these independent sets can be enumerated in $\mathrm{O}(k\,(ed)^{12k/\pi})$ time. The lemma follows. □

Recall that a *separation* of a graph G is a pair $\{A, B\}$ such that $A \cup B = V(G)$ and there is no path in G from $A - B$ to $B - A$.

Lemma 6.3.5 $\bigcup_{b\in\mathbb{Z}} I^b$ *is an independent set.*

Proof: In general, it is easy to see that the following is true. If G is a graph and $\{A, B\}$ is any separation of G, then given independent sets $I^A \subseteq A - B$ and $I^B \subseteq B - A$, $I^A \cup I^B$ is an independent set of G. By observing that $\{\mathcal{D}^b \cup S^b, \mathcal{D} - \mathcal{D}^b\}$ induces a separation for any $b \in \mathbb{Z}$ and recursively applying the preceding observation, we prove the lemma. □

Now apply the shifting technique. For each integer $0 \leq a \leq k-1$ (the *shifting parameter*), we define a strip decomposition as follows. The b-th strip consists of rows r_i with $bk + 1 + a \leq i \leq (b+1)k - 1 + a$, i.e. rows with $i \equiv a \pmod{k}$ are not in any strip. This induces a strip decomposition as before (note that for $a = 0$, it actually is the same). Hence we can use Lemma 6.3.4 to compute a maximum independent set of these strips.

For each integer $0 \leq a \leq k-1$ and $b \in \mathbb{Z}$, let $\mathcal{D}_a^b$ denote the set of disks contained in the b-th strip induced by shifting parameter a and let I_a^b be the independent set returned by the algorithm of Lemma 6.3.4 in this case. Let $I_a = \bigcup_{b\in\mathbb{Z}} I_a^b$ and let $I_{\max}$ denote a largest such set.

Lemma 6.3.6 $|I_{\max}| \geq (1 - 1/k) \cdot |\mathcal{I}|$, *where $\mathcal{I}$ is a maximum independent set of G.*

Proof: Because I_a^b is a maximum independent set of $\mathcal{D}_a^b$, $|I_a^b| \geq |\mathcal{I} \cap \mathcal{D}_a^b|$. Let $\mathcal{D}_a = \bigcup_{b\in\mathbb{Z}} \mathcal{D}_a^b$. Then $|I_a| \geq |\mathcal{I} \cap \mathcal{D}_a|$. Observe that a disk is in $\mathcal{D}_a$ for precisely $k-1$ values of a. Hence

$$k \cdot |I_{\max}| \;\geq\; \sum_{a=0}^{k-1} |I_a| \;\geq\; \sum_{a=0}^{k-1} |\mathcal{I} \cap D_a| \;=\; (k-1) \cdot |\mathcal{I}|,$$

and thus $|I_{\max}| \geq (1 - 1/k) \cdot |\mathcal{I}|$. □

Combining Lemma 6.3.4 and Lemma 6.3.6, we obtain the following.

Lemma 6.3.7 *For any $k \geq 2$, one can obtain a $(1 - 1/k)$-approximation for Maximum Independent Set on n-vertex unit disk graphs G with a known representation $\mathcal{D}$ of density d in $\mathrm{O}(k^3 n^2 d\,(2ed)^{12k/\pi})$ time.*

Proof: There are at most n nonempty strips for each of the k values of a. Hence one can compute $I_{\max}$ in $\mathrm{O}(k^3 n^2 d\,(ed)^{12k/\pi})$ time. □

Since the notion of density is more general than the notion of λ-precision, the scheme presented above is more general than the scheme given by Hunt et al. [154] on unit disk graphs of constant precision. Moreover, the above scheme has a better running time.

Theorem 6.3.8 *There is an eptas for Maximum Independent Set on unit disk graphs with n vertices and bounded density, i.e. density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$.*

Proof: Consider any $\epsilon > 0$. Choose k as the largest integer such that $(12k/\pi) \cdot \log(ed) \leq \log n$. If $k < 2$, output any single vertex. Otherwise, apply the algorithm of Lemma 6.3.7 and compute $I_{\max}$ in $\mathrm{O}(n^4 \log^3 n)$ time. Furthermore, if $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$, there is a c_ϵ such that $k \geq 1/\epsilon$ and $k \geq 2$ for all $n \geq c_\epsilon$. Therefore, if $n \geq c_\epsilon$, it follows from Lemma 6.3.6 and the choice of k that $I_{\max}$ is a $(1 - \epsilon)$-approximation of the optimum. Hence there is a fiptas$^\omega$ for Maximum Independent Set on n-vertex unit disk graphs of bounded density, i.e. of density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$. The theorem follows from Theorem 2.2.4. □

Observe that d is always bounded by n. Hence the worst-case running time of the algorithm described in Lemma 6.3.7 is $\mathrm{O}(k^3 n^3\,(en)^{12k/\pi})$.

Theorem 6.3.9 *There is a ptas for Maximum Independent Set on unit disk graphs.*

The ptas given here matches the $n^{\mathrm{O}(1/\epsilon)}$-time ptas given by Hunt et al. [154].

6.3.2 Minimum Vertex Cover

There are (at least) two ways to give an approximation scheme for Minimum Vertex Cover on unit disk graphs. We can a) transfer the ideas of the previous paragraph to Minimum Vertex Cover or b) use the approximation scheme for Maximum Independent Set as a black box. We present both approaches.

We first use the scheme for Maximum Independent Set as a black box. Recall that an independent set is the complement of a vertex cover.

Lemma 6.3.10 *For some $m > 1$, let G be a nonempty graph with no isolated vertices and no $K_{1,m}$ induced subgraph. For any $k \geq 1$, if $\mathcal{C}$ is a minimum vertex cover of G, $\mathcal{I}$ is a maximum independent set of G, and I any independent set of G for which $|I| \geq \left(1 - \frac{1}{(2m-1)k}\right) \cdot |\mathcal{I}|$, then $|V(G) - I| \leq (1 + 1/k) \cdot |\mathcal{C}|$.*

Proof: The proof is essentially due to Wiese and Kranakis [270]. We claim that $|V(G)| \leq 2m|\mathcal{C}|$. Let M be a maximal matching of G and $V(M)$ the set of its endpoints. Consider $V(G)-V(M)$. Since M is a maximal matching, no two vertices in $V(G) - V(M)$ are adjacent. Hence, as G is $K_{1,m}$-free, no vertex in $V(M)$ is adjacent to more than $m-1$ vertices in $V(G)-V(M)$. It follows that $|V(G)-V(M)| \leq (m-1)|V(M)|$ and thus $|V(G)| = |V(G)-V(M)|+|V(M)| \leq m \cdot |V(M)|$. Now observe that any vertex cover of G must contain at least one endpoint of each edge in M. Then $|V(G)| \leq m \cdot |V(M)| \leq 2m \cdot |\mathcal{C}|$.

Using this claim,

$$\begin{aligned}
|V(G) - I| &\leq |V(G)| - \left(1 - \frac{1}{(2m-1)k}\right) \cdot |\mathcal{I}| \\
&= |V(G)| - \left(1 - \frac{1}{(2m-1)k}\right) \cdot (|V(G)| - |\mathcal{C}|) \\
&= |\mathcal{C}| + \frac{1}{(2m-1)k} \cdot (|V(G)| - |\mathcal{C}|) \\
&\leq |\mathcal{C}| + \frac{1}{(2m-1)k} \cdot (2m \cdot |\mathcal{C}| - |\mathcal{C}|) \\
&= (1 + 1/k) \cdot |\mathcal{C}|
\end{aligned}$$

The lemma follows. $\square$

We know that unit disk graphs have no $K_{1,6}$ induced subgraph. Combining this observation with Lemma 6.3.7 and Lemma 6.3.10, we obtain the following.

Lemma 6.3.11 *For any $k \geq 1$, one can obtain a $(1+1/k)$-approximation for Minimum Vertex Cover on n-vertex unit disk graphs with a known representation $\mathcal{D}$ of density d in $\mathrm{O}(k^3 n^2 d\,(ed)^{132k/\pi})$ time.*

Even though this already leads to an approximation scheme, we can improve on the running time of the scheme by transferring the ideas of the previous paragraph to Minimum Vertex Cover.

Let $k \geq 2$ be an integer. For each integer $0 \leq a \leq k-1$ and $b \in \mathbb{Z}$, the b-th strip consists of rows r_i with $bk + a \leq i \leq (b+1)k + a$, i.e. rows with $i \equiv a \pmod{k}$ are in two strips. Define $\mathcal{D}_a^b$ to be the set of disks in the b-th strip induced by shifting parameter a and $S_a^b = (\mathcal{D}_a^{b-1} \cap \mathcal{D}_a^b) \cup (\mathcal{D}_a^b \cap \mathcal{D}_a^{b+1})$ as the set of disks in row r_{bk+a} or $r_{(b+1)k+a}$.

Following Lemma 6.3.4, one can compute a minimum vertex cover C_a^b of $\mathcal{D}_a^b$ in $\mathrm{O}(k^2 nd\,(ed)^{12(k+2)/\pi})$ time. Let $C_a = \bigcup_{b \in \mathbb{Z}} C_a^b$ and let $C_{\min}$ be a smallest such set.

Lemma 6.3.12 *Any C_a is a vertex cover of G and $|C_{\min}| \leq (1+1/k) \cdot |\mathcal{C}|$, where $\mathcal{C}$ is a minimum vertex cover of G.*

Proof: The following is true in general. If $\{A, B\}$ is a separation of a graph G, then given vertex covers $C^A \subseteq A$ and $C^B \subseteq B$ of A and B respectively, $C^A \cup C^B$

is a vertex cover of G. Observing that $\{D_a^b, \mathcal{D} - (\mathcal{D}_a^b - S_a^b)\}$ is a separation of G for any $b \in \mathbb{Z}$ and recursively applying the preceding observation, we prove that C_a is a vertex cover of G for any value of a.

Because C_a^b is a minimum vertex cover of $\mathcal{D}_a^b$,

$$|C_a| \quad \leq \quad \sum_{b \in \mathbb{Z}} |C_a^b| \quad \leq \quad \sum_{b \in \mathbb{Z}} |\mathcal{D}_a^b \cap \mathcal{C}| \quad = \quad |\mathcal{C}| + \tfrac{1}{2} \sum_{b \in \mathbb{Z}} |S_a^b \cap \mathcal{C}|.$$

A vertex is in $\bigcup_{b \in \mathbb{Z}} S_a^b$ for precisely one value of a. For this value of a, it is in S_a^b for precisely two values of b. Hence

$$k \cdot |C_{\min}| \quad \leq \quad \sum_{a=0}^{k-1} |C_a| \quad \leq \quad k|\mathcal{C}| + \tfrac{1}{2} \sum_{a=0}^{k-1} \sum_{b \in \mathbb{Z}} |S_a^b \cap \mathcal{C}| \quad = \quad k|\mathcal{C}| + |\mathcal{C}|$$

and thus $|C_{\min}| \leq (1 + 1/k) \cdot |\mathcal{C}|$. □

We can now offer the following improvement on Lemma 6.3.11.

Lemma 6.3.13 *For any $k \geq 2$, one can obtain a $(1 + 1/k)$-approximation for Minimum Vertex Cover on n-vertex unit disk graphs G with a known representation $\mathcal{D}$ of density d in $\mathrm{O}(k^3 n^2 d\,(ed)^{12(k+2)/\pi})$ time.*

As in Theorem 6.3.8 and Theorem 6.3.9, we can now prove the existence of an eptas for Minimum Vertex Cover on unit disk graphs of bounded density and a ptas on arbitrary unit disk graphs. The scheme on unit disk graphs of bounded density generalizes the scheme by Hunt et al. [154] on unit disk graphs of bounded precision. Moreover, we attain a better running time.

We can prove a better result however, owing to an idea by Marx [202].

Lemma 6.3.14 *For any $k \geq 1$, let G_k be the graph obtained from G by iteratively removing all cliques with at least $k + 1$ vertices. If C if a $(1 + 1/k)$-approximation for Minimum Vertex Cover on G_k, then $C \cup (V(G) - V(G_k))$ is a $(1 + 1/k)$-approximation on G.*

Proof: Observe that for any clique K of G, any vertex cover of G must contain either $|V(K)|$ or $|V(K) - 1|$ vertices of K. Hence if $|V(K)| \geq k + 1$, then

$$|V(K)| \quad \leq \quad (1 + 1/k) \cdot (|V(K)| - 1) \quad \leq \quad (1 + 1/k) \cdot |\mathcal{C} \cap V(K)|$$

for a minimum vertex cover $\mathcal{C}$ of G. Let $\mathcal{K} = \{K_1, \ldots, K_p\}$ be any sequence of cliques with at least $k + 1$ vertices whose sequential removal results in G_k. As

$$|C| \quad \leq \quad (1 + 1/k) \cdot |\mathcal{C}(G_k)| \quad \leq \quad (1 + 1/k) \cdot |\mathcal{C} \cap G_k|,$$

it follows that

$$|C \cup (V(G) - V(G_k))|$$

$$
\begin{aligned}
&= |C| + \sum_{K \in \mathcal{K}} |V(K)| \\
&\le (1 + 1/k) \cdot |\mathcal{C} \cap G_k| + \sum_{K \in \mathcal{K}} \big((1 + 1/k) \cdot |\mathcal{C} \cap V(K)|\big) \\
&= (1 + 1/k) \cdot |\mathcal{C}|.
\end{aligned}
$$

Note that $V(G_k), V(K_1), \ldots, V(K_p)$ are pairwise disjoint sets. Moreover, for any clique K of G, C is a vertex cover of $G - K$ if and only if $C \cup V(K)$ is a vertex cover of G. The lemma follows. □

Clark, Colbourn, and Johnson [67] showed that Maximum Clique can be solved in $\mathrm{O}(n^{9/2})$ time on unit disk graphs. Hence we can reduce a unit disk graph G to a graph G_k (for any $k \ge 1$) in $\mathrm{O}(n^{11/2}/k)$ time.

Theorem 6.3.15 *There is an eptas for Minimum Vertex Cover on unit disk graphs.*

Proof: For any $\epsilon > 0$, let $k = \lceil 1/\epsilon \rceil$. Let G be a unit disk graph and reduce it to G_k. Following Lemma 6.3.13, one can obtain a $(1 + 1/k)$-approximation for Minimum Vertex Cover on G_k in $\mathrm{O}(k^3 n^2 d\, (ed)^{12(k+2)/\pi})$ time. As G_k contains only cliques of size k or less, the density is at most $4k$. Applying Lemma 6.3.14, one can obtain a $(1 + 1/k)$-approximation for Minimum Vertex Cover on G in $\mathrm{O}(k^4 n^2\, (4ek)^{12(k+2)/\pi} + n^{11/2}/k)$ time. □

The above $2^{\mathrm{O}(\epsilon^{-1} \log \epsilon^{-1})}$-time scheme improves on the earlier $2^{\mathrm{O}(\epsilon^{-2})}$-time scheme by Marx [202].

6.3.3 Minimum Dominating Set

The analysis for the scheme for Minimum Dominating Set is slightly more involved. Let $k \ge 3$ be an integer. For each integer $0 \le a \le k - 1$ and $b \in \mathbb{Z}$, the b-th strip consists of rows r_i with $bk + a \le i \le (b + 1)k + a + 1$, i.e. rows with $i \equiv a \pmod{k}$ and $i \equiv a + 1 \pmod{k}$ are in two strips. Define $\mathcal{D}_a^b$ to be the set of disks in the b-th strip induced by shifting parameter a, $S_a^b = (\mathcal{D}_a^{b-1} \cap \mathcal{D}_a^b) \cup (\mathcal{D}_a^b \cap \mathcal{D}_a^{b+1})$, and N_a^b as the set of disks in r_{bk+a} or $r_{(b+1)k+a+1}$.

Lemma 6.3.16 *For any $0 \le a \le k - 1$ and any $b \in \mathbb{Z}$, one can compute a minimum set $C_a^b \subseteq \mathcal{D}_a^b$ dominating $\mathcal{D}_a^b - N_a^b$ in $\mathrm{O}(k^2 nd\, (ed)^{24(k+3)/\pi})$ time.*

Proof: We use the algorithm described in Theorem 5.3.6. Observe that it suffices to enumerate for three consecutive slabs all possible sets of disks that can be in a minimum dominating set. Disks in these slabs dominate vertices in five consecutive slabs, adding the slab to the left and the one to the right of the original three slabs. As any maximal independent set is a dominating set, a set of disks in the three slabs in a minimum dominating set should not have cardinality more than $24(k+3)/\pi$, according to Lemma 6.3.3. As three consecutive slabs contain at most $3(k+2)d$ vertices, it follows from Lemma 6.3.2 and Theorem 5.3.6 that the algorithm takes $\mathrm{O}(k^2 nd\,(ed)^{24(k+3)/\pi})$ time. $\square$

Let $C_a = \bigcup_{b\in\mathbb{Z}} C_a^b$ for any value of a and $C_{\min}$ a smallest such set. We have to show that $C_{\min}$ is a dominating set that approximates the optimum well.

Definition 6.3.17 *The pair $\{A, B\}$ is a* double separation *of a graph G if $A \cup B = V(G)$ and there is no 1- or 2-edge path in G from $A - B$ to $B - A$.*

Lemma 6.3.18 *Any C_a is a dominating set and $|C_{\min}| \leq (1+2/k)\cdot|\mathcal{C}|$, where $\mathcal{C}$ is a minimum dominating set of G.*

Proof: The following is true in general. If $\{A, B\}$ is a double separation of G, then given sets $C^A \subseteq A$ and $C^B \subseteq B$ dominating $(A - B) \cup M$ and $(B-A)\cup((A\cap B)-M)$ respectively for some subset $M \subseteq A\cap B$, $C^A \cup C^B$ is a dominating set of G. Observing that $\{\mathcal{D}_a^b, \mathcal{D}-(\mathcal{D}_a^b - S_a^b)\}$ is a double separation of G for any $b \in \mathbb{Z}$ and recursively applying the preceding observation, we prove that C_a is a dominating set of G for any value of a.

Because $C_a^b \subseteq \mathcal{D}_a^b$ is a smallest set dominating $\mathcal{D}_a^b - N_a^b$,

$$|C_a| \;\leq\; \sum_{b\in\mathbb{Z}} |C_a^b| \;\leq\; \sum_{b\in\mathbb{Z}} |\mathcal{C}\cap\mathcal{D}_a^b| \;=\; |\mathcal{C}| + \tfrac{1}{2}\sum_{b\in\mathbb{Z}} |\mathcal{C}\cap S_a^b|.$$

A vertex is in $\bigcup_{b\in\mathbb{Z}} S_a^b$ for precisely two values of a. For these values of a, it is in S_a^b for precisely two values of b. Hence

$$k\cdot|C_{\min}| \;\leq\; \sum_{a=0}^{k-1} |C_a| \;\leq\; k|\mathcal{C}| + \tfrac{1}{2}\sum_{a=0}^{k-1}\sum_{b\in\mathbb{Z}} |\mathcal{C}\cap S_a^b| \;=\; k|\mathcal{C}| + 2|\mathcal{C}|$$

and thus $|C_{\min}| \leq (1+2/k)\cdot|\mathcal{C}|$. $\square$

We can conclude the following.

Lemma 6.3.19 *For any $k \geq 3$, one can obtain a $(1+2/k)$-approximation for Minimum Dominating Set on n-vertex unit disk graphs G with a known representation $\mathcal{D}$ of density d in $\mathrm{O}(k^3 n^2 d\,(ed)^{24(k+3)/\pi})$ time.*

This scheme generalizes the scheme by Hunt et al. [154] on unit disk graphs of bounded precision. Moreover, we attain a better running time.

Theorem 6.3.20 *There is an eptas for Minimum Dominating Set on unit disk graphs with n vertices and bounded density, i.e. density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$.*

Proof: Consider any number $\epsilon > 0$. Choose k as the largest integer such that $(24(k+3)/\pi)\cdot\log(ed) \leq \log n$. If $k < 3$, output $V(G)$. Otherwise, apply the algorithm of Lemma 6.3.19 and compute $C_{\min}$ in $\mathrm{O}(n^4 \log^3 n)$ time.

Furthermore, if $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$, there is a c_ϵ such that $k \geq 2/\epsilon$ and $k \geq 3$ for all $n \geq c_\epsilon$. Therefore, if $n \geq c_\epsilon$, it follows from Lemma 6.3.18 and the choice of k that $C_{\min}$ is a $(1+\epsilon)$-approximation of the optimum. Hence there is a fiptas$^\omega$ for Minimum Dominating Set on n-vertex unit disk graphs of bounded density, i.e. of density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$. The theorem now follows from Theorem 2.2.4. $\square$

Observe that d is always bounded by n. Hence the worst-case running time of the algorithm described in Lemma 6.3.19 is $\mathrm{O}(k^3 n^3 \,(en)^{24(k+3)/\pi})$.

Theorem 6.3.21 *There is a ptas for Minimum Dominating Set on unit disk graphs.*

The ptas given here improves on the $n^{\mathrm{O}(\epsilon^{-2})}$-time ptas given by Hunt et al. [154].

6.3.4 Minimum Connected Dominating Set

The problems treated thus far are very much local problems, where a solution can be verified by just considering the neighborhood of each vertex. Connectivity is a global constraint on the solution and hence tougher to satisfy. We show however that in the case of Minimum Connected Dominating Set, this global property can be dealt with efficiently.

We start by proving some auxiliary results, which hold not only on unit disk graphs, but on arbitrary graphs as well. Throughout this entire section, we assume graphs to be connected.

Definition 6.3.22 *The pair $\{A, B\}$ is a* quadruple separation *of a graph G if $A \cup B = V(G)$, and there is no* 1-, 2-, 3-, *or* 4-*edge path in G from $A - B$ to $B - A$.*

Lemma 6.3.23 *Let $\{A, B\}$ be a quadruple separation of some graph G. Let $C^A \subseteq A$ and $C^B \subseteq B$ form a set dominating $A - N(B-A)$ and $B - N(A-B)$ respectively such that $W \cap C^A$ and $X \cap C^B$ are connected for each connected component W and X of respectively A and B. Then $C^A \cup C^B$ is a connected dominating set of G.*

Proof: As $\{A, B\}$ is a quadruple separation, $(A - N(B-A)) \cup (B - N(A-B)) = V(G)$ and thus $C^A \cup C^B$ is a dominating set of G. Suppose that $C^A \cup C^B$ is not connected and let Y and Z be two distinct connected components of $C^A \cup C^B$. Consider a shortest Y–Z path $P = p_1 \dots p_m$ such that $p_1 \in Y$. By assumption $m \geq 3$. Since $\{A, B\}$ is a quadruple separation, either $p_1, p_2, p_3 \in A - N(B-A)$ or $p_1, p_2, p_3 \in B - N(A-B)$. Without loss of generality, assume that $p_1, p_2, p_3 \in A - N(B-A)$. Then there is a vertex $v \in C^A$ dominating p_3. Moreover $v \in Y$, as p_1 and v belong to the same connected component of A. Therefore $vp_3 \dots p_m$ is a shorter Y–Z path than P, a contradiction. Hence $C^A \cup C^B$ is connected. $\square$

Suppose that C^A has minimum cardinality under the above constraints. We show that there is an upper bound to the cardinality of C^A in terms of the cardinality of a minimum connected dominating set.

Proposition 6.3.24 *If G is a connected graph and S an arbitrary dominating set of G such that S has c connected components, then G has a connected dominating set of cardinality at most $|S| + 2(c-1)$.*

Proof: The case $c = 1$ is trivial. If $c > 1$, then since S is a dominating set, there exist two connected components of S that can be connected by adding at most two vertices to the set. Now apply induction. $\square$

Lemma 6.3.25 *Let $\{A, B\}$ be a quadruple separation of a graph G and let $C^A \subseteq A$ be a smallest set dominating $A - N(B - A)$ such that $C^A \cap Z$ is connected for each connected component Z of A. If C is any connected dominating set of G, then $|C^A| \leq |C \cap A| + 2 \cdot |N(B - A) \cap C|$.*

Proof: Clearly, $C \cap A$ is a dominating set of $A - N(B - A)$. However, for each connected component Z of A, $C \cap Z$ might consists of several connected components. Observe that each such connected component must intersect $N(B - A)$, as $\{N(B - A), V(G)\}$ is a separation of G. Hence the number of connected components of $C \cap Z$ is at most $|N(B - A) \cap C \cap Z|$. By Proposition 6.3.24, we can augment $C \cap A$ to C' such that $C' \cap Z$ is connected by adding at most $2|N(B - A) \cap C \cap Z|$ vertices. Applying this to each connected component of A, we obtain a set $C' \subseteq A$ dominating $A - N(B - A)$ such that $C' \cap Z$ is connected for each connected component Z of A and

$$\begin{aligned} |C'| &\leq |C \cap A| + 2 \cdot \textstyle\sum_Z |N(B - A) \cap C \cap Z| \\ &\leq |C \cap A| + 2 \cdot |N(B - A) \cap C|. \end{aligned}$$

But then

$$|C^A| \quad \leq \quad |C'| \quad \leq \quad |C \cap A| + 2 \cdot |N(B - A) \cap C|$$

and the lemma follows. $\square$

Lemma 6.3.26 *Let $U \subseteq V(G)$ for some graph G. If C is a connected dominating set of G and $C_U \subseteq N[U]$ a set dominating $N[U]$ such that $C_U \cap Z$ is connected for each connected component Z of $N[U]$, then $C' = (C - U) \cup C_U$ is a connected dominating set.*

Proof: Observe that $C - U$ is a dominating set of $V(G) - N[U]$. As C_U dominates $N[U]$, C' is a dominating set. It remains to prove that C' is connected. To this end, we prove the following claim: If $s, t \in C$, then for any $s' \in N[s] \cap C'$ and $t' \in N[t] \cap C'$, there is an s'–t' path in C'.

Note that C contains an s–t path Q. Consider $Q \cap N[U]$. If this is nonempty, it consists of one or more subpaths $Q_1, \ldots, Q_m$. For each such path Q_i, consider its start and end vertices s_i, t_i. Because C_U dominates $N[U]$, there exist vertices $s'_i \in N[s_i] \cap C_U$ and $t'_i \in N[t_i] \cap C_U$ (if possible, let $s'_i = s'$ and $t'_i = t'$). As $Q_i \subseteq N[U]$, s'_i and t'_i are in the same connected component of C_U. Hence there is an s'_i–t'_i path Q'_i in C_U. Let Q''_i be the path induced by s_i (if $s_i \notin U$), Q'_i, and t_i (if $t_i \notin U$). Replace Q_i by Q''_i. This gives an s'–t' path in C'.

Suppose that C' is not connected. Consider two distinct connected components Y and Z of C' and let $s' \in Y$ and $t' \in Z$. Because C is a dominating set, there exist vertices $s \in N[s'] \cap C$ and $t \in N[t'] \cap C$. But then it follows from the above claim that C' has an s'–t' path, contradicting that Y and Z are distinct connected components of C'. The lemma follows. □

Now we apply these ideas to unit disk graphs, together with the shifting technique. Let $k \geq 5$ be an integer. For each integer $0 \leq a \leq k-1$ and $b \in \mathbb{Z}$, the b-th strip consists of rows r_i with $bk + a \leq i \leq (b+1)k + a + 3$, i.e. rows with $i \equiv a \pmod{k}$, $i \equiv (a+1) \pmod{k}$, $i \equiv (a+2) \pmod{k}$, and $i \equiv (a+3) \pmod{k}$ are in two strips. Define $\mathcal{D}_a^b$ to be the set of disks in the b-th strip induced by shifting parameter a. Let $S_a^b = (\mathcal{D}_a^{b-1} \cap \mathcal{D}_a^b) \cup (\mathcal{D}_a^b \cap \mathcal{D}_a^{b+1})$ and let N_a^b be the set of disks in r_{bk+a} or $r_{(b+1)k+a+3}$.

Lemma 6.3.27 *For any $0 \leq a \leq k-1$ and any $b \in \mathbb{Z}$, one can compute a minimum set C_a^b dominating $\mathcal{D}_a^b - N_a^b$ such that $C_a^b \cap Z$ is connected for each connected component Z of $\mathcal{D}_a^b$ in $\mathrm{O}(k^3 n\,(ed)^{72(k+5)/\pi} 2^{24(k+5)/\pi})$ time.*

Proof: We will apply the algorithm described in Theorem 5.4.9. Similar to Lemma 6.3.16, one needs to bound the maximum number of disks of a minimum connected dominating set appearing in three consecutive slabs, by considering the slabs to the left and to the right of these three slabs. Following Lemma 6.3.3, Lemma 6.3.16, and Proposition 6.3.24, a dominating set C of cardinality $3 \cdot 24(k+5)/\pi$ exists for these five slabs such that $C \cap Z$ is connected for each connected component Z of these slabs. As three consecutive slabs contain at most $3(k+4)d$ disks, Lemma 6.3.26 and Lemma 6.3.2 show that one needs to consider at most $\mathrm{O}(k\,(ed)^{72(k+5)/\pi})$ different subsets.

Observe furthermore that the number of connected components of a subset of the disks in a single slab is bounded by the maximum cardinality of an independent set. Using Lemma 6.3.3, this number is at most $8(k+5)/\pi$. The lemma now follows from Theorem 5.4.9. □

Applying Lemma 6.3.23, we can show that for any $0 \leq a \leq k-1$, $C_a = \bigcup_{b \in \mathbb{Z}} C_a^b$ is a connected dominating set, since $\{\mathcal{D}_a^b, \mathcal{D} - (\mathcal{D}_a^b - S_a^b)\}$ is a quadruple separation. Let $C_{\min}$ be a smallest such set.

Lemma 6.3.28 $|C_{\min}| \leq (1 + 8/k) \cdot |\mathcal{C}|$*, where $\mathcal{C}$ is a minimum connected dominating set of G.*

Proof: It follows from Lemma 6.3.25 that for any $0 \leq a \leq k-1$ and any $b \in \mathbb{Z}$, $|C_a^b| \leq |\mathcal{C} \cap \mathcal{D}_a^b| + 2|\mathcal{C} \cap N_a^b|$. As $C_a = \bigcup_{b \in \mathbb{Z}} C_a^b$,

$$|C_a| \;\leq\; \sum_{b \in \mathbb{Z}} \left(|\mathcal{C} \cap \mathcal{D}_a^b| + 2|\mathcal{C} \cap N_a^b|\right) \;\leq\; |\mathcal{C}| + \tfrac{1}{2} \sum_{b \in \mathbb{Z}} \left(|\mathcal{C} \cap S_a^b| + 4|\mathcal{C} \cap N_a^b|\right).$$

A disk is in $\bigcup_{b\in\mathbb{Z}} S_a^b$ for four values of a (and then for two values of b) and in $\bigcup_{b\in\mathbb{Z}} N_a^b$ for two values of a. Then

$$\begin{aligned} k\cdot|C_{\min}| \quad &\leq \quad \sum_{a=0}^{k-1}|C_a| \quad \leq \quad k|\mathcal{C}| + \tfrac{1}{2}\sum_{a=0}^{k-1}\sum_{b\in\mathbb{Z}}\left(|\mathcal{C}\cap S_a^b| + 4|\mathcal{C}\cap N_a^b|\right) \\ &\leq \quad k|\mathcal{C}| + 4|\mathcal{C}| + 4|\mathcal{C}| \end{aligned}$$

and thus $|C_{\min}| \leq (1+8/k)\cdot|\mathcal{C}|$. $\square$

Lemma 6.3.29 *For any $k \geq 5$, one can obtain a $(1+8/k)$-approximation for Minimum Connected Dominating Set on n-vertex unit disk graphs G with a known representation $\mathcal{D}$ of density d in $\mathrm{O}(k^4 n^2\,(ed)^{72(k+5)/\pi}2^{24(k+5)/\pi})$ time.*

We can now give an eptas in a manner similar as we did for the other problems in this chapter.

Theorem 6.3.30 *There is an eptas for Minimum Connected Dominating Set on unit disk graphs with n vertices and bounded density, i.e. density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$.*

Proof: Consider any number $\epsilon > 0$. Choose k as the largest integer such that $(72(k+5)/\pi)\cdot\log(ed) \leq \log n$. If $k < 5$, output $V(G)$. Otherwise, apply the algorithm of Lemma 6.3.29 and compute $C_{\min}$ in $\mathrm{O}(n^4\log^4 n)$ time. Furthermore, if $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$, there is a c_ϵ such that $k \geq 8/\epsilon$ and $k \geq 5$ for all $n \geq c_\epsilon$. Therefore, if $n \geq c_\epsilon$, it follows from Lemma 6.3.28 and the choice of k that $C_{\min}$ is a $(1+\epsilon)$-approximation of the optimum. Hence there is a fiptas$^\omega$ for Minimum Connected Dominating Set on n-vertex unit disk graphs of bounded density, i.e. of density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$. The theorem now follows from Theorem 2.2.4. $\square$

Recall that d is always bounded by n. Hence the worst-case running time of the algorithm described in Lemma 6.3.29 is $\mathrm{O}(k^4 n^2\,(en)^{72(k+5)/\pi}2^{24(k+5)/\pi})$.

Theorem 6.3.31 *There is a ptas for Minimum Connected Dominating Set on unit disk graphs.*

The ptas given here improves on the $n^{\mathrm{O}(\epsilon^{-2}\log^2\epsilon^{-1})}$-time ptas given by Cheng et al. [64] and the $n^{\mathrm{O}(\epsilon^{-2})}$-time ptas by Zhang et al. [278].

6.3.5 Generalizations

We considered Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set on unit disk graphs. For all of these problems, we obtained a ptas in general and an eptas if the density of the given representation is $\mathrm{O}(n^{\mathrm{o}(1)})$. For Minimum Vertex Cover, we even have an eptas on arbitrary unit disk graphs. These schemes extend to any constant dimension. It is easy to extend the notion of density to any finite dimension. Then consider boxes of infinite width and all other sides of length $\approx k$, a natural extension of strips. This yields approximation schemes with running time $\mathrm{O}(\mathrm{poly}(n, 1/\epsilon)\, d^{\mathrm{O}(\epsilon^{-l-1})})$ on unit ball graphs in $\mathbb{R}^l$ with density d.

Observe that extension to unit ball graphs in any dimension is not possible. Any n-vertex graph can be embedded as a constant-density unit ball graph in $(n-1)$-dimensional space [198, 147]. Hence Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set in $(n-1)$-dimensional space are as hard as in general.

Furthermore, we can extend the schemes to intersection graphs of other geometric objects than unit disks, for instance unit squares, unit triangles, etc., as long as the unit object is sufficiently 'disk-like'. In other words, the object should be *fat*. Many formal definitions of 'fat' exist, but as an example, it is easy to see that the algorithms extend to translated copies of any α-fat object. A convex subset s of $\mathbb{R}^l$ is *α-fat* for some $\alpha \geq 1$ if the ratio between the radii of the smallest sphere enclosing s and the largest sphere inscribed in s is at most α [97].

If we generalize these problems further, for instance when considering their weighted case, the worst-case analysis worsens. In the presence of (arbitrary) weights on the vertices of the graph, the presented schemes are extendable to an eptas for Minimum-Weight Vertex Cover and Maximum-Weight Independent Set if the density is bounded by $\mathrm{O}(n^{\mathrm{o}(1)})$. They are a ptas on general unit disk graphs and extend to fat objects and to any constant dimension. Unfortunately, the idea behind Lemma 6.3.14 that reduces the density of the unit disk graph for the minimum vertex cover problem does not seem to carry over to the weighted case, so the existence of an eptas in this case is open.

For Minimum Dominating Set on unit disk graphs, we used in the analysis that the cardinality of a maximum independent set yields a linear upper bound to the cardinality of a dominating set. This property is not transferable to the weighted case and hence we lose the upper bound implied by Lemma 6.3.2. Hence the scheme of Lemma 6.3.19 now has a worst-case running time of $\mathrm{O}(\mathrm{poly}(n, 1/\epsilon)\, 2^{\mathrm{O}(d/\epsilon)})$ on unit disk graphs of density d. Therefore we have (analogously to Theorem 6.3.20) an eptas for Minimum-Weight Dominating Set, but only if the density is $\mathrm{o}(\log n)$. Moreover, the scheme does not extend to a ptas on general unit disk graphs. Although a constant-factor approximation algorithm exists in this case [13, 153, 74], the existence of a ptas is open.

The connected dominating set problem on weighted unit disk graphs is even harder. We inherit the difficulties described above for Minimum-Weight Dominating Set, but now Lemma 6.3.25 also fails. At the moment, it is unclear whether a result similar to Lemma 6.3.25 applies to Minimum-Weight Connected Dominating Set. This would immediately yield an eptas for this problem on unit disk graphs of density $d = d(n) = \mathrm{o}(\log n)$.

Instead of considering unit objects, one can also extend the schemes to intersection graphs of objects of bounded size, e.g. unit disks of bounded radius. If the ratio between the smallest and the largest object is constant, the usual analysis holds and we obtain $O(\text{poly}(n, 1/\epsilon)\ d^{O(1/\epsilon)})$-time schemes, where the hidden constants depend on the radius ratio. This yields an eptas if the density is bounded by $O(n^{o(1)})$, and a ptas in the general case.

If only an upper bound to the size is known, one can no longer bound the size of a maximum independent set in a slab by a number independent of d as in Lemma 6.3.3. Hence we obtain $O(\text{poly}(n, 1/\epsilon)\, 2^{O(d/\epsilon)})$-time schemes, where the hidden constants depend on the maximum object size. This gives an eptas on intersection graphs of fat objects if the density is $o(\log n)$. These schemes extend to the weighted case, except for Minimum Connected Dominating Set.

Finally, we consider subgraphs of disk graphs of bounded radius. In other words, we consider geometric graphs of bounded edge length. If the subgraph actually is a ρ-quasi unit disk graph for some constant ρ, we again obtain $O(\text{poly}(n, 1/\epsilon)\ d^{O(1/\epsilon)})$-time approximation schemes, improving on the schemes implied by Nieberg, Hurink, and Kern [219]. If ρ cannot be bounded by a constant, then we obtain $O(\text{poly}(n, 1/\epsilon)\, 2^{O(d/\epsilon)})$-time schemes, where the hidden constants depend on the maximum edge length. Moreover, these schemes also apply to the weighted case (except for Minimum-Weight Connected Dominating Set). This gives an eptas on geometric graphs if the density is $o(\log n)$. This generalizes results of Hunt et al. [154], who showed that such schemes exist on civilized graphs, which are geometric graphs of bounded edge length and bounded precision (recall that density is a more general notion than precision).

The extension to disk graphs with disks of arbitrary ratio requires new techniques and is considered in Chapter 7 and 8.

6.4 Optimality

Beyond these generalizations, an important question is whether one can improve on the algorithms given in this chapter. We show that the schemes given here are optimal, up to constants. This result follows essentially from close inspection of work by Marx [204].

The optimality results given here are under the condition of the *exponential time hypothesis*, which states that n-variable 3SAT cannot be decided in $2^{o(n)}$ time. Using probabilistically checkable proof systems, one can show the following. An m-clause SAT formula is called *α-satisfiable* for some $0 \leq \alpha \leq 1$ if there is a truth setting such that at least αm clauses are satisfied.

Lemma 6.4.1 (Marx [204]) *There is a constant $0 < \alpha < 1$ such that if there is an algorithm that can distinguish in $2^{O(m)^{1-\beta}}$ time for some $\beta > 0$ whether an m-clause 3SAT formula is satisfiable or not α-satisfiable, then the exponential time hypothesis is false.*

One can show that 3SAT formulas are reducible to instances of Maximum Independent Set and Minimum Dominating Set on unit disk graphs.

Lemma 6.4.2 (Marx [204]) *Given an m-clause 3SAT formula φ and an integer k, there is an instance x of Maximum Independent Set on unit disk graphs of density $d = \mathrm{O}(3^{m/k})$ such that for every $0 < \alpha < 1$:*

- *if φ is satisfiable, then $m^*(x) = f(k)$*
- *if φ is not α-satisfiable, then $m^*(x) < f(k) - k(1-\alpha)/2 + 1$,*

where $f(k) = \Theta(k^2)$ is a polynomial. Moreover, this instance x can be computed in time polynomial in m, d, and k.

These two lemmas can be used to prove the following theorem.

Theorem 6.4.3 *If there exist constants $\delta \geq 1$, $0 < \beta < 1$ such that Maximum Independent Set on unit disk graphs of density d has a ptas with running time $2^{\mathrm{O}(1/\epsilon)^\delta} d^{\mathrm{O}(1/\epsilon)^{1-\beta}} n^{\mathrm{O}(1)}$, then the exponential time hypothesis is false.*

Proof: We show that if a ptas as in the theorem statement exists, then an algorithm as in the statement of Lemma 6.4.1 exists. Let φ be an m-clause 3SAT formula. Set $k = \lceil m^{1/(2\delta+1)} \rceil$ and apply Lemma 6.4.2 to obtain an instance x of Maximum Independent Set on unit disk graphs with density $\mathrm{O}(3^{m/k}) = \mathrm{O}(3^{m^{1-1/(2\delta+1)}})$. This takes $2^{\mathrm{O}(m)^{1-1/(2\delta+1)}}$ time.

Now let α be as in Lemma 6.4.1 and choose $\epsilon = \frac{k(1-\alpha)/2-1}{f(k)}$, with f as in Lemma 6.4.2. If φ is satisfiable, then $m^*(x) = f(k)$. If φ is not α-satisfiable, then $m^*(x) < f(k) - k(1-\alpha)/2 + 1 = (1-\epsilon) \cdot f(k)$. As the ptas gives a solution y for which $m(x,y) \geq (1-\epsilon) \cdot m^*(x)$, the choice of ϵ is sufficient to distinguish whether φ is satisfiable or not α-satisfiable. For this choice of ϵ, the ptas runs in time

$$2^{\mathrm{O}(1/\epsilon)^\delta} d^{\mathrm{O}(1/\epsilon)^{1-\beta}} n^{\mathrm{O}(1)} \quad = \quad 2^{\mathrm{O}(1/\epsilon)^\delta + \mathrm{O}(1/\epsilon)^{1-\beta} \log d + \mathrm{O}(\log mdk)}.$$

Since

$$\begin{aligned}
&\mathrm{O}(1/\epsilon)^\delta + \mathrm{O}(1/\epsilon)^{1-\beta} \log d + \mathrm{O}(\log d) + \mathrm{O}(\log mk) \\
&\quad = \mathrm{O}(f(k)/k)^\delta + \mathrm{O}(f(k)/k)^{1-\beta} \cdot \mathrm{O}(m/k) + \mathrm{O}(m/k) + \mathrm{O}(\log m) \\
&\quad = \mathrm{O}(k)^\delta + \mathrm{O}(k)^{1-\beta} \cdot \mathrm{O}(m)^{1-1/(2\delta+1)} \\
&\quad = \mathrm{O}(m)^{\delta/(2\delta+1)} + \mathrm{O}(m)^{(1-\beta)/(2\delta+1)} \cdot \mathrm{O}(m)^{2\delta/(2\delta+1)} \\
&\quad = \mathrm{O}(m)^{1-(\delta+1)/(2\delta+1)} + \mathrm{O}(m)^{((1-\beta)+2\delta)/(2\delta+1)} \\
&\quad = \mathrm{O}(m)^{1-(\delta+1)/(2\delta+1)} + \mathrm{O}(m)^{1-\beta/(2\delta+1)} \\
&\quad = \mathrm{O}(m)^{1-\beta/(2\delta+1)}
\end{aligned}$$

But then one can distinguish whether φ is satisfiable or not α-satisfiable in $2^{\mathrm{O}(m)^{1-\beta/(2\delta+1)}}$ time. As $0 < \beta/(2\delta+1) < 1$, according to Lemma 6.4.1, this implies that the exponential time hypothesis is false. $\square$

Marx [204] gives a reduction similar to Lemma 6.4.2 for Minimum Dominating Set. This implies the following result.

Theorem 6.4.4 *If there exist constants $\delta \geq 1$, $0 < \beta < 1$ such that Minimum Dominating Set on unit disk graphs of density d has a ptas with running time $2^{O(1/\epsilon)^{\delta}} d^{O(1/\epsilon)^{1-\beta}} n^{O(1)}$, then the exponential time hypothesis is false.*

Therefore the approximation schemes for Maximum Independent Set and Minimum Dominating Set given in Lemma 6.3.7 and Lemma 6.3.19 are optimal, up to constants, unless the exponential time hypothesis is false.

The scheme we presented for Minimum Vertex Cover is also optimal, but a slightly different idea is needed to prove it. In this case, we start out from 2SAT formulas.

Lemma 6.4.5 (Marx [204]) *There are constants $0 < \alpha_2 < \alpha_1 < 1$ such that if there is an algorithm that can distinguish in $2^{O(m)^{1-\beta}}$ time for some constant $\beta > 0$ whether an m-clause 2SAT formula is α_1-satisfiable or not α_2-satisfiable, then the exponential time hypothesis is false.*

Actually, one may assume that the 2SAT formula is simple (meaning that it contains no duplicate clauses and each clause is satisfiable) and a variable appears in a constant number of clauses. We call this a *basic* 2SAT formula.

We now reduce from 2SAT formulas to instances of Minimum Vertex Cover on unit disk graphs.

Lemma 6.4.6 (Marx [204]) *There is a constant d_0 such that given an m-clause basic 2SAT formula φ there is an instance x of Minimum Vertex Cover on unit disk graphs of density at most d_0, such that for every $0 < \alpha_2 < \alpha_1 < 1$:*

- *if φ is α_1-satisfiable, then $m^*(x) \geq f(k) + (1 - \alpha_1)m$,*
- *if φ is not α_2-satisfiable, then $m^*(x) < f(k) + (1 - \alpha_2)m$,*

where $k = \Theta(m)$ and $f(k) = \Theta(k^2)$. Moreover, this instance x can be computed in time polynomial in m.

Theorem 6.4.7 *If there is a constant $0 < \beta < 1$ such that Minimum Vertex Cover on unit disk graphs of density at most d_0 has an eptas with running time $2^{O(1/\epsilon)^{1-\beta}} n^{O(1)}$, then the exponential time hypothesis is false.*

Proof: Suppose that an eptas as in the theorem statement does exist. Let φ be a basic m-clause 2SAT formula and use Lemma 6.4.6 to construct an instance x of Minimum Vertex Cover on unit disk graphs of density at most d_0. If we set $\epsilon = (\alpha_1 - \alpha_2)m/f(k)$ with k and f as in Lemma 6.4.6, then

$$
\begin{aligned}
&(1+\epsilon) \cdot (f(k) - (1-\alpha_2)m) \\
&\qquad \leq \Big(f(k) - (1-\alpha_2)m\Big) + \Big(f(k) - (1-\alpha_2)m\Big) \cdot \frac{(\alpha_1 - \alpha_2)m}{f(k) - (1-\alpha_2)m} \\
&\qquad = f(k) - (1-\alpha_1)m
\end{aligned}
$$

and thus the eptas applied to x and ϵ can distinguish whether φ is α_1-satisfiable or not α_2-satisfiable. The running time is

$$2^{\mathrm{O}(1/\epsilon)^{1-\beta}} n^{\mathrm{O}(1)} \quad = \quad 2^{\mathrm{O}(k)^{1-\beta}} k^{\mathrm{O}(1)} \quad = \quad 2^{\mathrm{O}(m)^{1-\beta}}.$$

Following Lemma 6.4.5, the exponential time hypothesis is false. □

Observe that Lemma 6.3.13 yields an $2^{\mathrm{O}(1/\epsilon)} n^{\mathrm{O}(1)}$-time eptas for Minimum Vertex Cover on unit disk graphs of any constant density. But then the algorithm given in Lemma 6.3.13 is optimal, up to constants, unless the exponential time hypothesis is false.

The exponential time hypothesis is not very frequently used in proving hardness of approximation results. If we settle for slightly worse results, we can use more familiar complexity conditions. Marx [202, 203] showed that Maximum Independent Set and Minimum Dominating Set are W[1]-hard on unit disk graphs. As the standard parameterization of any problem permitting an eptas must be in FPT [26, 53], one has the following result.

Theorem 6.4.8 (Marx [202, 203]) *Maximum Independent Set and Minimum Dominating Set on unit disk graphs have no eptas, unless FPT=W[1].*

The constructions used in the W[1]-hardness proofs have high density. Using a small trick, we can obtain the following strengthening of Theorem 6.4.8.

Theorem 6.4.9 *Maximum Independent Set and Minimum Dominating Set on n-vertex unit disk graphs of density $d = d(n) = \Omega(n^{\alpha})$ for some constant $0 < \alpha \leq 1$ cannot have an eptas, unless FPT=W[1].*

Proof: Suppose that such an eptas does exist. Then consider any set $\mathcal{D}$ of unit disks, which clearly has density at most n. Take $\lceil n^{(1-\alpha)/\alpha} \rceil$ disjoint copies of $\mathcal{D}$ and let $\mathcal{D}'$ denote the resulting set of disks. The density of $\mathcal{D}'$ is at most $n = (n \cdot n^{(1-\alpha)/\alpha})^{\alpha} \leq |\mathcal{D}'|^{\alpha}$. Now run the eptas on $\mathcal{D}'$ and return the best solution over all copies of $\mathcal{D}$. This construction gives an eptas on arbitrary unit disk graphs, which is impossible by Theorem 6.4.8. □

The bound of Theorem 6.4.9 is a precise match with Theorem 6.3.8 and Theorem 6.3.20, where we showed that Maximum Independent Set and Minimum Dominating Set have an eptas on n-vertex unit disk graphs of density $d = d(n) = \mathrm{O}(n^{o(1)})$. Hence no better approximation scheme is possible then given in these theorems, unless FPT=W[1].

Note that this last result does not imply anything about the actual running time of the schemes, and hence it is slightly weaker than Theorem 6.4.3 and Theorem 6.4.4. Also, we know that Minimum Vertex Cover does have an eptas and hence we cannot say much about it using FPT versus W[1]. In fact, using classic complexity theory, we can (at the moment) only say that Minimum Vertex Cover does not have an fptas on unit disk graphs, unless P=NP.

Similarly, it is hard to give an optimality result for Minimum Connected Dominating Set. We conjecture that Lemma 6.3.29 is optimal and that Minimum Connected Dominating Set on unit disk graphs has no eptas. To prove this, one could give an L-reduction from Maximum Independent Set or Minimum Dominating Set to Minimum Connected Dominating Set on unit disk graphs to apply Lemma 6.4.1 in the manner of Theorem 6.4.3 and Theorem 6.4.4, or extend the reduction of Theorem 6.4.8.

Finally, we note that since the required auxiliary results of Marx carry over to unit square graphs [202, 203, 204], all results of this section also apply.

6.5 Connected Dominating Set on Graphs Excluding a Minor

Although this chapter has solely focused on unit disk graphs, we end the chapter with a small aside on minor-closed classes of graphs excluding a fixed minor. Note that such graphs are not a generalization of unit disk graphs, as unit disk graphs are not minor closed (see Section 3.2).

To be precise, we will study minor-closed classes of graphs excluding an apex graph as a minor, in short, *apex-minor-free* graph classes. An *apex graph* is a graph H that possesses a vertex v (the *apex*) such that $H-v$ is planar. Examples of apex-minor-free graph classes are planar graphs, graphs of bounded genus, and single-crossing-minor-free graphs [99, 85].

We direct the attention to Minimum Connected Dominating Set on apex-minor-free graph classes. Demaine and Hajiaghayi [84] proved that Minimum Connected Dominating Set has an $n^{O((1/\epsilon)\log(1/\epsilon)\log\log n)}$-time approximation scheme (an *almost-ptas*) on apex-minor-free graphs, a ptas on single-crossing-minor-free graphs, and an eptas on planar graphs. The last two schemes follow from a generic approach to approximating so-called bidimensional problems. The first scheme is based on a generalization of Baker's shifting technique [22] in the way proposed by Eppstein [99] and Grohe [131]. We show that this generalization can actually be used in a way that improves on all these schemes.

The goal of this section is to prove the following theorem.

Theorem 6.5.1 *There is a $2^{O_H(1/\epsilon)}\, n^{O(1)}$-time eptas for Minimum Connected Dominating Set on the minor-closed class of graphs excluding some fixed apex graph H as a minor.*

Here $O_H(1/\epsilon)$ means that the hidden constant depends on (the number of vertices of) the excluded graph H.

The scheme employs the shifting technique as proposed by Eppstein [99] and Grohe [131] and relies heavily on the ideas developed in Section 6.3.4.

The basic idea is the following. Let G be a connected graph from a minor-closed class of graphs that excludes some fixed apex graph H. Fix a vertex $v_0 \in V(G)$ and consider the layers of G with respect to v_0. We say that $u \in L_i$ (i.e. u is in layer i) if the shortest v_0–u path has length i. Clearly, every vertex of G is in some layer and there are at most $n = |V(G)|$ layers. One may group layers together, such that $L_{i,j} = \bigcup_{h=i}^{j} L_h$. For simplicity, we assume that $L_i = \emptyset$ for all $i < 0$.

We first show how the shifting technique can be applied to obtain a $(1+\epsilon)$-approximation. Let $k \geq 5$ be an integer. For each integer $0 \leq a \leq k-1$ and $b \in \mathbb{Z}$, let $L_a^b = L_{bk+a,(b+1)k+a+3}$. Let $S_a^b = (L_a^{b-1} \cap L_a^b) \cup (L_a^b \cap L_a^{b+1})$ and $N_a^b = L_{bk+a} \cup L_{(b+1)k+a+3}$. Suppose that we can compute for each $0 \leq a \leq k-1$ and $b \in \mathbb{Z}$ a minimum set $C_a^b \subseteq L_a^b$ dominating $L_a^b - N_a^b$ such that $C_a^b \cap Z$ is connected for each connected component Z of L_a^b. Observe that $\{L_a^b, V(G) - (L_a^b - S_a^b)\}$ is a quadruple separation of G. Hence, according to Lemma 6.3.23, $C_a = \bigcup_{b \in \mathbb{Z}} C_a^b$ is a connected dominating set of G. Let $C_{\min}$ be a set C_a of minimum cardinality. Then Lemma 6.3.28 proves that $|C_{\min}| \leq (1 + 8/k) \cdot |\mathcal{C}|$, where $\mathcal{C}$ is a minimum connected dominating set.

It now remains to show that the sets C_a^b can be computed in $2^{O_H(k)}\, n^{O(1)}$ time. We prove this in two steps. Step one is to bound the treewidth of L_a^b.

Lemma 6.5.2 $\operatorname{tw}(L_a^b) = O_H(k)$ *for any* $0 \leq a \leq k-1$ *and any* $b \in \mathbb{Z}$.

Proof: We may assume that $L_a^b \neq \emptyset$. Eppstein [99] showed that a minor-closed class of graphs does not contain all apex graphs if and only if every graph in this class that has diameter D has treewidth at most $f(D)$ for some function f. In fact, Demaine and Hajiaghayi [85] strengthened Eppstein's result and proved that this function f is always linear in D. As G is from a minor-closed class of graphs excluding the apex graph H as a minor, it follows that any minor of G of diameter D has treewidth $O_H(D)$.

Now consider $L_a^b = L_{bk+a,(b+1)k+a+3}$. Let $\hat{L}_a^b$ be the minor of G obtained by contracting layers $L_0, \ldots, L_{bk+a-1}$ into v_0 and removing layers L_i for all $i > (b+1)k + a + 3$. By the construction of the layers, $\hat{L}_a^b$ has diameter at most $k+5$. From the above results, this implies that $\operatorname{tw}(\hat{L}_a^b) = O_H(k)$. As L_a^b is a minor of $\hat{L}_a^b$, $\operatorname{tw}(L_a^b) \leq \operatorname{tw}(\hat{L}_a^b)$ and thus $\operatorname{tw}(L_a^b) = O_H(k)$. $\square$

Step two is to use this bound in an algorithm for Minimum Connected Dominating Set.

Lemma 6.5.3 *For any* $0 \leq a \leq k-1$ *and* $b \in \mathbb{Z}$, C_a^b *can be computed in* $2^{O_H(k)}\, n^{O(1)}$ *time.*

Proof: Dorn, Fomin, and Thilikos [91] present an algorithm that given an integer w and a graph G from some H-minor-free graph class either certifies that $\operatorname{bw}(G) \geq w$ or returns a branch decomposition with Catalan structure of width $O_H(w)$. Moreover, the algorithm runs in $O_H(1)\, n^{O(1)}$ time. By a branch decomposition with Catalan structure, we mean that for any middle set M the number of equivalence classes induced by $M \cap \mathcal{Z}$ is at most $2^{O_H(|M|)}$, where $\mathcal{Z}$ ranges over all families of connected subgraphs of G. (Actually, the result by Dorn, Fomin, and Thilikos is stronger, but the above will be sufficient here.) We are now in a similar case as in Theorem 5.4.9, but on branch decompositions. Adapting the algorithm of Theorem 5.4.9 or using the faster algorithms by Dorn [89], one can solve Minimum Connected Dominating Set in $2^{O_H(w)}\, n^{\mathrm{O}(1)}$ time.

Because Lemma 6.5.2 shows that $\mathrm{bw}(L_a^b) \leq \mathrm{tw}(L_a^b) + 1 = O_H(k)$, one can find a branch decomposition with Catalan structure of width $O_H(k)$ in $O_H(1)\, n^{\mathrm{O}(1)}$ time. Then we use a slight variation of the above algorithm to compute C_a^b using this branch decomposition. Hence C_a^b can be computed in $2^{O_H(k)}\, n^{\mathrm{O}(1)}$ time. $\square$

Proof of Theorem 6.5.1: Given any $0 < \epsilon < 1$, choose $k = \max\{5, \lceil 8/\epsilon \rceil\}$. Compute $C_{\min}$ in $2^{O_H(1/\epsilon)}\, n^{\mathrm{O}(1)}$ time. Then $C_{\min}$ is a $(1+\epsilon)$-approximation of the optimum. $\square$

We conjecture that Theorem 6.5.1 can be extended to a ptas or even an eptas on H-minor-free graphs for arbitrary graphs H by using the techniques developed by Grohe [131]. We leave this to future research.

Baker [22] showed, using similar ideas as presented above, that Maximum Independent Set, Minimum Vertex Cover, and Minimum Dominating Set have an $2^{\mathrm{O}(1/\epsilon)}\, n^{\mathrm{O}(1)}$-time eptas on planar graphs. These were extended by Eppstein [99] to apex-minor-free graphs. Recently, Marx [204] showed that under the exponential time hypothesis, these schemes are essentially optimal, meaning that they have no $2^{\mathrm{O}(1/\epsilon)^{1-\beta}}\, n^{\mathrm{O}(1)}$-time eptas for any $\beta > 0$. We conjecture that the above scheme is also optimal under the exponential time hypothesis. A proof direction would be to find an L-reduction from Maximum Independent Set, Minimum Vertex Cover, or Minimum Dominating Set on planar graphs to Minimum Connected Dominating Set on planar graphs. Combined with the results of Marx [204], this would prove the optimality of Theorem 6.5.1.