Optimization and approximation on systems of geometric objects

van Leeuwen, E.J.

Publication date
2009

Document Version
Final published version

Optimization and Approximation

on

Systems of Geometric Objects

Erik Jan van Leeuwen

This research has been carried out at the Centrum Wiskunde & Informatica in Amsterdam.

Part of this research has been funded by the Dutch BSIK/BRICKS project.

Printed and bound by Ipskamp Drukkers B.V., The Netherlands.

ISBN 978-90-9024317-7

Optimization and Approximation

on

Systems of Geometric Objects

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. D.C. van den Boom
ten overstaan van een door het college voor promoties
ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op dinsdag 16 juni 2009, te 10:00 uur

door

Erik Jan van Leeuwen

geboren te Utrecht

Promotiecommissie:

Promotor:
prof. dr. A. Schrijver

Overige leden:
prof. dr. M.T. de Berg
prof. dr. H.M. Buhrman
prof. dr. T. Erlebach
prof. dr. G.J. Woeginger
dr. L. Torenvliet

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Do not disturb my circles!

– last words of Archimedes
(287 BC – 212 BC)

Contents

Chapter 1

Introduction

Can a combinatorial optimization problem be approximated better if it is determined by a system of geometric objects? Many combinatorial optimization problems are NP-hard to solve and frequently even hard to approximate. On systems of geometric objects however, these same problems can usually be solved efficiently or approximated well.

The computational complexity of such geometric optimization problems highly depends on the underlying system of objects. We consider objects that arise naturally in diverse applications. For instance, the broadcasting range of a wireless device can be seen as a disk, an atom is just a sphere, and a label on a map is a rectangle. In this thesis, we consider hard combinatorial optimization problems on systems of geometric objects, motivated by applications in for example wireless networks, computational biology, and map labeling.

1.1 Optimization Problems and Systems of Geometric Objects

Given a combinatorial optimization problem and its objective function, one preferably is able to compute the optimum objective value and a corresponding solution in an efficient way. Unfortunately, depending on one's notion of efficient, this is not always possible. The measure we use to decide whether an algorithm is efficient is polynomial running time. That is, whether the algorithm returns the optimum objective value (and possibly a corresponding solution) using a number of 'basic steps' that is polynomial in the size of the problem instance. An optimization problem having such an algorithm is called *polynomial-time solvable*.

Many optimization problems are solvable in polynomial time. However, there are also many optimization problems that are not known to be solvable in polynomial time. In fact, there is very strong belief that these problems cannot be solved in polynomial time. To support this idea, a class of hard optimization problems has been identified, namely the class of *NP-hard* optimization problems.

The most sensible way to work around NP-hardness while still restricting to polynomial-time algorithms is to try and find approximate solutions to such optimization problems. An *approximation algorithm* returns an objective value (and possibly a corresponding solution) within a certain (additive or multiplicative) factor of the optimum objective value. Approximation algorithms have been developed for many different optimization problems and are widely used in practice. We should note however that sometimes one can even show that approximating a problem within a certain factor is NP-hard.

In this field of hard optimization problems, we try to ascertain the influence on the computational complexity of these problems when using a system of geometric objects as the underlying combinatorial structure. Since many optimization problems use graphs as their underlying structure, it might be no surprise that one of the central notions of this thesis is a graph induced by a system of geometric objects: a *geometric intersection graph*.

Given a system of geometric objects, the vertices of the induced graph correspond to the given objects and there is an edge between two vertices if the corresponding objects intersect. There are many famous examples of geometric intersection graphs, for instance the classic *interval graph* is an intersection graph of intervals on the real line. For this thesis however, we are concerned with intersection graphs of objects in two- and higher-dimensional space, such as *(unit) disk graphs*, which are intersection graphs of (equally-sized) disks.

Most classical optimization problems are still NP-hard on (intersection graphs of) systems of two-dimensional geometric objects [17, 67, 110, 157, 267]. However, there is considerable evidence (e.g. [4, 68, 74, 103, 150, 154, 201, 279]) to support the claim that *approximating* these problems is easier, meaning that better approximation algorithms exist, on systems of geometric objects than in general. The goal of this thesis is to further investigate this claim.

> **In this thesis we study the approximability of hard combinatorial optimization problems on (intersection graphs of) systems of geometric objects.**

As part of this study, we give both approximation algorithms and inapproximability results. In the process, we show that the approximability highly depends on the shape and size of the geometric objects under consideration.

1.2 Application Areas

It turns out that optimization problems on systems of geometric objects, and particularly on geometric intersection graphs, occur in many application areas. We consider some of the foremost application areas and associated problems and systems of objects below.

1.2.1 Wireless Networks

It is very common for wireless networks to be modeled as geometric intersection graphs. In 1961 already, Gilbert [122] used the following idea to model wireless networks. Each wireless device is assumed to be a point in the plane and an edge is drawn if the distance between two points is at most some constant δ.

In other words, the network is modeled by the intersection graph of a set of disks of radius $\delta/2$. This graph would later be called a *unit disk graph* [136].

By now, unit disk graphs have easily become the prevalent model for wireless networks (although more advanced models exist, see Section 3.2.1). The model implicitly assumes that all devices in the network have the same broadcasting range. This is a perfect fit to an increasingly popular network type, called a *(mobile) ad hoc network* [82, 166, 225, 232, 252]. A mobile ad hoc network is an autonomous collection of mobile devices that communicate over wireless channels. The network is self-organizing and self-reliant. In contrast to classic wireless networks, no fixed infrastructure in the form of base stations is present and messages are routed from source to destination through multiple hops. Mobile ad hoc networks already exist since the late 1970s [1]. With the ongoing miniaturization of chips and wireless transceivers, these networks attained renewed interest in the last few years. This led for instance to the advent of *wireless sensor networks* [5, 121, 133, 274].

Several well-known optimization problems are relevant to (ad hoc) wireless networks, and have consequently been studied on unit disk graphs and its generalizations. A *maximum independent set* in the graph corresponds to a largest set of devices that can transmit simultaneously without causing interference. The *minimum dominating set* problem has been studied to find a collection of emergency transmitters, capable of reaching every device in the network. Alternatively, it can be used to construct a routing backbone. This is also where a *minimum connected dominating set* comes in.

One of the first studies into unit disk graphs was on its *chromatic number* [136]. The minimum number of colors needed is equal to the minimum number of channels needed in the network to communicate without causing interference. Hence solving the coloring problem actually solves the frequency assignment problem.

1.2.2 Wireless Network Planning

Classical wireless networks consist of a number of powerful base stations that provide wireless service to smaller devices. Well-known examples are cellular networks (i.e. GSM networks) and wi-fi networks (i.e. 802.11 networks).

Because these networks rely on base stations, it is imperative to position them carefully. The common model (see e.g. [124]) is to view base stations as disks, where the radius of each disk corresponds to the range of the wireless signal broadcast by the base station, and the smaller wireless devices as points. The question then is how to place as few disks (base stations) as possible, but still cover all points, i.e. provide wireless service to all devices. This is the geometric version of the well-known *minimum set cover* problem. Several variations of this geometric set cover problem exist, depending for instance on whether we are free to choose the location of the disk or should choose from a given set of potential locations, or on how strict one is in insisting that all points are covered. See Part III of this thesis for more variants.

1.2.3 Computational Biology

Optimization problems on systems of geometric objects occur in several biological applications. Armitage [15] described in 1949 the need to find the number and the size of clumps of particles in order to properly count the total number of particles under a microscope. This clearly corresponds to problems on cliques in geometric intersection graphs. Armitage considered several models for the particles, including (unit) disks and rectangles of bounded aspect-ratio.

Kaufmann et al. [160] consider the alignment of DNA-sequences, which can be described by a maximum clique problem on max-tolerance graphs, a generalization of interval graphs. Interestingly, max-tolerance graphs are equivalent to intersection graphs of isosceles right triangles. Xu and Berger [272] use a geometric model to study the problem of attaching or assigning side-chains of a protein to an existing backbone while maximizing system energy.

1.2.4 Map Labeling

In general, map labeling is the problem of placing labels on a map, such that the labeling satisfies certain properties. For example, a label should be close to its corresponding item on the map and the texts of different labels should not overlap. Commonly, the labels are seen as rectangles, the items as points, and the boundary of the rectangle of a certain label should overlap the point modeling the corresponding item. Then one aims to maximize the number of labels that can be placed, without any overlaps among the rectangles. If the positions in which the rectangle may be placed are discretized, this is just the maximum independent set problem on an intersection graph of rectangles [4]. The continuous case is more complex [102].

1.2.5 Further Applications

One of the most commonly cited algorithmic results for optimization problems on systems of geometric objects applies to (among others) a problem in VLSI [150]. In the *geometric packing* problem, one wants to pack the largest number of objects of a certain prescribed shape into a larger object. This corresponds to maximizing yield when cutting chips from a large chip wafer.

A slightly morbid application is bombing. Garwood [117] described in 1947 the problem of minimizing the number of bombs needed to destroy points of interest in a certain area. Assuming that the bombs have a circular area of destruction upon impact, one wants to minimize the number of disks needed to cover certain other geometric objects, e.g. rectangles or points, representing buildings or matériel.

1.3 Thesis Overview

The thesis is comprised of three parts and eleven chapters. Below is an overview of their contents.

Part I: Foundations

We introduce and expand on the basic notions needed to understand the other parts of the thesis.

The topic of this thesis is made from two main ingredients. The first is approximation algorithms for optimization problems. Chapter 2 formally defines the type of optimization problems we consider here and what an approximation algorithm is. We define several classes of optimization problems, each admitting a particular type of approximation algorithm. The algorithms given in this thesis gave rise to new classes of optimization problems. We consider the relation of these classes to classic problem classes. The results of this chapter were obtained in joint work with J. van Leeuwen [263].

The second ingredient of this thesis is systems of geometric objects and geometric intersection graphs. We survey the main results on structural aspects of geometric intersection graphs in Chapter 3. We also present a new way to look at a fundamental question surrounding geometric intersection graphs: do geometric intersection graphs have a representation of polynomial size? In Chapter 4, we prove that this is equivalent to the question whether or not geometric intersection graphs have a representation that is polynomially separated. Chapter 4 is an extended version of joint work with J. van Leeuwen [262].

Part II: Approximating Optimization Problems on Geometric Intersection Graphs

We describe exact algorithms and approximation schemes for optimization problems on geometric intersection graphs. In particular, we consider Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set. As most of these problems are motivated by wireless networks, the majority of the graphs studied in this part are disk graphs.

We start by considering unit disk graphs in Chapter 5. If a unit disk graph has a special property called bounded thickness, then many optimization problems (such as Maximum Independent Set and Minimum Connected Dominating Set) can be solved exactly in polynomial time. To this end, we define a new graph decomposition, called a relaxed tree decomposition, which might be interesting on its own. We provide exact algorithms on such decompositions that are applicable to general graphs. If the graph is a unit disk graph, the running time of these algorithms can be expressed in terms of the thickness. Moreover, the bound on the worst-case running time of the algorithm for Minimum Connected Dominating Set is significantly lower on unit disk graphs of bounded thickness than might be expected from the bound in the general case.

In Chapter 6, we use the algorithms of Chapter 5 to give new, better approximation schemes for these optimization problems on general unit disk graphs by using the so-called shifting technique. The schemes are a ptas on general unit disk graphs and an eptas if the density of the set of disks is bounded. They improve on (the running time of) previous approximation schemes. For Minimum Vertex Cover, we give an improved eptas on arbitrary unit disk graphs. We generalize to intersection graphs of unit fat objects in constant dimension and the weighted case. Furthermore, we prove that the algorithms are optimal (up to constants), unless the exponential time hypothesis fails. The ideas we give for the minimum connected dominating set problem are used to give the first eptas for this problem on apex-minor-free graphs. Chapter 5 and Chapter 6 are based on a revised and extended version of [259].

In Chapter 7 we generalize these ideas to general disk graphs and intersection graphs of fat objects. The crucial idea is to consider systems of fat objects that have bounded ply. We subsequently present new, improved approximation schemes for Minimum Vertex Cover and Maximum Independent Set using the multi-level shifting technique. The approximation scheme for Minimum Vertex Cover is the first eptas for this problem on disk graphs. Most of the results of this chapter were presented in [260].

It seems difficult to generalize the results of Chapter 6 for Minimum (Connected) Dominating Set in the same way. Chapter 8 is devoted to the approximability of these problems on intersection graphs of systems of geometric objects of different sizes. The shifting technique yields a constant-factor approximation algorithm on intersection graphs of any set of fat objects, if the ply is bounded. The foremost innovation however is a general theorem to approximate Minimum Dominating Set on intersection graphs. We apply this to obtain the first constant-factor approximation algorithm for Minimum Dominating Set on intersection graphs of homothetic convex polygons and several other object types. We prove however that these methods cannot extend to intersection graphs of fat objects of arbitrary ply, of convex polygons, or of homothetic polygons by giving a strong approximation hardness result. We also prove APX-hardness on intersection graphs of arbitrary rectangles. This chapter contains an extended version of joint work with T. Erlebach [105].

Part III: Approximating Geometric Coverage Problems

We describe algorithms for the geometric version of (variants of) the minimum set cover problem. That is, the input consists of a set of geometric objects and a set of points, and we are asked to find a subset of the objects covering all points. In Chapter 9 we give the first polynomial-time approximation scheme for this problem on unit squares. We extend this to a ptas for Geometric Budgeted Maximum Coverage on unit squares. Moreover, we show that Geometric Set Cover is (very) hard to approximate on systems of fat objects and APX-hard on several other systems of two-dimensional objects. The chapter is based on yet unpublished joint work with T. Erlebach.

When we no longer insist that all points are covered, we arrive at several interesting new problems, which are variants of the unique coverage problem. In this problem, we are asked to cover a maximum number of points uniquely. Chapter 10 gives the first constant-factor approximation algorithms for this problem on unit disks and on unit squares by combining the shifting technique with complex dynamic programming algorithms. The multi-level shifting technique then generalizes these results to fat objects of bounded ply.

We again present hardness results to prove that the restriction to bounded ply is necessary. We also consider the approximability of the geometric version of a generalization of Minimum Set Cover, called Minimum Membership Set Cover. Chapter 10 is based on joint work with T. Erlebach [104].

Finally, Chapter 11 presents the conclusion and an outlook to future work.

1.3.1 Published Papers

This thesis is partially based on the following four (refereed) papers.

[1] van Leeuwen, E.J., "Approximation Algorithms for Unit Disk Graphs" in Kratsch, D. (ed.) *Graph-Theoretic Concepts in Computer Science, 31st International Workshop, WG 2005, Metz, France, June 23-25, 2005, Revised Selected Papers*, Lecture Notes in Computer Science 3787, Springer-Verlag, Berlin, 2005, pp. 351–361.

[2] van Leeuwen, E.J., "Better Approximation Schemes for Disk Graphs" in Arge, L., Freivalds, R. (eds.) *Algorithm Theory - SWAT 2006, 10th Scandinavian Workshop on Algorithm Theory, Riga, Latvia, July 6-8, 2006, Proceedings*, Lecture Notes in Computer Science 4059, Springer-Verlag, Berlin, 2006, pp. 316–327.

[3] Erlebach, T., van Leeuwen, E.J., "Approximating Geometric Coverage Problems" in Teng, S.H. (ed.) *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, Association for Computing Machinery, 2008, pp. 1267–1276.

[4] Erlebach, T., van Leeuwen, E.J., "Domination in Geometric Intersection Graphs" in Laber, E.S., Bornstein, C.F., Nogueira, L.T., Faria, L. (eds.) *LATIN 2008: Theoretical Informatics, 8th Latin American Symposium, Búzios, Brazil, April 7-11, 2008, Proceedings*, Lecture Notes in Computer Science 4957, Springer-Verlag, Berlin, 2008, pp. 747–758.

Part I

Foundations

Chapter 2

Primer on Optimization and Approximation

A key ingredient of this thesis is approximation algorithms for certain optimization problems. Hence it is useful to formally define what an optimization problem is and what types of approximation algorithms we distinguish.

In addition to defining the classic types of approximation schemes, we propose a new type of approximation scheme, the asymptotic approximation scheme. Many of the approximation algorithms that we will encounter later are essentially an asymptotic approximation scheme. We prove however that two classes of optimization problems having such a scheme, namely FPTAS$^\omega$ and FIPTAS$^\omega$, both coincide with the well-known class EPTAS. Hence instead of placing problems in a class of problems with an asymptotic approximation scheme, we can place these problems in the more familiar EPTAS class as well.

2.1 Classic Notions

To make formal statements about (equivalences among) classes of approximation schemes, we have to be precise about the machine model that we use, the type of problems that are considered, and the definitions of the studied classes. Throughout, we assume the random access machine model with logarithmic costs and representations in bits. This machine model is polynomially equivalent to the classic Turing machine and thus defines equivalent complexity classes up to polynomial factors. Furthermore, all numbers in this chapter are assumed to be rationals, unless otherwise specified.

Using this model, we study optimization problems following the definitions as can be found for example in Ausiello et al. [19].

Definition 2.1.1 *An* optimization problem P *is characterized by four properties:*

- *a set of* instances *(bitstrings)* I_P*;*
- *a function* S_P *that maps instances of* P *to (nonempty) sets of* feasible solutions *(bitstrings) for these instances;*

- *an* objective function m_P *that gives for each pair* (x, y) *consisting of instance* $x \in I_P$ *and solution* $y \in S_P(x)$ *a positive integer* $m_P(x, y)$*, the* objective value*;*
- *a* goal $\text{goal}_P \in \{\min, \max\}$ *depending on whether* P *is a minimization or a maximization problem.*

We denote by $S^*_P(x) \subseteq S_P(x)$ the set of optimal solutions for an instance $x \in I_P$, i.e. for every $y^* \in S^*_P(x)$,

$$m_P(x, y^*) = \text{goal}_P\{m_P(x, y) \mid y \in S_P(x)\}.$$

The objective value attained by an optimal solution for an instance x is denoted by $m^*_P(x)$.

Definition 2.1.2 *An optimization problem* P *is in the class* NPO *if*

- *the set of instances* I_P *can be recognized in polynomial time;*
- *there is a (monotone nondecreasing) polynomial (say* q_P*) such that* $|y| \leq q_P(|x|)$ *for any instance* $x \in I_P$ *and any feasible solution* $y \in S_P(x)$*;*
- *for any instance* $x \in I_P$ *and any* y *with* $|y| \leq q_P(|x|)$*, one can decide in polynomial time whether* $y \in S_P(x)$*;*
- *there is a (monotone nondecreasing) polynomial (say* r_P*) such that the objective function* m_P *is computable in* $r_P(|x|, |y|)$ *time for any* $x \in I_P$ *and* $y \in S_P(x)$*.*

All problems considered below will be in NPO and all considered classes will be subclasses of NPO.

Note that for any problem $P \in$ NPO and any $n \in \mathbb{N}$ the maximum objective value of instances of size n, i.e. $\max\{m_P(x, y) \mid x \in I_P, |x| = n, y \in S_P(x)\}$, is bounded by $2^{r_P(n, q_P(n))}$, as the objective function value of any $x \in I_P$ and $y \in S_P(x)$ can be represented by at most $r_P(|x|, |y|) \leq r_P(|x|, q_P(|x|))$ bits. Let $M_P(n) = 2^{r_P(n, q_P(n))}$.

If one compares NPO to the class NP, then PO is the 'equivalent' of P. *PO* is the class of problems in NPO for which an optimal solution $y^* \in S^*_P(x)$ can be computed in time polynomial in $|x|$ for any $x \in I_P$. Paz and Moran [221] proved that P=NP implies PO=NPO and vice versa. Because it is not expected that all problems in NPO also fall in PO, several classes have been defined that contain NPO-problems for which an *approximate* solution can be found in polynomial time. Approximation algorithms are classified by two properties: their running time and their approximation ratio.

Definition 2.1.3 ([19, 115]) *For an optimization problem* $P \in NPO$*, any* $x \in I_P$*, and any* $y \in S_P(x)$*, the* approximation ratio *achieved by* y *for* x *is*

$$R(x, y) = \max\left\{\frac{m_P(x, y)}{m^*_P(x)}, \frac{m^*_P(x)}{m_P(x, y)}\right\}.$$

Problem class	Running time	Approx. ratio
APX	Polynomial in $\|x\|$	c
PTAS	Polynomial in $\|x\|$ (for every fixed ϵ)	$(1+\epsilon)$
FPTAS	Polynomial in $\|x\|$ and $1/\epsilon$	$(1+\epsilon)$
FIPTAS	Polynomial in $\|x\|$	$(1+\epsilon)$
PO	Polynomial in $\|x\|$	1

Table 2.1: Problem classes and the distinguishing properties of the approximation algorithms admitted by problems in a particular class.

We say y is within (a factor) r *of* $m^*_P(x)$ *if* $R(x,y) \leq r$. *The* approximation ratio *of an algorithm $\mathcal{A}$ is defined as*

$$R_{\mathcal{A}} = \max\{R(x, \mathcal{A}(x)) \mid x \in I_P\}.$$

Observe that irrespective of whether $\text{goal}_P = \min$ or $\text{goal}_P = \max$, the approximation ratio is a number that is at least 1. Sometimes however, in the case where $\text{goal}_P = \max$, we will instead use

$$R'(x,y) = \frac{1}{R(x,y)} = \frac{m_P(x,y)}{m^*_P(x)},$$

which is at most 1. To keep the exposition simple, we will only use $R(x,y)$ in this chapter.

Any textbook on approximation algorithms covers at least the classes of Table 2.1. The table should be interpreted as follows: *PTAS* for instance is the class of optimization problems P in NPO having a *ptas*, i.e. having an algorithm $\mathcal{A}$ such that for any instance $x \in I_P$ and any $\epsilon > 0$, $\mathcal{A}(x,\epsilon)$ runs in time polynomial in $|x|$ for every fixed ϵ and the solution output by $\mathcal{A}(x,\epsilon)$ has approximation ratio $(1+\epsilon)$. We use lowercase for a scheme name and uppercase for the name of the corresponding class (i.e. ptas and PTAS).

APX is the class of problems having a *constant-factor approximation algorithm*, meaning a polynomial-time algorithm that returns a solution of approximation ratio c, for some fixed constant c.

The class FIPTAS (Fully Input-Polynomial-Time Approximation Scheme) in Table 2.1 is a new class. Clearly, FIPTAS=PO (use $\epsilon = 1/M_P(|x|)$), but the reason for defining this class will become apparent later.

A relatively new class that is of increasing interest is EPTAS [26, 53].

Definition 2.1.4 *Algorithm $\mathcal{A}$ is an* efficient polynomial-time approximation scheme (eptas) *for problem P if there is a computable function $f : \mathbb{Q}^+ \to \mathbb{N}$ such that for any $x \in I_P$ and any $\epsilon > 0$, $\mathcal{A}(x,\epsilon)$ runs in time $f(1/\epsilon)$ times a polynomial in $|x|$ and the solution output by $\mathcal{A}(x,\epsilon)$ has approximation ratio $(1+\epsilon)$. An NPO-problem is in the class* EPTAS *if and only if it has an eptas.*

The popularity of eptas is not only due to the separate dependence on $1/\epsilon$ and instance size in the running time, but also to the beautiful relation to the widely researched class FPT: any problem admitting an eptas is also in FPT w.r.t. its standard parameterization [26, 53]. An intriguing exploration of the type of problems that admit an eptas may be found in Cai et al. [48].

It is well-known that PO $\subseteq$ FPTAS $\subseteq$ EPTAS $\subseteq$ PTAS $\subseteq$ APX $\subseteq$ NPO. In most cases, the inclusion is strict (unless P=NP), except that EPTAS $\subset$ PTAS unless FPT=W[1] [26, 53]. The question whether FPT=W[1] is an open problem for fixed-parameter complexity theory akin to the question whether P=NP for classic complexity theory (see e.g. Downey and Fellows [92]).

2.2 Asymptotic Approximation Schemes

We introduce a new type of approximation scheme, the asymptotic approximation scheme.

Definition 2.2.1 *An approximation scheme $\mathcal{A}$ for an optimization problem $P \in NPO$ is* asymptotic *if there is a computable function $a : \mathbb{Q}^+ \to \mathbb{N}$ (the* threshold function*) such that for any $\epsilon > 0$ and any $x \in I_P$, it returns a $y \in S_P(x)$ and if $|x| \geq a(1/\epsilon)$, then y is within $(1+\epsilon)$ of $m^*_P(x)$.*

This definition leads to the following classes of asymptotic approximation schemes.

Class	Running time	Approx. ratio
PTAS$^\omega$	Polynomial in $\lvert x\rvert$ (for every fixed ϵ)	$(1+\epsilon)$ if $\lvert x\rvert \geq a(1/\epsilon)$
FPTAS$^\omega$	Polynomial in $\lvert x\rvert$ and $1/\epsilon$	$(1+\epsilon)$ if $\lvert x\rvert \geq a(1/\epsilon)$
FIPTAS$^\omega$	Polynomial in $\lvert x\rvert$	$(1+\epsilon)$ if $\lvert x\rvert \geq a(1/\epsilon)$

Example 2.2.2 *Maximum Independent Set has a fiptas$^\omega$ on bounded ply disk graphs (see the proof of Theorem 7.3.9). Disk graphs are intersection graphs of disks in the plane. A set of disks has ply γ if γ is the smallest integer such that any point of the plane is strictly contained in at most γ disks. One can find in $O(n^{10}\log^4 n)$ time an independent set of a graph induced by n disks. If for some $\epsilon > 0$ an odd integer k can be chosen such that $\max\{5, 4(1+\epsilon)/\epsilon\} \leq k \leq c_1 \log n/\log(c_2\gamma)$ (where c_1, c_2 are fixed constants), then this independent set will be within $(1+\epsilon)$ of the optimum. If $\gamma = \gamma(n) = O(n^{o(1)})$, such an integer exists if $|x| \geq n \geq a(1/\epsilon)$ for some function a.*

We start with some easy observations about the asymptotic classes.

Proposition 2.2.3 *The following relations hold:*

- *FIPTAS$^\omega$ $\subseteq$ FPTAS$^\omega$ $\subseteq$ PTAS$^\omega$ and*
- *FIPTAS $\subseteq$ FIPTAS$^\omega$, FPTAS $\subseteq$ FPTAS$^\omega$, PTAS $\subseteq$ PTAS$^\omega$.*

The relations given by this proposition are straightforward and one might expect that the inclusions are strict under some hardness condition. However, this turns out not to be true for all of them. We can in fact prove very interesting equivalences and at the same time tie these new classes to existing approximation classes, in particular to EPTAS.

Theorem 2.2.4 *EPTAS = FPTAS$^\omega$ = FIPTAS$^\omega$.*

Proof: We first show that EPTAS $\subseteq$ FIPTAS$^\omega$. Let $P \in$ EPTAS and let $\mathcal{A}$ be an eptas for P with running time at most $p(|x|) \cdot f(1/\epsilon)$ for some computable function f and polynomial p. Construct a fiptas$^\omega$ for P as follows. Given an arbitrary instance $x \in I_P$ and an arbitrary $\epsilon > 0$, run $\mathcal{A}(x, \epsilon)$ for $p(|x|) \cdot |x|$ time steps. If $\mathcal{A}(x, \epsilon)$ finishes, return the solution given by $\mathcal{A}(x, \epsilon)$. Otherwise, return $\mathcal{A}(x, 1/2)$. This algorithm clearly runs in time polynomial in $|x|$ and always returns a feasible solution. Furthermore if $|x| \geq f(1/\epsilon)$, $\mathcal{A}(x, \epsilon)$ always finishes and returns a feasible solution with approximation ratio $(1+\epsilon)$. Hence we constructed a fiptas$^\omega$ for P with $a = f$.

We next prove that FPTAS$^\omega$ $\subseteq$ EPTAS. Let $P \in$ FPTAS$^\omega$ and let $\mathcal{A}$ be an fptas$^\omega$ for P with threshold function a. Construct an eptas as follows. Given an arbitrary instance $x \in I_P$ and an arbitrary $\epsilon > 0$, compute $a(1/\epsilon)$. By assumption, $a(1/\epsilon)$ is computable. The amount of time it takes to compute $a(1/\epsilon)$ is some computable function depending on $1/\epsilon$. If $|x| \geq a(1/\epsilon)$, simply compute and return $\mathcal{A}(x, \epsilon)$ in time polynomial in $|x|$ and $1/\epsilon$. If $|x| < a(1/\epsilon)$, proceed as follows. As FPTAS$^\omega$ $\subseteq$ NPO, any feasible solution for x has size at most $q(|x|)$ for some polynomial q. Furthermore, given any y with $|y| \leq q(|x|)$, one can determine in polynomial time whether $y \in S_P(x)$. The objective value of a feasible solution can also be computed in polynomial time. Hence by employing exhaustive search, one can find a $y^* \in S^*_P(x)$ in time

$$2^{q(|x|)} \cdot r_P(|x|, q(|x|)) = 2^{q(a(1/\epsilon))} \cdot \mathrm{poly}(a(1/\epsilon)).$$

The result is an eptas for P with appropriately defined function f.

Since FIPTAS$^\omega$ $\subseteq$ FPTAS$^\omega$, we have EPTAS $\subseteq$ FIPTAS$^\omega$ $\subseteq$ FPTAS$^\omega$ $\subseteq$ EPTAS, and hence all classes must be equal. □

The exponential increase in running time in the reduction from an fptas$^\omega$ to an eptas might be reduced by using an exact or fixed-parameter algorithm specific to the problem.

The equivalence of F(I)PTAS$^\omega$ and EPTAS allows an indirect proof of the existence of an eptas for a problem, where a direct proof seems more difficult. For instance, Maximum Independent Set on disk graphs of bounded ply has a fiptas$^\omega$ (Example 2.2.2) and thus, as an immediate consequence of Theorem 2.2.4, it also has an eptas.

We now show that PTAS$^\omega$ and PTAS are equivalent.

Theorem 2.2.5 *PTAS = PTAS$^\omega$.*

Proof: By Proposition 2.2.3 it suffices to prove that $\text{PTAS}^\omega \subseteq \text{PTAS}$. Let $P \in \text{PTAS}^\omega$ and let $\mathcal{A}$ be a ptas^ω for P. For an arbitrary instance $x \in I_P$ and an arbitrary $\epsilon > 0$, compute $a(1/\epsilon)$. If $|x| \geq a(1/\epsilon)$, compute and return $\mathcal{A}(x, \epsilon)$. Otherwise, apply the same exhaustive search trick as in the proof of Theorem 2.2.4. The result is a ptas for P. $\square$

This implies that $\text{FPTAS}^\omega \subset \text{PTAS}^w$, unless FPT=W[1].

Chapter 3

Guide to Geometric Intersection Graphs

The introduction of this thesis presented geometric intersection graphs in the context of their many application areas, including wireless networks, computational biology, and map labeling. This chapter aims to provide a more formal and thorough introduction to these interesting graph classes. In particular, we will be concerned with the structure of (geometric) intersection graphs and the complexity of recognizing such graphs.

First however, we give a formal definition of intersection graphs and geometric intersection graphs. These are supported by examples of classes of (geometric) intersection graphs being studied in the literature. We describe the inclusion relations of these classes. We give special attention to some classes that are being studied in relation to wireless communication networks. We also expose connections between geometric intersection graphs and both general graphs and planar graphs. General graphs are shown to be intersection graphs of convex objects in $\mathbb{R}^3$ [268] and planar graphs are the intersection graphs of internally disjoint boxes in $\mathbb{R}^3$ [251] and of internally disjoint triangles [80], disks [169], or smooth convex objects in $\mathbb{R}^2$ [237].

We then consider questions related to recognizing geometric intersection graphs. We present several positive results on the recognition of certain classes of geometric intersection graphs, such as interval graphs, which can be recognized in polynomial time [38]. Most of these results however are on intersection graphs of one-dimensional objects. Recognizing intersection graphs of two-dimensional objects, such as intersection graphs of disks or polygons, is often NP-hard [41, 174, 42]. In fact, for intersection graphs of disks of equal radius, one can even give an approximation hardness result [184]. Several recognition problems are NP-hard and not NP-complete, as membership of NP is not known for these problems. A full discussion of this particular issue is deferred to Chapter 4.

3.1 Intersection Graphs

We define intersection graphs as follows. All considered graphs are simple, finite, and undirected, unless otherwise stated.

Definition 3.1.1 *Given a universe* $\mathbb{U}$ *and some finite collection* $\mathcal{S}$ *of subsets of* $\mathbb{U}$*, the* intersection graph $G = G[\mathcal{S}]$ *of* $\mathcal{S}$ *is the graph where each* $\mathcal{S}(u) \in \mathcal{S}$ *corresponds to a unique vertex* $u \in V(G)$ *and there is an edge between two vertices if and only if the corresponding subsets intersect (i.e.* $(u, v) \in E(G)$ *if and only if* $\mathcal{S}(u) \cap \mathcal{S}(v) \neq \emptyset$*).*

If a graph G is (isomorphic to) the intersection graph of some set system $(\mathbb{U}, \mathcal{S})$, then this set system is a *representation* of G, whereas G is said to be *induced* by $(\mathbb{U}, \mathcal{S})$. Usually, we will not distinguish between the sets of the set system and the vertices they correspond to.

It is easy to prove that any graph is an intersection graph of some set system [248]. Given a graph G, take $\mathbb{U} = E(G)$ and let $\mathcal{S} = \{S_v \mid v \in V(G)\}$, where $S_v = \{e = (u, v) \in E(G)\}$. One can however prove more interesting theorems. For instance, *chordal graphs* (the class of graphs having no induced cycle of length greater than three) are precisely the intersection graphs of subtrees of a fixed tree (see e.g. [135, 46, 118, 265]). A nice overview of these and other results can be found in the books by McKee and McMorris [208] and Spinrad [246] and the survey by Kozyrev and Yushmanov [171].

It is not easy to give a precise definition of geometric intersection graphs, as 'geometry' is not easily defined. For the purpose of this thesis however, the following definition is used.

Definition 3.1.2 *Given some finite collection* $\mathcal{S}$ *of subsets of* $\mathbb{R}^d$ *for some* $d \geq 0$*, the intersection graph* $G[\mathcal{S}]$ *is called a* geometric intersection graph*.*

Commonly, there is a restriction on the nature of the sets as well. This is expressed in the following definition.

Definition 3.1.3 *Let* $\mathcal{A}$ *be a set of subsets of* $\mathbb{R}^d$ *for some* $d > 0$*. Then* G *is an* $\mathcal{A}$-intersection graph *if it is (isomorphic to) the intersection graph* $G[\mathcal{S}]$ *for some collection* $\mathcal{S}$ *of translated copies of objects in* $\mathcal{A}$*.*

It is these restrictions that we will be most interested in. We note here that we only consider sets $\mathcal{A}$ of objects that are either all closed sets or all open sets. We prohibit mixing open and closed sets to simplify the presentation.

3.1.1 Interval Graphs and Generalizations

Perhaps the first and most frequently studied class of geometric intersection graphs are *interval graphs*. Here the universe is $\mathbb{R}^1$ and the sets are *intervals* or *segments* of the real line (i.e. connected subsets of $\mathbb{R}^1$). Sometimes these intervals are forced to have unit length (i.e. equal length, usually assumed to be 1), leading to *unit interval graphs* (or *indifference graphs*).

Several characterizations of interval graphs are known. Lekkerkerker and Boland [189] proved that G is an interval graph if and only if it is a chordal graph and has no asteroidal triple (a set of three vertices of the graph, any two of which are connected by a path containing no vertex of the neighborhood of the third vertex). Gilmore and Hoffman [123] proved that G is an interval graph if and only if it is a chordal graph and a cocomparability graph (a graph with a partial order on the vertices where two vertices are adjacent if and only if they are not related in the ordering). Unit interval graphs are precisely the interval graphs with no $K_{1,3}$ induced subgraph [268].

Interval graphs are recognizable in linear time [38]. Several optimization problems that are NP-hard on general graphs are easily polynomial-time solvable on interval graphs, such as Maximum Clique and Minimum Vertex Cover. An important structural property of interval graphs is that interval graphs are perfect graphs. Further properties and characterizations of (unit) interval graphs may be found in the books by Golumbic [126] and McKee and McMorris [208], and the paper by Kozyrev and Yushmanov [171].

Most classes of geometric intersection graphs can be viewed as generalizations of interval graphs, either because they stick to the idea of segments or because the considered geometric objects are intervals when restricted to $\mathbb{R}^1$. We first consider other classes of intersection graphs of segments.

In *multi-interval graphs*, we allow the sets $\mathcal{S}(u)$ to consist of multiple intervals on the real line. Clearly, any graph is a multi-interval graph by taking sufficiently many intervals for each vertex. However, determining whether the minimum number of intervals needed per vertex (the *interval number*) is at most k is an NP-complete problem for any fixed integer $k \geq 2$ [269].

Tolerance (interval) graphs place a restriction on the nature of the intersection when determining whether two vertices are adjacent. Each vertex is assigned a positive value (its *tolerance*) and two vertices are adjacent if the value of some function of the intersection of their intervals is at least the value of some function of their tolerances. For instance, in a *max-tolerance graph*, the length of the intersection should be at least the maximum of the tolerances. Note that the idea behind tolerance graphs is not necessarily restricted to interval graphs, but can be applied to any intersection graph [208]. The book by Golumbic and Trenk [127] provides a good overview of this subject. Here however we consider only tolerance interval graphs.

If instead of intervals on the real line, one considers intervals (arcs) on a circle, *circular-arc graphs* are obtained. Even though they are only a slight generalization of interval graphs, their structure is fundamentally different. Circular-arc graphs are not necessarily chordal or perfect, as they can contain induced cycles of any length. In fact, no characterization in terms of forbidden subgraphs is known [208]. They can however be recognized in linear time [207]. If all intervals have equal length (*unit circular-arc graphs*), a structural characterization does exist [253] and recognition is possible in linear time [195].

If we stick to line segments, but in $\mathbb{R}^2$, we arrive at *k-DIR graphs*, which are intersection graphs of line segments that can point in one of k directions. Recognizing graphs in this class is NP-complete for any fixed $k \geq 2$ [173]. Intersection graphs of piecewise linear curves consisting of at most k line segments (*k-segment intersection graphs*) are NP-hard to recognize for fixed $k \geq 2$ [179].

It is not known whether the recognition problem is in NP. Demonstrating membership of NP seems hard, as there are 1-segment intersection graphs where the coordinates of the endpoints of the segments must be double exponential integers [179].

Generalizing further, we have intersection graphs of arbitrary simple curves (*string graphs*). String graphs are NP-hard to recognize [172]. Recognizing string graphs surprisingly is in NP [233] and thus NP-complete. Note that every graph is the intersection graph of simple curves in $\mathbb{R}^3$. This is not true however in $\mathbb{R}^2$ (i.e. not every graph is a string graph) [241, 242, 98].

A relatively recent overview of the results in this area and a description of further classes can be found in the course notes of Kratochvíl [175].

3.1.2 Intersection Graphs of Higher Dimensional Objects

Another way to generalize interval graphs, more in line with further topics of this thesis, is to consider d-dimensional objects that are intervals if $d = 1$. A good example are *intersection graphs of d-dimensional axis-parallel boxes*. A d-dimensional axis-parallel box is simply the Cartesian product of d orthogonal intervals, e.g. $[a_1, b_1] \times \cdots \times [a_d, b_d]$ for numbers $a_i < b_i$. For the case $d = 2$, these are also known as *rectangle intersection graphs*.

An important related property of general graphs is their *boxicity*, the minimum number d such that the given graph is isomorphic to an intersection graph of d-dimensional axis-parallel boxes. Roberts [228] (who was the first to study boxicity) proved that any n-vertex graph has boxicity at most $\lfloor n/2 \rfloor$. Determining the boxicity of a graph is NP-hard [71]. Recognizing whether the boxicity is at most d is NP-complete for any fixed $d \geq 2$ [173, 273, 196]. Note that for $d = 1$, the recognition problem is equal to the problem of recognizing interval graphs, which is in P.

Given the above, it is not hard to imagine a natural generalization of unit interval graphs. This leads to *intersection graphs of d-dimensional axis-parallel unit cubes*, which are d-dimensional axis-parallel boxes with side length equal to one. For $d = 2$, these are called *unit square intersection graphs*, or simply *unit square graphs*. The notion of *cubicity* can be defined analogously to boxicity. Roberts [228] bounded the cubicity of n-vertex graphs by $\lfloor 2n/3 \rfloor$, the tight example being complete k-partite graphs with $k = \lfloor n/3 \rfloor$. Recognizing graphs of cubicity d is NP-complete [273, 40, 70, 73] for any fixed $d \geq 2$. The case $d = 1$ is in P, as these are precisely the unit interval graphs.

As an intermediate step between unit cubes and rectangles, one could consider axis-parallel cubes of arbitrary side length. In two dimensions, these are *square (intersection) graphs*. There seem to be no results on the minimum dimension d needed for a graph to be isomorphic to an intersection graph of d-dimensional axis-parallel cubes. Breu [40] showed that testing whether a graph is an intersection graph of squares where the ratio between the size of the largest and of the smallest square is some fixed constant $\rho \geq 1$, is NP-hard. The recognition problem of intersection graphs of arbitrary-sized squares is

NP-hard as well [182].

A further relevant class of geometric intersection graphs are *triangle intersection graphs*. It is known that any planar graph is an intersection graph of internally disjoint triangles [80] and conjectured that any planar graph is an intersection graph of homothetic triangles [176]. The class of intersection graphs of isosceles right triangles is also interesting, as this class was shown to be equivalent to max-tolerance interval graphs [160].

The definitions of box, cube, and triangle intersection graphs invite to generalizations to intersection graphs of other geometric objects. One could for instance consider *intersection graphs of convex objects*. Again, for $d = 1$, this corresponds to interval graphs and hence they can be recognized in linear time. For $d = 2$, recognition is NP-hard [178]. Proving membership of NP is likely to be very difficult, as Pergel (see [176]) proved that convex object intersection graphs exist for which any integer representation requires double exponential integers. However, recognizing the intersection graph of scaled and translated copies of a fixed convex polygon is both NP-hard [182] and in NP [261], and thus NP-complete. The recognition problem for higher dimensions is trivial, because as shown later in Theorem 3.3.1, any graph is the intersection graph of three-dimensional convex objects.

Finally, we treat a subclass of planar convex object intersection graphs, namely the class of *polygon-circle graphs*. These are the intersection graphs of polygons inscribed in a circle, i.e. all corners of the polygons should lie on one given circle. Recognizing these graphs is NP-complete [180, 224], but polynomial if the girth is greater than four [181]. It can be proved by a simple argument that chordal graphs are polygon-circle graphs (see Corollary 3.3.3).

3.2 Disk Graphs and Ball Graphs

A different generalization of interval graphs are *ball graphs*, intersection graphs of d-dimensional balls. A d-dimensional ball is given by its center and consists of all points within a certain distance. For $d = 2$, these are the well-known disk graphs.

Definition 3.2.1 *A graph isomorphic to an intersection graph of two-dimensional balls (i.e. disks) is called a* disk graph*.*

We emphasize disk graphs as they motivated most of the research in this thesis.

Recognizing disk graphs is NP-hard [174], even if the ratio between the radii of the largest and smallest disk is bounded by a constant [41]. The complexity of recognizing intersection graphs of higher dimensional balls is unknown, but expected to be NP-hard [42].

When generalizing unit interval graphs, we get *unit ball graphs*, i.e. intersection graphs of d-dimensional balls of equal radius.

Definition 3.2.2 *A graph isomorphic to an intersection graph of two-dimensional balls (i.e. disks) of radius* $1/2$ *is called a* unit disk graph*.*

Although we define unit disk graphs to have a representation with disks of radius 1/2, this number is mostly chosen for convenience. The most important property is that all disks have equal radius. By scaling, one can always assume this common radius to be 1/2.

We should note here that disk graphs and unit disk graphs are really different graph classes. For instance, unit disk graphs cannot have a $K_{1,6}$ or $K_{2,3}$ induced subgraph [257], whereas these graphs are disk graphs. An example of a graph that is not a disk graph is $K_{3,3}$. In fact, all triangle-free disk graphs must be planar [199].

Instead of focusing on intersections of d-dimensional balls, another way to define unit ball graphs is to place n points in $\mathbb{R}^d$ and say that two vertices are adjacent if and only if they are at distance at most 1. If $d = 1$, this definition corresponds to the definition of indifference graphs.

Unit disk graphs are sometimes also called *geometric graphs* (see for example DeWitt and Krieger [86]). This should not be confused with the currently used definition of geometric graphs, namely graphs where each vertex is assigned a point in $\mathbb{R}^d$ and edges are drawn as straight lines between the points. Although under this modern definition, a representation of a unit disk/ball graph induces a geometric graph, it can be readily observed that not every geometric graph is a unit disk/ball graph.

The *(unit) sphericity* of a graph is the minimum d such that the graph is isomorphic to an intersection graph of d-dimensional unit balls [144, 145, 109]. Maehara [198] proved that the sphericity of any n-vertex graph G is at most $n - \omega(G)$, where $\omega(G)$ denotes the size of the largest clique of G. Moreover, this bound is essentially tight. Recognizing unit disk graphs (i.e. graphs of sphericity two) is NP-hard, even if the graph is planar or has sphericity at most three [42]. Recognizing graphs of sphericity at most three is also NP-hard, but the complexity for higher constants is unknown (though conjectured to be NP-hard) [42].

Kuhn, Moscibroda, and Wattenhofer [184] provide a strengthening of the NP-hardness result in two dimensions. Define

$$q(G, c) = \frac{\max_{(u,v)\in E(G)} \|c_u - c_v\|}{\min_{(u,v)\notin E(G)} \|c_u - c_v\|}$$

as the *quality* of a mapping $c : V(G) \to \mathbb{R}^2$ of a unit disk graph G, where c_u (c_v) denotes the location of the center of the disk corresponding to u (v). Note that any (nontrivial) unit disk graph by definition has a mapping of quality less than one. Kuhn, Moscibroda, and Wattenhofer [184] show however that it is NP-hard to decide if a mapping of quality at most $\sqrt{3/2} - \epsilon$ exists, where ϵ tends to 0 as the number of vertices of the graph approaches infinity. On the positive side, Moscibroda et al. [215] give a polynomial-time algorithm yielding a mapping of quality $O((\log^{5/2} n) \cdot \sqrt{\log \log n})$. This clearly leaves a large gap and a major open question.

An important special case of (unit) ball graphs are intersection graphs of internally disjoint (unit) balls, called *(unit) ball touching graphs* or *(unit)* *ball contact graphs*. In two dimensions, ball contact graphs are also called *disk contact graphs* or *coin graphs* [231]. Coin graphs are interesting, as they coincide with the class of all planar graphs [169] (see also Section 3.3.1). Hence these graphs are recognizable in linear time [152]. If however the ratio of the radii of the largest and smallest disk is any fixed constant, then the recognition problem becomes NP-hard [41]. The complexity for recognition of ball contact graphs in higher dimensions is open. For $d = 1$, (unit) ball contact graphs are disjoint unions of paths and are thus recognizable in linear time.

In the case of unit ball contact graphs, we know a bit more. Any n-vertex graph is a unit ball contact graph in dimension $n - 1$ [147]. Recognizing unit ball contact graphs is known to be NP-hard for dimension 2 [42], 3, 4, 5 [146], 8, 9, 24, and 25 [147]. The hardness proofs for dimensions 5, 9, and 25 follow from a construction of Kirkpatrick and Rote (see Hliněný [146]) who showed that a graph G is isomorphic to a unit d-ball contact graph if and only if $G \oplus K_2$ is isomorphic to a unit $(d + 1)$-ball contact graph, where $G \oplus K_2$ is obtained from the disjoint union of G and K_2 by adding all edges between the vertices of the summands.

Another generalization of disk graphs are *intersection graphs of noncrossing arc-connected sets* [174]. A set is *arc-connected* if between any two points of the set an arc can be drawn containing only points of the set. The class of intersection graphs of arc-connected sets in the plane coincides with the class of string graphs [174]. Two arc-connected sets X and Y are said to be *noncrossing* if both $X - Y$ and $Y - X$ are arc-connected. Intersection graphs of noncrossing arc-connected sets are not equivalent to string graphs, as $K_{3,3}$, which is a string graph, is not an intersection graph of noncrossing arc-connected sets [174]. Recognizing intersection graphs of noncrossing arc-connected sets in the plane is NP-hard [174]. However, each graph is an intersection graph of three-dimensional noncrossing arc-connected sets.

Noncrossing arc-connected sets in the plane are essentially *k-admissible regions* for some even integer k, which are a collection of noncrossing arc-connected sets each bounded by a simple closed Jordan curve, such that each pair of curves intersects at most k' times, for some even $k' \leq k$ [226]. We call a collection of 2-admissible regions a collection of *pseudo-disks*.

We should note that several of the NP-hard recognition problems described in this chapter that are not in P, such as the problem of recognizing disk graphs, are not known to be in NP. In particular, we know of no polynomially-sized representation for these graph classes. However, one can prove that the recognition problems are in PSPACE [179, 147, 51].

3.2.1 Models for Wireless Networks

In the introduction (Chapter 1), we mentioned wireless networks as one of the main application areas of geometric intersection graphs. Particularly (unit) disk graphs are frequently used as a model in this setting. To bring these models even closer to the situations encountered in practice, several more sophisticated graph models have been proposed. We describe some of them. Further models can be found in the survey by Schmid and Wattenhofer [236].

A restriction of disk graphs is the following. Suppose that the radii of the disks model broadcasting ranges. Then u can hear v if and only if u is within v's broadcasting range, i.e. if u lies within the disk centered on v. Use this to determine the adjacency of vertices in the graph (so that u and v are adjacent if and only if u lies within the disk centered on v or v lies within the disk centered on u) and one obtains a *containment disk graph* [129, 199]. Malesińska [199] proved that the class of containment disk graphs is not contained in the class of disk graphs, as $K_{3,3}$ is a containment disk graph, but not a disk graph (recall that triangle-free disk graphs are planar). It is unclear whether a disk graph exists that is not a containment disk graph. However, any unit disk graph is a containment disk graph.

We presented the containment disk graph as an undirected graph. Given its motivation however, it makes more sense to define it as a directed graph. In this case, there is a directed edge from v to u if and only if u is contained in the disk centered on the location of v. This graph is called a *directed disk graph* [93], but in many cases it is also referred to as a *(directed) geometric radio network* [69].

A graph class that generalizes both disk graphs and containment disk graphs is the class of *double disk graphs* [129, 199, 94]. As the name implies, centered on the location of a vertex are two disks, s and b, such that the radius of b is at least the radius of s. Then two vertices u and v are adjacent if and only if $s(u)$ and $b(v)$ intersect or $s(v)$ and $b(u)$ intersect. The idea behind this graph model is that any wireless device has a range within which it can communicate with other devices and a larger range within which its signal interferes with the signals of other devices.

Another generalization of unit disk graphs is the quasi unit disk graph. In a unit disk graph, two vertices are adjacent if and only if the distance between their locations is at most one. Usually however, the probability of successfully connecting to another device decreases as it is further away from the source of the signal. In a *quasi unit disk graph*, given some $\rho \in [0, 1]$, two vertices are adjacent if they are within distance ρ, can be adjacent if they are within distance more than ρ but at most one, and are not adjacent if their distance is more than one [187]. Note that the behavior is undefined if the distance is in $(\rho, 1]$. This could be determined by an adversary or some probabilistic model. Kuhn, Moscibroda, and Wattenhofer [184] prove that recognizing ρ-quasi unit disk graphs with $\rho \geq \sqrt{1/2}$ is NP-hard.

Both unit disk graphs and quasi unit disk graphs have the property that the size of any independent set of the r-neighborhood of any vertex is polynomial in r, where the r-neighborhood of a vertex u consists of all vertices having a path of length at most r to u. This behavior can be used to define a class of graphs. A *bounded independence graph* or a *graph of polynomially-bounded growth* is a graph where for any r and any vertex u, all independent sets in the r-neighborhood of u have cardinality polynomial in r [219].

3.3 Relation to Other Graph Classes

Geometric intersection graphs have many relations to other well-known graph classes, for which it might be surprising that they are contained in a particular class of geometric intersection graphs. Some of these relations were already discussed in the previous section. Here we expand on these results and (where possible) sketch a proof. In particular, we will discuss the strong connections between geometric intersection graphs and planar graphs.

We begin by showing that any graph is an intersection graph of (internally disjoint) three-dimensional convex polytopes. This result is frequently attributed to Wegner [268]. Wegner himself [268, p. 28] however attributes it to Grünbaum, but this proof seems to be unpublished. Kalinin [158] also gives a proof. Here we follow Wegner's proof [268].

Theorem 3.3.1 *Any graph is the intersection graph of a set of (internally disjoint) three-dimensional convex polytopes.*

Proof: For any integer n, there exists a family $\mathcal{S}$ of n internally disjoint three-dimensional convex polytopes with nonempty interior, any two of which intersect. Moreover, the intersection of polytopes u and v is a (two-dimensional) facet of u or v. Finding such a family is known as Crum's problem [33, 227] and was solved by Besicovitch [33] and Rado [227].

Let G be any n-vertex graph and $\mathcal{S}$ a family as described above. Now let $\mathcal{S}'$ be obtained from $\mathcal{S}$ by taking for each $s \in \mathcal{S}$ a convex subset of the interior of s. Furthermore, for any $(u, v) \in E(G)$, choose a point p_{uv} in the interior of the intersection of $\mathcal{S}(u)$ and $\mathcal{S}(v)$. Let P_u denote the set of such points involving vertex u. Let $\widetilde{\mathcal{S}}$ be the set obtained by taking for each $u \in V(G)$ the convex hull of $\mathcal{S}'(u)$ and P_u. If H is the graph induced by $\widetilde{\mathcal{S}}$, then clearly $E(H) \supseteq E(G)$. Suppose that $(u, v) \in E(H) - E(G)$. Without loss of generality, the intersection of $\mathcal{S}(u)$ and $\mathcal{S}(v)$ is a facet of $\mathcal{S}(u)$. By the choice of P_u and $\mathcal{S}'(u)$, there is a hyperplane separating $\widetilde{\mathcal{S}}(u)$ and this facet. As $\widetilde{\mathcal{S}}(u)$ ($\widetilde{\mathcal{S}}(v)$) is a convex subset of $\mathcal{S}(u)$ ($\mathcal{S}(v)$), $\widetilde{\mathcal{S}}(u)$ and $\widetilde{\mathcal{S}}(v)$ cannot intersect. This contradicts that $(u, v) \in E(H)$. Hence $E(H) = E(G)$. □

Note that the constructed polytopes in fact have at most one point in common.

So what about intersection graphs of two-dimensional convex polytopes? If the polytopes are internally disjoint, they can be fully characterized (see Theorem 3.3.5). If we allow arbitrary intersections however, no characterization is known. Wegner [268, p. 25] showed that K_5 with each edge bisected is not the intersection graph of convex two-dimensional polytopes, since this would imply a planar drawing of K_5. (This is also implied by a result of Sinden [241, 242] and Ehrlich, Even, and Tarjan [98], who showed that this bisection of K_5 is not a string graph.)

We can give the following positive result. A planar graph is *outerplanar* if it has a planar embedding in which each vertex lies on the boundary of the outer face. We call this an *outerplanar embedding*.

Theorem 3.3.2 *The intersection graph G of a collection $\mathcal{H}$ of connected subgraphs of a fixed outerplanar graph H is a polygon-circle graph.*

Proof: We may assume that H is connected, otherwise we apply the proof given below to each connected component of H. Consider an outerplanar embedding of H. Going along the entire boundary of the outer face, let $v_1, \ldots, v_k$ be the vertices consecutively encountered. This induces an ordering $u_1, \ldots, u_n$ on the vertices of $V(H)$, where $n = |V(H)|$.

Now place n points $p_1, \ldots, p_n$ on (any arc of) a circle and map u_i to p_i. By the definition of the ordering on $V(H)$, this induces an outerplanar embedding of H. For any connected subgraph $\mathcal{H}(w) = \{u_{i_1}, \ldots, u_{i_k}\} \subseteq V(H)$ for $w \in V(G)$, let $C(w)$ be the convex hull of $p_{i_1}, \ldots, p_{i_k}$. One can verify that for $v, w \in V(G)$, $C(v)$ intersects $C(w)$ if and only if $\mathcal{H}(v)$ intersects $\mathcal{H}(w)$. □

As a corollary, we obtain a result by Duchet [95].

Corollary 3.3.3 *Any chordal graph is a polygon-circle graph.*

Proof: It is known that any chordal graph G is the intersection graph of a family $\mathcal{T}$ of subtrees of a fixed tree T. (This result is attributed to Surányi in [135]. Proofs can be found in [46, 118, 265]). Trees are clearly outerplanar. Now apply Theorem 3.3.2. □

3.3.1 Relation to Planar Graphs

Most results about the structure of geometric intersection graphs are related to planar graphs. Sometimes one can prove that planar graphs form a subclass of some class of geometric intersection graphs (or vice versa), but often one can give a full characterization. We will see examples of both.

Recall that k-interval graphs are intersection graphs of unions of k intervals. The interval number of a graph G is the minimum number k such that G is isomorphic to a k-interval graph. Scheinerman and West [235] proved that any planar graph has interval number at most three.

For higher dimensions, the following famous result is known.

Theorem 3.3.4 (Koebe [169]) *A graph G is planar if and only if G is isomorphic to a disk contact graph (coin graph).*

This result was rediscovered several times (see Sachs [231] for a history).

The radii of the disks in a disk contact representation of a planar graph are not necessarily polynomially bounded integers. The radii might differ by an exponential factor. Hansen [139] (see Malitz and Papakostas [200]) shows that there exist wheels for which disks of exponentially large radius are needed in any disk contact representation. Also, Breu and Kirkpatrick [41] showed that it is NP-hard to test whether such large disks are necessary.

Moreover, one cannot expect the radii of the disks to be integers. This would imply that any planar graph has a straight line embedding such that all edges have integer length. Although Geelen, Guo, and McKinnon [119] proved that any planar graph has an integer straight line embedding, Brightwell and Scheinerman [43] demonstrated that there exist planar graphs for which no integer straight line embedding can be induced by a disk contact representation. Otherwise one could trisect an angle of $\pi/3$ using ruler and compass.

An approximation to the coordinates and radii of a disk contact representation can be given though for 3-connected planar graphs. In fact, this holds for an even more general representation. Brightwell and Scheinerman [43] proved that for any n-vertex planar graph G, both G and the dual of G have a disk contact representation, $\mathcal{S}$ and $\mathcal{S}'$ respectively, such that for any edge $e = (u, v) \in E(G)$ and its dual edge $e^* = (f^*, g^*)$, the intersection point of $\mathcal{S}(u)$ and $\mathcal{S}(v)$ coincides with the intersection point of $\mathcal{S}'(f^*)$ and $\mathcal{S}'(g^*)$. Moreover, the line through the centers of $\mathcal{S}(u)$ and $\mathcal{S}(v)$ is perpendicular to the line through the centers of $\mathcal{S}'(f^*)$ and $\mathcal{S}'(g^*)$. Mohar [211, 212] gives an algorithm to determine the centers and radii of such a primal-dual representation to a precision of ϵ in time polynomial in n and $\max\{\log 1/\epsilon, 1\}$. A similar result is proved by Smith [243].

A consequence of Koebe's result is that planar graphs are string graphs, where any pair of curves is nowhere tangent and intersects at most twice. Chalopin, Gonçalves, and Ochem [56] showed that in fact one intersection suffices. This was a step forward in proving the following conjecture, sometimes referred to as *Scheinerman's conjecture* [234]: any planar graph is isomorphic to a 1-segment intersection graph. The conjecture was previously shown to hold for triangle-free [77], bipartite [141, 79], and several other types of planar graphs [78]. Recently, Chalopin and Gonçalves [55] managed to prove the conjecture, i.e. any planar graph is indeed isomorphic to a 1-segment intersection graph.

Disk graphs are generally not planar. However, Malesińska [199] proves that triangle-free disk graphs are planar. This also holds for triangle-free intersection graphs of pseudo-disks [174]. As planar graphs are disk intersection graphs, triangle-free (pseudo-)disk graphs are recognizable in polynomial time.

Characterizations similar to Koebe's theorem have been proved for other convex objects. The following result is implied by Koebe's theorem, but has a relatively easy proof due to Wegner [268], which we give below.

Theorem 3.3.5 *A graph is planar if and only if it is isomorphic to the intersection graph of a set of internally disjoint two-dimensional convex polytopes.*

Proof: The if-part is trivial. For the converse, let G be a planar graph. Augment G to an (edge) maximal planar graph G' by adding edges. Add a dummy vertex z in the unbounded face of G' and connect it to all vertices on the unbounded face of G'. Call the resulting graph G''. Since G'' is maximal planar, it is 3-connected and hence its dual is planar and 3-connected as well [213]. Following Stein [247] (see also Tutte [254] and Kelmans [162]), this implies that the dual of G'' has a straight line embedding such that each bounded face is convex and the face corresponding to z is the unbounded face. Hence we obtain a collection of internally disjoint two-dimensional convex polytopes whose intersection graph is G'. Using the same idea as in the proof of Theorem 3.3.1, we can remove unwanted edges to obtain a representation of G. □

A result of Thomassen [250] implies that finding such a representation actually takes linear time. Interestingly, Kratochvíl and Kuběna [177] showed that the complement of a planar graph is the intersection graph of a set of two-dimensional convex polytopes as well.

From the above proof, it seems that one might need polytopes with an arbitrary number of corners, but this is not necessarily the case. Some planar graphs are rectangle contact graphs. Thomassen [251] proved that G is a rectangle contact graph if and only if G is a proper subgraph of a 4-connected planar triangulation. Bipartite planar graphs are also rectangle intersection graphs [79, 141]. Generalizing in this direction, Thomassen [251] showed that any planar graph is the intersection graph of internally disjoint three-dimensional axis-parallel boxes.

De Fraysseix, Ossona de Mendez, and Rosenstiehl [80] proved that G is planar if and only if G is a triangle contact graph. Moreover, they gave a polynomial time algorithm to construct a representation by internally disjoint triangles. Although triangles are sufficient, the shapes of the triangles can be very different. Can one prove that they must be similar somehow?

Definition 3.3.6 *We say that two geometric objects are* homothetic *if one can be obtained from the other by only scaling and translating.*

There are planar graphs that are not the intersection graphs of internally disjoint homothetic triangles, although Kratochvíl and Pergel [176] conjecture that planar graphs are homothetic triangle graphs, i.e. without the constraint that the triangles should touch.

Observe that Koebe's result states that any planar graph is an intersection graph of internally disjoint homothetic disks. Schramm [237, 238] generalizes Koebe's theorem to homothetic copies of arbitrary convex planar bodies with smooth boundaries. The result actually is slightly more general.

Theorem 3.3.7 (Schramm [237, 238]) *Let G be any n-vertex planar graph and $\mathcal{A} = \{A_v \mid v \in V\}$ a collection of n planar convex bodies with smooth boundaries. Then G is the intersection graph of $\mathcal{S} = \{S_v \mid v \in V\}$, where S_v is a homothetic copy of A_v. Moreover, the objects in $\mathcal{S}$ are internally disjoint.*

Note that Theorem 3.3.7 requires the convex objects to have smooth boundaries and thus does not contradict the statement that planar graphs are not the intersection graphs of internally disjoint homothetic triangles.

Chapter 4

Geometric Intersection Graphs and Their Representation

As is clear from the previous chapter, a fundamental problem for most classes of geometric intersection graphs is how to recognize such graphs. The recognition problem is often NP-hard and in PSPACE, but membership of NP is not always known. One way of proving membership of NP is to find a representation of the graph that uses polynomially many bits (polynomial in the number of vertices of the graph). This is a *polynomial representation*.

For several classes of geometric intersection graphs, bounds on the number of bits needed to represent each object are known. These were already implicitly mentioned in Chapter 3. We mention some of them explicitly in Table 4.1. The main graph class missing in this table are (unit) disk graphs. We know of no polynomial or finite representation for this class.

This chapter gives new insight into whether polynomial representations exist for intersection graphs of any of a large class of geometric objects, called *scalable objects*, which includes convex objects.

We prove that any intersection graph of scalable objects has a representation using finitely many bits, i.e. using rationals. The main tool in this proof is the notion of ϵ-separation, a measure of the relative degree of overlap or disjointness of two objects. For several types of scalable objects (including disks and squares), we show that an intersection graph of such objects has a polynomial representation if and only if it has a representation that is polynomially separated (i.e. ϵ-separated where $\epsilon = 2^{-q(n)}$ for some polynomial q in the number of vertices n of the graph). We can even give an algorithm showing that the two are computationally equivalent as well. This equivalence might give a new way to prove or disprove the existence of polynomial representations for these classes of geometric intersection graphs.

4.1 Scalable and ϵ-Separated Objects

We start by formally defining the above notions. We then prove that any intersection graph of closed scalable objects has an ϵ-separated representation. The same holds for open scalable objects. Consequently, the classes of intersection graphs of closed scalable objects and their open counterparts coincide.

Graph class	Bound on largest coordinate	
Interval graphs	$2n$	(trivial)
Unit interval graphs	$\mathrm{O}(n^2)$	[70]
Rectangle int. graphs	$\mathrm{O}(n)$	[196]
Unit square graphs	$\mathrm{O}(n^2)$	[70, 73]
Homothetic convex polygon int. graphs	$2^{\mathrm{O}(n^4)}$	[261]

Table 4.1: The table gives a bound on the value of the largest object coordinate in some representation of a graph in the given class, where we assume that all object coordinates are integers greater than 0.

Throughout, we will assume objects to be either open or closed. For an object s, let $\mathrm{int}(s)$, $\mathrm{cl}(s)$, and $\mathrm{bd}(s)$ respectively denote the interior, the closure, and the boundary of s.

Definition 4.1.1 *A* scaling *of the space $\mathbb{R}^d$ by some $\tau > 0$ maps any point $p \in \mathbb{R}^d$ to $\tau \cdot p$. An object s is said to be* scaled around a point p *by $\tau > 0$ if (assuming s is the only object in the space) s is translated by $-p$, the space is scaled by τ, and then s is translated by p.*

Scaling an object around a point gives more control over the scaling, which is needed in the following definition.

Definition 4.1.2 *An object s is* scalable *if there is a point $p \in \mathrm{int}(s)$ such that for any $\tau \in \mathbb{R}_{>0} \backslash \{1\}$, scaling s around p by τ gives an object s' for which $\mathrm{cl}(s) \subseteq \mathrm{int}(s')$ or $\mathrm{cl}(s') \subseteq \mathrm{int}(s)$. If s is scalable, fix such a point p and call it* the scaling point *of s, denoted by c_s.*

Alternatively, we could demand that the distance between $\mathrm{bd}(s)$ and $\mathrm{bd}(s')$ is greater than zero. This would yield an equivalent definition.

The constraint that $c_s \in \mathrm{int}(s)$ is there for convenience. All results of this section would also hold if $c_s \notin \mathrm{cl}(s)$. The mixing of both object types is not considered here. Hence we restrict to $c_s \in \mathrm{int}(s)$.

We can easily determine which objects are scalable and which are not. An object s is said to be *strongly star-shaped* if there is a point $t_s \in \mathrm{int}(s)$ such that for any point $p \in s$ the straight line segment $\overline{t_s p}$ is contained in s, but does not contain any point of $\mathrm{bd}(s)$, except possibly p.

Proposition 4.1.3 *An object s is scalable if and only if it is strongly star-shaped.*

Proof: Suppose that s is scalable and has scaling point c_s. We claim that s is strongly star-shaped with $t_s = c_s$. For suppose there is a point $p \in s$ for which the straight line segment $\overline{c_s p}$ contains a point $z \in \mathrm{bd}(s)$, where $z \neq p$. By appropriately scaling (shrinking) s around c_s, we can map p to the position of

z. But then for this scaled object s', neither $\text{cl}(s) \subseteq \text{int}(s')$, since we shrunk s to obtain s', nor $\text{cl}(s') \subseteq \text{int}(s)$, by the preceding observation. This contradicts that s is scalable.

Now suppose that s is strongly star-shaped for some $t_s \in \text{int}(s)$. We claim that s is scalable with scaling point $c_s = t_s$. For let $\tau \in (0, 1)$ and let s' be the object obtained when scaling s around t_s by τ. Then any point $p \in s$ gets mapped to a point p' on the straight line segment $\overline{t_s p}$. Because $p' \neq p$ (unless $p = t_s$), $p' \in \text{int}(s)$. Hence $\text{cl}(s') \subseteq \text{int}(s)$. The case $\tau \in \mathbb{R}_{>1}$ is similar. □

Using this proposition, it is easy to see that for instance convex objects are scalable, as are L-shaped objects. Donut-shaped objects for example, such as a torus, are not scalable.

All objects we consider below are assumed to be scalable and hence we will not always mention this explicitly. Also, when scaling a scalable object, we implicitly mean scaling it around its scaling point.

We now define two measures of the degree of overlap or disjointness of a collection of objects.

Definition 4.1.4 *Two objects s and s' are* ϵ-distant *for some $\epsilon \geq 0$ if for any two vectors $\vec{a}$ and $\vec{b}$ with $\|\vec{a}\| = \|\vec{b}\| \leq \epsilon$, $s + \vec{a}$ and $s' + \vec{b}$ intersect if and only if s and s' intersect.*

Definition 4.1.5 *Two objects s and s' are* ϵ-separated *for some $0 \leq \epsilon < 1$ if for any τ with $1 - \epsilon \leq \tau \leq 1 + \epsilon$, s_τ and s'_τ intersect if and only if s and s' intersect, where s_τ (s'_τ) denotes the scaling of s (s') by τ.*

A collection of objects $\mathcal{S}$ is *ϵ-distant* (*ϵ-separated*) if the objects of $\mathcal{S}$ are pairwise ϵ-distant (ϵ-separated). Observe that any ϵ-distant (ϵ-separated) collection of objects is also ϵ'-distant (ϵ'-separated) for any $0 \leq \epsilon' \leq \epsilon$.

Lemma 4.1.6 *A collection of objects $\mathcal{S}$ is ϵ-distant for some $\epsilon > 0$ if and only if it is ϵ'-separated for some $\epsilon' > 0$.*

Proof: Suppose that $\mathcal{S}$ is ϵ-distant for some $\epsilon > 0$. As the scaling points of the objects in $\mathcal{S}$ are fixed, one can scale each object $s \in \mathcal{S}$ by some τ_s to an object s' such that any point of $\text{bd}(s')$ is within distance ϵ of $\text{bd}(s)$. Let $\epsilon' > 0$ be any number such that $1 - \epsilon' \leq \tau_s \leq 1 + \epsilon'$ for all $s \in \mathcal{S}$. Then $\mathcal{S}$ is ϵ'-separated.

Suppose that $\mathcal{S}$ is ϵ'-separated for some $\epsilon' > 0$. Let d be the smallest distance between the scaling point of s and the boundary of s for any $s \in \mathcal{S}$. Clearly $d > 0$, because the definition of scalable ensures that the scaling point of an object cannot lie on the object boundary. Then, when scaling an object $s \in \mathcal{S}$ by τ_s with $1 - \epsilon' \leq \tau_s \leq 1 + \epsilon'$ to an object s', the distance between $\text{bd}(s)$ and $\text{bd}(s')$ is at least $|\tau_s - 1| \cdot d$. Hence $\mathcal{S}$ is ϵ-distant with $\epsilon \geq \epsilon' \cdot d$. □

In spite of Lemma 4.1.6, we think of ϵ-separated as being a slightly more general notion, since the property of being ϵ-separated is invariant under a scaling of the space.

We now show that any intersection graph of scalable objects has an ϵ-separated representation.

Theorem 4.1.7 *For a family $\mathcal{A}$ of closed scalable objects, any $\mathcal{A}$-intersection graph has an ϵ-separated representation for some $\epsilon > 0$.*

Proof: Let G be an $\mathcal{A}$-intersection graph and $\mathcal{S}$ any representation of G. We prove that $\mathcal{S}$ can be turned into an ϵ-separated representation of G.

For any $u, v \in \mathcal{S}$, let δ_{uv} be maximal such that u and v are δ_{uv}-separated. Let δ be the smallest of the nonzero δ_{uv}, or 1 if all δ_{uv} are zero. Because the objects are scalable and closed, for any $s \in \mathcal{S}$ and any τ_s with $1 \leq \tau_s \leq 1+\delta$, we can scale s by τ_s and the resulting set $\mathcal{S}'$ still induces G. Choose α and ϵ such that $0 < \alpha < \delta$, $0 < \epsilon < 1$, $1 \leq (1-\epsilon)\cdot(1+\alpha)$, and $(1+\epsilon)\cdot(1+\alpha) \leq 1+\delta$, for instance $\alpha = \frac{1}{2}\delta$ and $\epsilon = \frac{1}{4}\delta$. Scale any object in $\mathcal{S}$ by $(1+\alpha)$ and denote the resulting set by $\mathcal{S}'$. By the choice of α and ϵ, $\mathcal{S}'$ is ϵ-separated. Furthermore, $\mathcal{S}'$ still induces G.

Finally, scale the space by $1/(1+\alpha)$. Then the objects of $\mathcal{S}'$ regain their original size, i.e. they are translates of the objects in $\mathcal{S}$. Hence the resulting set $\mathcal{S}''$ contains *only translated copies* of members of $\mathcal{A}$. As separation is invariant under a scaling of the space, $\mathcal{S}''$ is an ϵ-separated representation of G. □

The same theorem holds for open scalable objects.

Theorem 4.1.8 *Given a family $\mathcal{A}$ of open scalable objects, any $\mathcal{A}$-intersection graph has an ϵ-separated representation for some $\epsilon > 0$.*

Proof: The proof is essentially the same as the proof of the previous theorem, except now we are free to scale any $s \in \mathcal{S}$ by τ_s with $1 - \delta \leq \tau_s \leq 1$. Choose α and ϵ such that $0 < \alpha < \delta$, $0 < \epsilon < 1$, $1 - \delta \leq (1-\epsilon)\cdot(1-\alpha)$, and $(1+\epsilon)\cdot(1-\alpha) \leq 1$, for instance $\alpha = \frac{1}{2}\delta$ and $\epsilon = \frac{1}{2}\delta$. Scale the objects by $(1-\alpha)$ and the space by $1/(1-\alpha)$. The resulting collection of objects is ϵ-separated and is a representation of the graph. □

By Lemma 4.1.6, these two theorems imply the following corollary.

Corollary 4.1.9 *Given a family $\mathcal{A}$ of closed or of open scalable objects, any $\mathcal{A}$-intersection graph has an ϵ-distant representation for some $\epsilon > 0$.*

This fact is useful when proving the following corollary.

Theorem 4.1.10 *Let $\mathcal{A}$ be any family of closed scalable objects and $\mathcal{A}' = \{\text{int}(s) \mid s \in \mathcal{A}\}$ the family of their interiors. Then the class of $\mathcal{A}$-intersection graphs equals the class of $\mathcal{A}'$-intersection graphs.*

Proof: Given an $\mathcal{A}$-intersection graph G, let $\mathcal{S}$ be an ϵ-distant representation of G for some $\epsilon > 0$. Such a representation exists by Corollary 4.1.9. But then $\mathcal{S}' = \{\text{int}(s) \mid s \in \mathcal{S}\}$ also induces G. Moreover, $\mathcal{S}'$ uses only translated copies of members of $\mathcal{A}'$. Hence G is an $\mathcal{A}'$-intersection graph. The reverse relation is proved similarly. □

This implies for instance that the class of closed disk (square, triangle, ...) graphs equals the class of open disk (square, triangle, ...) graphs.

Because of the equivalence of open and closed graph classes, we will focus only on closed scalable objects from now on. Theorem 4.1.10 guarantees that the results translate to open scalable objects.

4.2 Finite Representation

With the ϵ-separated and ϵ-distant representations we know to exist now, we can prove that intersection graphs of closed scalable objects have a representation using rationals.

Assume that any object s contains a point d_s, its *distinguished point*. We show that the coordinates of this point can always be rational.

Theorem 4.2.1 *For a family $\mathcal{A}$ of closed scalable objects, any $\mathcal{A}$-intersection graph has a representation $\mathcal{S}$ such that the distinguished point of each object in $\mathcal{S}$ has rational coordinates.*

Proof: Let G be any $\mathcal{A}$-intersection graph and $\mathcal{S}'$ an ϵ-distant representation for G for some $\epsilon > 0$, which exists by Corollary 4.1.9. Because the rationals are dense in the reals, there exists for any $s \in \mathcal{S}'$ a vector $\vec{a}_s$ with $\|\vec{a}_s\| \leq \frac{1}{2}\epsilon$ such that $d_s + \vec{a}_s$ has rational coordinates.

Translate each $s \in \mathcal{S}'$ by $\vec{a}_s$ and let $\mathcal{S}$ be the resulting set of objects. Clearly, the distinguished point of each object in $\mathcal{S}$ has rational coordinates. As $\mathcal{S}'$ is ϵ-distant, it follows from the choice of the $\vec{a}_s$ that $\mathcal{S}$ still induces G. Moreover, $\mathcal{S}$ is $\frac{1}{2}\epsilon$-distant. □

A similar idea applied in the context of convex objects may be found in Czyzowicz et al. [73].

Besides having rational coordinates for the distinguished point of an object, we would like the objects in a representation to have rational size as well. This requires a precise definition of the size of an object.

Definition 4.2.2 *Associate with any object s two distinct points (the* size points *of s). Then the* size *of s is the distance between its two size points.*

Although it seems more natural to use the volume of the object here, this is much harder to work with and the volume might be infinite. Furthermore, this definition of object size captures the way many objects are specified. For instance, a disk is specified by its radius (the distance between the disk center and a point on the boundary) and a square by its side length (the distance between two corners).

The following theorem follows straightforwardly from Theorem 4.2.1.

Theorem 4.2.3 *Let $\mathcal{A}$ be a family of closed scalable objects, each of rational size. Then any $\mathcal{A}$-intersection graph has a representation $\mathcal{S}$ such that the* *distinguished point of each object in $\mathcal{S}$ has rational coordinates and all objects in $\mathcal{S}$ have rational size.*

For families containing objects of nonrational size, one needs to be more careful. We restrict the attention to families that are complete.

Definition 4.2.4 *A family $\mathcal{A}$ of scalable objects is* complete *if for any $s \in \mathcal{A}$ and for any $\tau > 0$ the scaling of s by τ is also in $\mathcal{A}$.*

The family of all disks or of all squares are good examples of complete families. We can now prove the following result.

Theorem 4.2.5 *Given a complete family $\mathcal{A}$ of closed scalable objects, any $\mathcal{A}$-intersection graph has a representation $\mathcal{S}$ such that the distinguished point of each object in $\mathcal{S}$ has rational coordinates and all objects in $\mathcal{S}$ have rational size.*

Proof: Let G be any $\mathcal{A}$-intersection graph and $\mathcal{S}'$ an ϵ-separated representation for G for some $\epsilon > 0$, which exists by Theorem 4.1.7. For any object $s \in \mathcal{S}'$ of size z_s, there exists some τ_s with $1 \leq \tau_s \leq 1 + \frac{\frac{1}{2}\epsilon}{1+\frac{1}{2}\epsilon}$ such that $z_s \cdot \tau_s$ is rational.

Scale each $s \in \mathcal{S}'$ by τ_s and let $\mathcal{S}$ be the resulting set of objects. Clearly, each object in $\mathcal{S}$ has rational size. As $\mathcal{S}'$ is ϵ-separated, it follows from the choice of the τ_s that $\mathcal{S}$ still induces G. Moreover, $\mathcal{S}$ is $\frac{1}{2}\epsilon$-separated.

Let $\mathcal{A}'$ be the family of objects in $\mathcal{A}$ having rational size. By the preceding argument, G is an $\mathcal{A}'$-intersection graph. The theorem now follows from Theorem 4.2.3. □

This implies for instance that a representation for any (unit) disk graph can be specified using only rationals. We should note that by Theorem 4.1.10 the above results also hold for families of open scalable objects. Furthermore, by scaling the space appropriately, we can replace 'rational' with 'integer' in the statement of Theorem 4.2.5.

4.3 Polynomial Representation and Separation

We proved that intersection graphs of scalable objects have a rational representation, i.e. a representation where both the coordinates of the distinguished point and (if the family is complete) the size of each object is rational. We now consider what happens when we require these rationals to have polynomial size, bringing the problem closer to what we want for the recognition problem.

A *q-bit rational* is a rational number where both the integer and fractional part of the rational are q-bit integers, i.e. there exist q-bit integers a and b such that the rational is $a + b/2^q$.

Definition 4.3.1 *For a family $\mathcal{A}$ of scalable objects, an $\mathcal{A}$-intersection graph G has a q-*representation *for some $q \geq 0$ if G has a representation $\mathcal{S}$ such that* *the distinguished point of each object in $\mathcal{S}$ has q-bit rational coordinates and each object has q-bit rational size.*

The class of $\mathcal{A}$-intersection graphs has a q-representation *for some function $q : \mathbb{N} \to \mathbb{N}$ if for each $n \in \mathbb{N}$ any n-vertex $\mathcal{A}$-intersection graph has a $q(n)$-representation. The class has a* polynomial representation *if it has a q-representation for some polynomially bounded function q.*

It is widely believed that (unit) disk graphs have a polynomial representation. As far as we know however, no function q (polynomial or exponential) is known for which (unit) disk graphs have a q-representation. Recall however that (unit) square graphs have a polynomial representation (see Table 4.1).

Definition 4.3.2 *For a family $\mathcal{A}$ of scalable objects, an $\mathcal{A}$-intersection graph G has a* q-separated (q-distant) representation *for $q \geq 0$ if G has a representation such that $\mathcal{S}$ is ϵ-separated (ϵ-distant) for some q-bit rational $\epsilon > 0$.*

The class of $\mathcal{A}$-intersection graphs has a q-separated (q-distant) representation *for some function $q : \mathbb{N} \to \mathbb{N}$ if for each $n \in \mathbb{N}$ any n-vertex $\mathcal{A}$-intersection graph has a $q(n)$-separated ($q(n)$-distant) representation. The class has a* polynomial separation *if it has a q-separated representation for some polynomially bounded function q.*

We would like to know which classes of geometric intersection graphs have a polynomial representation. In particular, we are interested in (unit) disk graphs and (unit) square graphs. To gain better insight into this question, we show that the existence of a q-representation implies the existence of a q'-separated representation for these graph classes. We prove that the converse holds as well.

4.3.1 From Representation to Separation

Throughout, we assume that the scaling point and the distinguished point of a disk or square coincide with its center. The size of a disk is its radius and the size of a square is its side length.

Theorem 4.3.3 *If a (unit) disk graphs has a q-representation for some $q \geq 0$, then it has a $(4q+6)$-separated representation.*

Proof: Let G be a (unit) disk graph and $\mathcal{S}$ a q-representation for G. Scale the space by 2^q such that all numbers of $\mathcal{S}$ are $2q$-bit integers. We claim that any two nontouching disks in $\mathcal{S}$ are δ-separated, where $\delta = 1/(2^{4q+4})$.

For any $u \in \mathcal{S}$, let $c_u = (x_u, y_u)$ denote the center and r_u the radius of disk u. Suppose two disks u and v intersect but do not touch. Then $\|c_u - c_v\| < r_u + r_v$ and thus $\|c_u - c_v\|^2 < (r_u + r_v)^2$. As $\|c_u - c_v\|^2 = (x_u - x_v)^2 + (y_u - y_v)^2$ and $(r_u + r_v)^2$ are both integral, $\|c_u - c_v\|^2 \leq (r_u + r_v)^2 - 1$. Hence

$$\|c_u - c_v\| \quad \leq \quad \sqrt{(r_u + r_v)^2 - 1}$$

$$\begin{aligned} &= (r_u + r_v) \cdot \sqrt{1 - \frac{1}{(r_u + r_v)^2}} \\ &\leq (r_u + r_v) \cdot \left(1 - \frac{1}{4(r_u + r_v)^2}\right) \\ &\leq (r_u + r_v) \cdot (1 - \delta) \end{aligned}$$

and thus u and v are δ-separated.

Suppose two disks u and v do not intersect. Following a similar argument,

$$\begin{aligned} \|c_u - c_v\| &\geq \sqrt{(r_u + r_v)^2 + 1} \\ &= (r_u + r_v) \cdot \sqrt{1 + \frac{1}{(r_u + r_v)^2}} \\ &\geq (r_u + r_v) \cdot \left(1 + \frac{1}{4(r_u + r_v)^2}\right) \\ &\geq (r_u + r_v) \cdot (1 + \delta) \end{aligned}$$

and thus u and v are δ-separated. It follows from the proof of Theorem 4.1.7 that $\mathcal{S}$ can be transformed (using only translations) into a $\frac{1}{4}\delta$-separated representation. This representation is q'-separated for $q' = 4q + 6$. □

Corollary 4.3.4 *If the class of (unit) disk graphs has a q-representation for some function $q : \mathbb{N} \to \mathbb{N}$, then it has a q'-separated representation, where $q'(n) = 4q(n) + 6$. In particular, polynomial representation implies polynomial separation.*

Theorem 4.3.5 *If a (unit) square graphs has a q-representation for some $q \geq 0$, then it has a $(2q + 3)$-separated representation.*

Proof: Let G be a (unit) square graph and $\mathcal{S}$ a q-representation for G. Scale the space by 2^q such that all numbers of $\mathcal{S}$ are $2q$-bit integers. Since the squares have side length at most $2^{2q} - 1$, any two nontouching squares in $\mathcal{S}$ are δ-separated, where $\delta = 1/(2^{2q+1})$. It follows from the proof of Theorem 4.1.7 that $\mathcal{S}$ can be transformed (using only translations) into a $\frac{1}{4}\delta$-separated representation. This representation is q'-separated for $q' = 2q + 3$. □

Corollary 4.3.6 *If the class of (unit) square graphs has a q-representation for some function $q : \mathbb{N} \to \mathbb{N}$, then it has a q'-separated representation, where $q'(n) = 2q(n) + 3$. In particular, polynomial representation implies polynomial separation.*

Recall that for unit square graphs a q-representation exists where $q(n) = \mathrm{O}(\log n)$ [70, 73] and for square graphs one exists where $q(n) = \mathrm{O}(n^4)$ [261].

Corollary 4.3.7 *Unit square graphs have a q'-separated representation where $q'(n) = \mathrm{O}(\log n)$. Square graphs have one where $q'(n) = \mathrm{O}(n^4)$.*

We believe that results similar to Theorem 4.3.5 can be proved for intersection graphs of other scalable objects. In particular, we conjecture that similar techniques apply to intersection graphs of (unit) regular hexagons.

Finally, observe that for the results in this section it does not matter if the disks or squares are open or closed.

4.3.2 From Separation to Representation

The above theorems were quite specific to the object type. We can prove that the converse holds in a more general setting. In the following, let z_s denote the size of an object s. Moreover, for a set of objects $\mathcal{S} = \{s_1, \ldots, s_n\}$, we assume that $z_{s_1} \leq \cdots \leq z_{s_n}$. We will sometimes use z_i as a shorthand for z_{s_i}.

Lemma 4.3.8 *Let $\mathcal{A}$ be a family of closed scalable objects and let $q \geq 0$. If an n-vertex $\mathcal{A}$-intersection graph G has a q-distant representation $\mathcal{S} = \{s_1, \ldots, s_n\}$ such that $z_1 = 1$, z_n is bounded by a q-bit rational, and the radius of the smallest enclosing sphere of any object is at most 2^q, then G has a q'-representation, where $q' = q + \lceil \log n \rceil + 2$.*

Proof: Suppose that $\mathcal{S}$ is ϵ-distant for some q-bit rational $0 < \epsilon < 1$. We may assume that $\epsilon = 2^{-q}$. Scale each object $s_i \in \mathcal{S}$ to s'_i such that $z'_i = z_i - \tau_i$ for the smallest value of $0 \leq \tau_i \leq \frac{1}{2}\epsilon$ for which z'_i is a multiple of $\frac{1}{2}\epsilon$. Let $\mathcal{S}' = \{s'_1, \ldots, s'_n\}$ be the resulting set of objects. By the choice of the τ_i, $1 = z'_1 \leq \cdots \leq z'_n$ and z'_n is bounded by a q-bit rational. Since each z'_i is a multiple of $\frac{1}{2}\epsilon$ and ϵ is a q-bit rational, each z'_i is a $(q+1)$-bit rational. Moreover, $\mathcal{S}'$ is an $\frac{1}{2}\epsilon$-distant representation of G.

By translating the objects if necessary, we may assume that there is no hyperplane h which intersects no objects of $\mathcal{S}'$ such that any two objects, one on each side of h, are ϵ'-distant for some $\epsilon' > \frac{1}{2}\epsilon$. This still is an $\frac{1}{2}\epsilon$-distant representation. As the radius of the smallest enclosing sphere of any object is (still) at most 2^q, all objects are contained in a box with sides of length at most

$$2n \cdot 2^q + (n-1)\epsilon < 2n(2^q + 1) < 2^{q + \lceil \log n \rceil + 2}.$$

Hence the integer part of the distinguished point of any object in $\mathcal{S}'$ needs at most $q + \lceil \log n \rceil + 2$ bits. Furthermore, following the proof of Theorem 4.2.1, the fractional part needs at most $q+2$ bits, by translating the objects slightly if necessary. The result is a q'-representation of G with $q' = q + \lceil \log n \rceil + 2$. □

Theorem 4.3.9 *Let $\mathcal{A}$ be a family of closed scalable objects, let $q \geq 0$, and let $0 < \delta < 1$. If an n-vertex $\mathcal{A}$-intersection graph G has a q-separated representation $\mathcal{S} = \{s_1, \ldots, s_n\}$ such that $z_1 = 1$, z_n is bounded by a q-bit rational, the radius of the smallest enclosing sphere of any object is at most 2^q, and all points within distance δ of the scaling point of any $s \in \mathcal{S}$ belong to s, then G has a q'-representation, where $q' = q + \lceil \log n \rceil + \lceil \log 1/\delta \rceil + 2$.*

Proof: Since all points within distance δ of the scaling point of s belong to s, any n-vertex $\mathcal{A}$-intersection graph has a q''-distant representation with $q'' = q + \lceil \log 1/\delta \rceil$ by the proof of Lemma 4.1.6. The theorem then follows immediately from Lemma 4.3.8. □

If G is a unit disk graph or a unit square graph with a q-separated representation for some $q \geq 0$, then G clearly has a $(q + \lceil \log n \rceil + 3)$-representation by Theorem 4.3.9.

Corollary 4.3.10 *If the class of unit disk graphs or of unit square graphs has a q-separated representation for some $q : \mathbb{N} \to \mathbb{N}$, then it has a q'-representation, where $q'(n) = q(n) + \lceil \log n \rceil + 3$. In particular, polynomial separation implies polynomial representation.*

Note that we can replace unit square here with any unit regular polygon.

For disk graphs or square graphs, a result as Corollary 4.3.10 is not immediate, as we have no bound (yet) on the size of the disks or squares in an ϵ-separated representation. (This size is constant, 1/2 and 1 respectively, in the unit case.) We prove such a bound for arbitrary disks and squares below.

Lemma 4.3.11 *Let G be an n-vertex disk graph with an ϵ-separated representation for some $\epsilon > 0$ for which $1/\epsilon$ is integer. Then G has a $\frac{1}{2}\epsilon$-separated representation in which all radii are at least 1 and at most $(256n/\epsilon)^{3n+1}$.*

The proof of this lemma uses trigonometry and linear programming and is quite long. It is given in [262]. Here we give a simpler inductive proof. Recall that the size of disk u is its radius and the scaling point c_u is its center.

Lemma 4.3.12 *Let G be an n-vertex disk graph with an ϵ-separated representation $\mathcal{S} = \{s_1, \ldots, s_n\}$ for some $0 < \epsilon < 1$. Then G has a $\frac{1}{2}\epsilon$-separated representation $\widetilde{\mathcal{S}} = \{\tilde{s}_1, \ldots, \tilde{s}_n\}$ such that $1 = \tilde{z}_1 \leq \cdots \leq \tilde{z}_n \leq (8n/\epsilon)^{n-1}$.*

Proof: We apply induction on n. If $n = 1$, the lemma is trivial. Suppose that $n > 1$. By scaling the space if necessary, we may assume that $z_n = (8n/\epsilon)^{n-1}$. Recall that separation is invariant under a scaling of the space.

If $z_i/z_{i-1} \leq (8n/\epsilon)$ for any $2 \leq i \leq n$, then $z_1 \geq 1$ and the lemma holds (by scaling the space if necessary). So let i be the largest index for which $z_i/z_{i-1} > (8n/\epsilon)$. Let $\mathcal{S}' = \{s_1, \ldots, s_{i-1}\}$ and $\mathcal{S}'' = \{s_i, \ldots, s_n\}$. Note that by the choice of i, $z_i \geq (8n/\epsilon)^{i-1} \geq 1$.

For each $u \in \mathcal{S}'$, let B_u denote the disk with radius $\frac{1}{4}\epsilon\, z_i$ centered at c_u. Furthermore, let

$$N_u = \{v \in \mathcal{S}'' \mid u \cap v \neq \emptyset\}.$$

Call $u, t \in \mathcal{S}'$ *equivalent* if $N_u = N_t$. Consider any equivalence class $\mathcal{E}$. By induction, $G[\mathcal{E}]$ has an $\frac{1}{2}\epsilon$-separated representation $\widetilde{\mathcal{E}} = \{\tilde{e}_1, \ldots, \tilde{e}_k\}$ with

$$1 = z_{\tilde{e}_1} \leq \cdots \leq z_{\tilde{e}_k} \leq (8k/\epsilon)^{k-1} \leq (8n/\epsilon)^{k-1} \leq (8n/\epsilon)^{n-1}.$$

Let $B_{\mathcal{E}} = B_u$ for some (fixed) $u \in \mathcal{E}$. Similar to Lemma 4.3.8, we can assume, by translating if necessary, that $\widetilde{\mathcal{E}}$ is contained in a disk of radius at most

$$\begin{aligned} k\,(8n/\epsilon)^{k-1} + (k-1)\epsilon\,(8n/\epsilon)^{k-1} &\leq (2k-1)\,(8n/\epsilon)^{k-1} \\ &\leq (2k-1)\,(8n/\epsilon)^{i-2} \\ &\leq (2k-1)\tfrac{1}{8n}\,\epsilon\, z_i. \end{aligned}$$

Hence we may assume that $\widetilde{\mathcal{E}}$ is contained in $B_{\mathcal{E}}$. Now replace $\mathcal{E}$ by $\widetilde{\mathcal{E}}$ and let $\widetilde{\mathcal{S}}$ denote the resulting set of disks.

We show that $\widetilde{\mathcal{S}}$ is a representation of G. We first give an auxiliary property. Let v_τ be the scaling of v by τ. We claim that for each $u \in \mathcal{S}'$,

(1) B_u is contained in $v_{1-\frac{1}{2}\epsilon}$ for each $v \in N_u$;

(2) B_u is disjoint from $v_{1+\frac{1}{2}\epsilon}$ for each $v \in \mathcal{S}'' - N_u$.

Suppose that $v \in N_u$. Since $v \in \mathcal{S}''$, $z_v \geq z_i$, and thus $z_u \leq \frac{1}{8n}\epsilon\, z_v$. Because $\mathcal{S}$ is ϵ-separated, $c_u \in v_{1-\frac{3}{4}\epsilon}$. But then $B_u \subseteq v_{1-\frac{1}{2}\epsilon}$. This proves (1).

Suppose that $v \in \mathcal{S}'' - N_u$. Because $\mathcal{S}$ is ϵ-separated, $c_u \notin v_{1+\epsilon}$. As $z_i \leq z_v$, B_u and $v_{1+\frac{1}{2}\epsilon}$ are disjoint. This proves (2).

Now by (1) and the definition of equivalent, $u \in \mathcal{S}$ and $v \in \mathcal{S}''$ intersect if and only if $\tilde{u}$ and $\tilde{v}$ intersect. Moreover, by induction, $u, t \in \mathcal{S}'$ in the same equivalence class intersect if and only if $\tilde{u}$ and $\tilde{t}$ intersect. We show below that if $u, t \in \mathcal{S}'$ are not in the same equivalence class, then B_u and B_t are disjoint. Since $u \subseteq B_u$ and $t \subseteq B_t$, u and t are disjoint. Moreover, by the choice of the $B_{\mathcal{E}}$, $\tilde{u}$ and $\tilde{t}$ are disjoint. This implies that $\widetilde{\mathcal{S}}$ is a representation of G.

We in fact prove a stronger statement, namely that if $N_u \neq N_t$ for $u, t \in \mathcal{S}'$, then B_u and B_t are disjoint and 1-separated. So assume that $N_u \neq N_t$ and w.l.o.g. that $|N_u - N_t| > 0$. Let $v \in N_u - N_t$. Then by (1), B_u is contained in $v_{1-\frac{1}{2}\epsilon}$. By (2), B_t is disjoint from $v_{1+\frac{1}{2}\epsilon}$. Hence B_u and B_t have distance at least $\epsilon\, z_v$. Because $\frac{1}{4}\epsilon\, z_i \leq \frac{1}{4}\epsilon\, z_v$, B_u and B_t are disjoint and 1-separated.

Hence $\widetilde{\mathcal{S}} = \{\tilde{s}_1, \ldots, \tilde{s}_n\}$ is a representation of G. By construction, $1 = \tilde{z}_1 \leq \cdots \leq \tilde{z}_n \leq (8n/\epsilon)^{n-1}$. From the construction of $\widetilde{\mathcal{S}}$, (1), and (2), $\widetilde{\mathcal{S}}$ is $\frac{1}{2}\epsilon$-separated. The lemma follows. □

Theorem 4.3.13 *If an n-vertex disk graph has a q-separated representation representation for some $q \geq 0$, then it has a q'-representation, where $q' = n(q + \lceil \log n \rceil + 3)$.*

Proof: Let G be a disk graph with an ϵ-separated representation for some q-bit rational $0 < \epsilon < 1$. Apply Lemma 4.3.12 to obtain a $\frac{1}{2}\epsilon$-separated representation $\widetilde{\mathcal{S}} = \{\tilde{s}_1, \ldots, \tilde{s}_n\}$ such that $1 = \tilde{z}_1 \leq \cdots \leq \tilde{z}_n \leq (8n/\epsilon)^{n-1}$. Note that $(8n/\epsilon)^{n-1} < (8n2^q)^{n-1}$ is an $((n-1)(q + \lceil \log n \rceil + 3))$-bit rational. Following Theorem 4.3.9, G has a q'-representation, where

$$q' = (n-1)(q + \lceil \log n \rceil + 3) + \lceil \log n \rceil + 3 \leq n(q + \lceil \log n \rceil + 3).$$

The theorem follows. □

Corollary 4.3.14 *If the class of disk graphs has a q-separated representation for some* $q : \mathbb{N} \to \mathbb{N}$*, then it has a* q'*-representation, where* $q'(n) = n(q(n) + \lceil \log n \rceil + 3)$*. In particular, polynomial separation implies polynomial representation.*

The same results hold, mutatis mutandis, for square graphs. Moreover, the theorems apply both to open and closed disks or squares.

We can now prove the following result.

Theorem 4.3.15 *The class of intersection graphs of closed (unit) disks has a polynomial representation if and only if the class of intersection graphs of open (unit) disks has a polynomial representation.*

Proof: Suppose that the class of intersection graphs of closed (unit) disks has a polynomial representation. By Corollary 4.3.4, it has a polynomial separation. Hence the class of intersection graphs of open (unit) disks has a polynomial separation. But then Corollary 4.3.14 implies that it has a polynomial representation. The reverse relation follows in a similar manner. □

This theorem also holds mutatis mutandis for (unit) squares. Hence when looking for a polynomial representation of the class of (unit) disk or of (unit) square graphs, it does not matter whether we consider disks or squares that are open or closed. This strengthens Theorem 4.1.10.

As a last observation, note that the proofs of this section are constructive. Hence there is an algorithm to transform a q-separated representation to a q'-representation. Furthermore, it is easy to see that the above corollaries imply the existence of a recognition algorithm for (unit) disk graphs and (unit) square graphs, provided that q-separated representations exist for finite q.

By $O^*(\cdot)$ we mean that polynomial terms are ignored.

Theorem 4.3.16 *If the class of disk graphs or of square graphs has a q-separated representation for some function* $q : \mathbb{N} \to \mathbb{N}$*, then it can be recognized in* $O^*(2^{6n^2(q(n)+\lceil \log n \rceil+3)})$ *time. In the unit case, the time bound improves to* $O^*(2^{6n(q(n)+\lceil \log n \rceil+3)})$.

Proof: We only consider (unit) disk graphs. The case for (unit) square graphs is similar. By Corollary 4.3.14, disk graphs have a q'-representation, where $q'(n) = n(q(n) + \lceil \log n \rceil + 3)$. Hence any n-vertex disk graph G has a representation by $3n\ q'(n)$-bit rationals. We now enumerate all possible representations and verify whether one induces G. The bound in the case of unit disk graphs follows from Corollary 4.3.10. □

Part II

Approximating Optimization Problems on Geometric Intersection Graphs

Overview

Wireless communication networks increase their influence on human interaction on a daily basis. Many people do away with their fixed landline phone and only use a cell phone, wireless Internet allows one to check e-mail everywhere, etcetera. Making such networks work has led to a lot of new problems and challenges, ranging from practical questions to theoretical conundrums.

The purpose of this part of the thesis is to advance insight into some of these theoretical problems. We address them by considering a model that is commonly used for wireless networks, the geometric intersection graph model, and in particular the (unit) disk graph model. We then study several optimization problems on wireless networks, which translate to well-known graph optimization problems (such as Maximum Independent Set and Minimum Connected Dominating Set) on geometric intersection graphs. As geometric intersection graphs have a geometric representation, we can show that these problems can be better solved or approximated on such graphs than on general graphs.

Problems

We consider various optimization problems on graphs that are relevant to geometric intersection graph models and specifically to (unit) disk graphs models of wireless communication networks.

Definition II.1 *Let G be a graph. A set $S \subseteq V(G)$ is an* independent set *if there are no $u, v \in S$ such that $(u, v) \in E(G)$. A set $S \subseteq V(G)$ is a* vertex cover *if for each $(u, v) \in E(G)$ it holds that $u \in S$ or $v \in S$.*

Observe that an independent set is the complement of a vertex cover (and vice versa) [115]. Furthermore, we are usually looking for a maximum independent set and a minimum vertex cover. An independent set is *maximum* if there is no independent set of greater cardinality. A vertex cover is *minimum* if there is no vertex cover of smaller cardinality. In the context of wireless communication networks, an independent set of a (unit) disk graph can be seen as a set of nodes that can transmit simultaneously without signal interferences. Vertex covers are mostly interesting from a theoretical point of view in this context.

Definition II.2 *Let G be a graph. A set $S \subseteq V(G)$ is a* dominating set *if for each vertex $v \in V$, $v \in S$ or there is a vertex $u \in S$ for which $(u, v) \in E(G)$.*

Definition II.3 *Let G be a graph. A set $S \subseteq V(G)$ is a* connected dominating set *if S is a dominating set and the subgraph of G induced by S is connected.*

A dominating set in a wireless communication network can be seen as a set of emergency transmitters capable of reaching every node in the network, or as central nodes in node clusters. A connected dominating set can be used as a backbone for easier and faster communications. The problem is to find a minimum (connected) dominating set, i.e. we look for (connected) dominating sets of minimum cardinality.

Previous Work

All problems mentioned above are NP-complete on general graphs (see Garey and Johnson [115]). Since (unit) disk graphs are a restricted class of graphs with a nice geometric interpretation, one might hope that these problems are better solvable. Unfortunately, these problems are NP-complete on unit disk graphs and other classes of intersection graphs of planar objects as well [267, 17, 67]. The NP-hardness holds even if the degree is at most 3 and (except for Maximum Independent Set and Minimum Vertex Cover) the graph is bipartite [65]. Therefore most research has focused on approximation algorithms.

We survey the previous work in this area, with strong emphasis on geometric intersection graphs and (unit) disk graphs in particular. Work specific to a particular chapter will be surveyed there. A detailed survey of available approximation algorithms and inapproximability results on general graphs can be found (for example) in the compendium of Ausiello et al. [19].

Maximum Independent Set

On general n-vertex graphs, Maximum Independent Set has a polynomial time $O(n/\log^2 n)$-approximation algorithm [39] and is not approximable within $O(n^{1-\epsilon})$ for any $\epsilon > 0$, unless NP=ZPP [142]. On geometric intersection graphs however, it is easy to give a constant-factor approximation algorithm. Many geometric intersection graphs have no $K_{1,m}$ induced subgraph for some $m > 1$. Unit disk graphs have no $K_{1,6}$ induced subgraph for example. In general, Maximum Independent Set has a polynomial-time $m/2$-approximation algorithm on graphs with no $K_{1,m}$ induced subgraph, even in the weighted case [137, 276, 21, 31]. This immediately gives a 3-approximation for Maximum (Weighted) Independent Set on unit disk graphs (see for example Yu, Kouvelis, and Luo [277]).

A similar result can be obtained by a greedy algorithm based on the representation of the unit disk graph. Marathe et al. [201] showed that greedily choosing the vertex corresponding to the leftmost disk gives a 3-approximation algorithm for the unweighted case. On general disk graphs, which can have a $K_{1,m}$ induced subgraph for any $m \geq 1$, greedily choosing the vertex with the smallest disk radius results in a 5-approximation algorithm [201].

Agarwal and Mustafa [2] provide a more general approach and consider the intersection graph of a family $\mathcal{S}$ of convex two-dimensional objects. If α is the cardinality of the maximum independent set of this graph, their algorithm returns an independent set of cardinality $(\alpha/(2\log(2n/\alpha)))^{\frac{1}{3}}$ in $\mathrm{O}(n^3 + \tau(\mathcal{S}))$ time, where $\tau(\mathcal{S})$ is the time needed to compute the left- and rightmost point of each object and test which objects intersect. A set of disks clearly is a set of convex two-dimensional objects and hence the algorithm of Agarwal and Mustafa also applies to (unit) disk graphs.

Maximum Independent Set on (unit) disk graphs also has a polynomial-time approximation scheme (ptas) using the so-called *shifting technique.* This is a general technique, independently discovered by Baker [22] and Hochbaum and Maass [150]. The basic idea is the following. A set of regularly spaced separators is used to decompose the problem into smaller, easier solvable subproblems. The solutions of the subproblems are merged to form a solution to the global problem. This is repeated for several placements of the separator set. The best solution over these placements is then selected as an approximation of the optimum. Moving the separator set can be regarded as shifting the set through the problem. Hence the name 'shifting technique'.

Since its discovery, the shifting technique has been used to solve various problems [4, 154, 103, 272]. In the context of (unit) disk graphs, Matsui [206] and Hunt et al. [154] both presented a ptas using the shifting technique, employing different proof ideas. Nieberg, Hurink, and Kern [219] give a (robust) ptas for Maximum Independent Set on unit disk graphs for which no disk representation is given. The ptas extends to graphs of polynomially bounded growth. A *robust algorithm* on unit disk graphs solves the problem correctly for every unit disk graph. For graphs that are not unit disk graphs, the algorithm may either produce the correct output for the problem, or provide a certificate that the input is not a unit disk graph.

Hunt et al. [154] also consider unit disk graphs with a representation where the disk centers are at least λ apart, so-called *λ-precision unit disk graphs.* Using the shifting technique, they give an eptas for Maximum Independent Set on unit disk graphs of constant precision.

Erlebach, Jansen, and Seidel [103] generalize the shifting technique to give a ptas for Maximum Independent Set on general disk graphs, which extends to intersection graphs of fat objects. Li and Wang [192] extend these ideas to several other disk models for wireless networks, such as the ones discussed in Section 3.2.1. Chan [57] presents a ptas for Maximum Independent Set on the intersection graph of a set of fat objects. The scheme uses polynomial space. Under the used definition, a set of disks is fat. Hence the presented scheme is a ptas for Maximum Independent Set on disk graphs. The above schemes extend to the weighted case. Chan and Har-Peled [59] generalize in a different direction and give a ptas for Maximum Independent Set on intersection graphs of pseudo-disks. For the weighted case, they give a constant-factor approximation algorithm.

For rectangle intersection graphs where the rectangles are noncrossing, Agarwal and Mustafa [2] give a constant-factor approximation algorithm. For intersection graphs of rectangles that have unit height, Agarwal, van Kreveld, and Suri [4] and Chan [58] give a ptas. If the rectangles have arbitrary height and are d-dimensional, a $O(\log_k^d n)$-approximation algorithm can be given that runs in $O(n^{O(k)})$ time [164, 32, 58] for any $k \geq 2$. Chalermsook and Chuzhoy [54] recently improved on this by presenting a polynomial-time $O(\log^{d-2} n \log\log n)$-approximation algorithm. This algorithm does not apply to the weighted case. For this case however, Chan and Har-Peled [59] give a $O(\log n / \log\log n)$-approximation algorithm. Chlebík and Chlebíková [65] prove that Maximum Independent Set is APX-hard on intersection graphs of d-dimensional axis-parallel boxes for any $d \geq 3$.

Minimum Vertex Cover

For general n-vertex graphs, Minimum Vertex Cover can be approximated in polynomial time within $2 - \frac{\log\log n}{2\log n}$ [214, 25] or (for dense graphs) within $2 - \frac{2\ln\ln n}{\ln n}$ [138], and cannot be approximated within 1.3606, unless P=NP [87].

On unit disk graphs, one can give a polynomial-time 3/2-approximation algorithm [201]. Malesińska [199] gives a constant-factor approximation algorithm for Minimum Vertex Cover on general disk graphs. More important however is the existence of a ptas. Hunt et al. [154] prove that Minimum Vertex Cover has a ptas on unit disk graphs, again using the shifting technique, and an eptas on constant-precision unit disk graphs. Marx [202] even managed to show that an eptas exists on arbitrary unit disk graphs. Nieberg, Hurink, and Kern [219] give a (robust) ptas for Minimum Vertex Cover on unit disk graphs for which no disk representation is given. The idea behind the scheme also works on graphs of polynomially bounded growth.

Erlebach, Jansen, and Seidel [103] use their multi-level shifting technique to give a ptas for Minimum Vertex Cover on general disk graphs, which extends to intersection graphs of fat objects. Li and Wang [192] give a ptas for other disk models, such as given in Section 3.2.1. It is interesting to note that the schemes also apply to the weighted case, but the eptas does not seem to do so.

Chlebík and Chlebíková [65] demonstrate that Minimum Vertex Cover is APX-hard on intersection graphs of d-dimensional axis-parallel boxes for $d \geq 3$.

Minimum (Connected) Dominating Set

Given any n-vertex graph, Minimum Dominating Set can be approximated within $1 + \ln n$ in polynomial time [156, 197, 66], but this problem has no polynomial-time algorithm achieving ratio $(1 - \epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME$(n^{O(\log\log n)})$ [108]. Minimum Connected Dominating Set has a polynomial time $(3 + \ln n)$-approximation algorithm, and similar to Minimum Dominating Set cannot be approximated within $(1 - \epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME$(n^{O(\log\log n)})$ [132]. Better results can be proved if one assumes that the maximum vertex degree is bounded by Δ, giving a ratio of $1 + \ln\Delta$ and $3 + \ln\Delta$ respectively.

Marathe et al. [201] propose constant-factor approximation algorithms for Minimum (Connected) Dominating Set on graphs without a $K_{1,m}$ induced subgraph, yielding approximation ratios of $m-1$ and $2(m-1)$ respectively. This gives a ratio of 5 and 10 respectively on unit disk graphs. Hunt et al. [154] give a ptas for Minimum Dominating Set on unit disk graphs and an eptas on constant-precision unit disk graphs. Nieberg, Hurink, and Kern [219] give a (robust) ptas for Minimum Dominating Set on unit disk graphs for which no disk representation is given. The idea behind the scheme also works on graphs of polynomially bounded growth.

A ptas for Minimum Connected Dominating Set on unit disk graphs was discovered by Cheng et al. [64]. It was recently improved by Zhang et al. [278]. Zhang et al. also give a ptas for Minimum Connected Dominating Set on three-dimensional unit ball graphs.

Chlebík and Chlebíková [65] show that Minimum (Connected) Dominating Set is APX-hard on intersection graphs of d-dimensional axis-parallel boxes for any $d \geq 3$.

The approximability of the weighted case of Minimum (Connected) Dominating Set on unit disk graphs was a longstanding open problem, until recently, when Ambühl et al. [13] gave a 72- and a 89-approximation algorithm respectively. Huang et al. [153] proposed a polynomial-time $(6+\epsilon)$-approximation algorithm for any (fixed) $\epsilon > 0$ for Minimum-Weight Dominating Set on unit disk graphs. This was subsequently improved upon by Dai and Yu [74], who presented a polynomial-time $(5+\epsilon)$-approximation algorithm for any (fixed) $\epsilon > 0$. Applying the 3.875-approximation algorithm for Node-Weighted Steiner Tree on unit disk graphs by Zou et al. [279], one immediately obtains a $(8.875+\epsilon)$-approximation for Minimum-Weight Connected Dominating Set.

Local (Distributed) Algorithms

Because many of the problems described here are motivated by applications in wireless networks, a large amount of research has focused on distributed algorithms, particularly on so-called *local algorithms*, where the state of a node depends only on the state of nodes at a constant distance.

Kuhn et al. [185] give a local ptas for Maximum Independent Set and Minimum Dominating Set on graphs of polynomially bounded growth using nodes at distance $O(\log^* n)$. Wiese and Kranakis [270] show that on unit disk graphs, a local ptas exists for Maximum Independent Set and Minimum Vertex Cover which only requires constant distance.

For Minimum Connected Dominating Set the first distributed approximation algorithm, attaining a ratio of 8, was given by Wan, Alzoubi, and Frieder [12, 266]. It has message complexity $O(n \log n)$ and time complexity $O(n)$. Czyzowicz et al. [72] presented the first local algorithm for Minimum Connected Dominating Set on unit disk graphs, yielding a $(7.453+\epsilon)$-approximation for any $\epsilon > 0$.

This is only a small portion of the known distributed algorithms for these problems. Since we study centralized algorithms in this thesis, we chose to survey only local algorithms here.

Other Optimization Problems

We briefly survey results on two optimization problems that are frequently studied on geometric intersection graphs, but are not studied in this thesis.

The *maximum clique* problem is to find a largest subset of pairwise connected vertices of a graph. It is the complement of the maximum independent set problem. Interestingly, this problem is polynomial-time solvable on rectangle intersection graphs [155], unit disk graphs [67], and intersection graphs of homothetic triangles [160]. On intersection graphs of two-dimensional convex polygons however, it is NP-hard [23] by reduction from Maximum Independent Set on planar graphs. It is APX-hard on intersection graphs of ellipses of eccentricity $0 < e < 1$ [14]. Intriguingly, the problem is still open on general disk graphs.

Another frequently studied problem on geometric intersection graphs is *Chromatic Number*, where one wants to determine the smallest number of colors needed for which each vertex can be assigned a color such that no two adjacent vertices receive the same color. On unit disk graphs it is NP-hard to decide if k colors are sufficient for any fixed $k \geq 3$ [130]. One can even show that Chromatic Number has no polynomial-time $(4/3)$-approximation algorithm [67]. Chromatic Number has a 3-approximation algorithm on unit disk graphs and a 5-approximation algorithm on disk graphs [222, 201, 199, 130, 101] by a simple greedy strategy. Kim, Kostochka, and Nakprasit [167] extend this to a 3-approximation algorithm on intersection graphs of translated copies of a fixed convex compact set and a 6-approximation algorithm on intersection graphs of homothetic copies of a fixed convex compact set. It is worth noting that these results follow from an appropriate upper bound on the coloring number in terms of the clique number of the intersection graph.

An interesting related problem is the geometric version of the *conflict-free coloring* problem, where a coloring has to be found such that each point in the covered part of the plane is overlapped by an object with a color that is unique among the colors of all objects overlapping this point. See for example [106, 140, 245] and the references therein.

Chapter 5

Algorithms on Unit Disk Graph Decompositions

Decompositions form an important way to optimally solve graph optimization problems. Well known decompositions include branch and tree decompositions. We describe a new graph decomposition, called a *relaxed tree decomposition*, and use it to define the relaxed treewidth of a graph. The relaxed treewidth of a graph is no larger than the (ordinary) treewidth and might be up to a factor of two smaller. Hence it could also be smaller than the branchwidth. We investigate the relation of relaxed treewidth to ordinary treewidth and to the strong treewidth of a graph.

Relaxed tree decompositions and tree decompositions have a different structure. We need new algorithms that take a relaxed tree decomposition as input. We give such algorithms for Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set. The performance of these algorithms matches that of algorithms using just a tree decomposition.

The motivation for studying relaxed tree decompositions is that they arise in a natural way in unit disk graphs. We propose a geometric parameter of unit disks, the so-called *thickness*, which translates into an upper bound on the relaxed treewidth. We use this to give algorithms for Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set on unit disk graphs of bounded thickness. We show that for Minimum Connected Dominating Set one can improve significantly on the analysis in the general case if the input graph is a unit disk graph. By applying noncrossing partitions instead of general partitions we prove that unit disk graphs of bounded thickness have a relaxed tree decomposition with Catalan structure.

We start with a description and analysis of the graph decompositions we use. The definition of thickness follows in Section 5.2. Section 5.3 gives algorithms on general graphs and relaxed tree decompositions. Improved analysis for unit disk graphs of bounded thickness is offered in Section 5.4.

5.1 Graph Decompositions

The algorithms of this chapter use various decompositions of the graph. Although the definitions of tree and branch decompositions are widely known by now, we give them here for completeness.

Definition 5.1.1 ([230]) *A* branch decomposition (T, l) *of a graph* G *is a ternary tree* T *and a bijection* l *between the edges of* $E(G)$ *and the leafs of* T*. Associated with every edge* $e \in E(T)$ *is the* middle set *of* e*, defined as the set of vertices in* $V(G)$ *for which there are two incident edges* e_1, e_2 *such that the leafs* $l(e_1)$ *and* $l(e_2)$ *are in different components of* $T - e$*. The* width *of a branch decomposition is the maximum cardinality of any middle set. The* branchwidth $\mathrm{bw}(G)$ *of a graph* G *is the minimum width of any branch decomposition of* G*.*

Definition 5.1.2 ([229]) *A* tree decomposition (T, X) *of a graph* G *is a tree* T *and a collection of* bags $X_t \subseteq V(G)$*, one for each vertex* $t \in V(T)$*, satisfying three conditions:*

(i). $\bigcup_{t \in V(T)} X_t = V(G)$,

(ii). for all $(u, v) \in E(G)$*, there is a vertex* $t \in V(T)$ *such that* $u, v \in X_t$*, and*

(iii). $X_t \cap X_{t''} \subseteq X_{t'}$ *for any* $t, t'' \in V(T)$ *and for any* $t' \in V(T)$ *on the* t*–*t'' *path in* T*.*

The width *of a tree decomposition is* $\max_{t \in V(T)}\{|X_t|\} - 1$*. The* treewidth $\mathrm{tw}(G)$ *of a graph* G *is the minimum width of any tree decomposition of* G*.*

If the tree of a tree decomposition is in fact a path, we speak of a *path decomposition*. The *pathwidth* $\mathrm{pw}(G)$ of a graph G is the minimum width of any path decomposition of G.

The notions of branch-, tree-, and pathwidth are closely related.

Theorem 5.1.3 ([230, 36]) *For any graph* G*,* $\max\{\mathrm{bw}(G), 2\} \leq \mathrm{tw}(G) + 1 \leq \max\{\lfloor \frac{3}{2} \cdot \mathrm{bw}(G) \rfloor, 2\}$ *and* $\mathrm{tw}(G) \leq \mathrm{pw}(G) \leq \mathrm{tw}(G) \cdot \log |V(G)|$*.*

See Bodlaender [36, 37] for surveys on these graph decompositions.

The problems discussed in this chapter all have algorithms that leverage bounds on the branch-, tree-, or pathwidth. We summarize the running times of the current best results in the following table.

	Branchwidth	Treewidth	Pathwidth
Max. Indep. Set	$2^{\mathrm{bw}(G)\omega/2}$ [89]	$2^{\mathrm{tw}(G)}$ [249]	$2^{\mathrm{pw}(G)}$ [249]
Min. Vertex Cover	$2^{\mathrm{bw}(G)\omega/2}$ [89]	$2^{\mathrm{tw}(G)}$ [249]	$2^{\mathrm{pw}(G)}$ [249]
Min. Dom. Set	$2^{2 \cdot \mathrm{bw}(G)}$ [89]	$2^{2 \cdot \mathrm{tw}(G)}$ [6, 8]	$3^{\mathrm{pw}(G)}$ [6, 8]
Min. Con. D. Set	$\mathrm{bw}(G)^{\mathrm{bw}(G)}$ [89]	$\mathrm{tw}(G)^{\mathrm{tw}(G)}$ [84]	$\mathrm{pw}(G)^{\mathrm{pw}(G)}$ [84]

Here $\omega \leq 2.376$ is the best known matrix multiplication exponent.

For unit disk graphs, it is usually better to consider the following variation of treewidth, which might be interesting in its own right.

Definition 5.1.4 *A* relaxed tree decomposition (T, X) *is a tree* T *and a collection of* bags $X_t \subseteq V(G)$, *one for each vertex* $t \in V(T)$, *satisfying:*

(i). $\bigcup_{t \in V(T)} X_t = V(G)$,

(ii). for all $(u, v) \in E(G)$, *there is a vertex* $t \in V(T)$ *such that* $u, v \in X_t$ *or there is an edge* $(t, t') \in E(T)$ *such that* $u \in X_t$ *and* $v \in X_{t'}$, *and*

(iii). $X_t \cap X_{t''} \subseteq X_{t'}$ *for any* $t, t'' \in V(T)$ *and for any* $t' \in V(T)$ *on the* t–t'' *path in* T.

The width *of a relaxed tree decomposition is* $\max_{t \in V(T)}\{|X_t|\}$. *The* relaxed treewidth $\text{rtw}(G)$ *of a graph* G *is the minimum width of any relaxed tree decomposition of* G.

If the tree of a relaxed tree decomposition is in fact a path, we speak of a *relaxed path decomposition*. The *relaxed pathwidth* $\text{rpw}(G)$ of a graph G is the minimum width of any relaxed path decomposition of G.

The difference between treewidth and relaxed treewidth lies in constraint *(ii)* of the underlying decomposition and the absence or presence of the -1 term in the definition of width (note that $\text{rtw}(T) = \text{tw}(T) = 1$ for any tree T). The treewidth and the relaxed treewidth of a graph are related though.

Proposition 5.1.5 *For any graph* G, $\frac{1}{2}(\text{tw}(G) + 1) \leq \text{rtw}(G) \leq \text{tw}(G) + 1$ *and* $\frac{1}{2}(\text{pw}(G) + 1) \leq \text{rpw}(G) \leq \text{pw}(G) + 1$. *Moreover, these bounds are tight.*

Proof: Let (T, X) be a relaxed tree decomposition of a graph G. For each edge $(t, t') \in E(T)$, replace it by a two-edge path with new vertex tt' and let $X_{tt'} = X_t \cup X_{t'}$. This yields a tree decomposition of G of width at most twice the width of (T, X) minus one. This bound is tight, as $\text{rtw}(K_n) = \lceil n/2 \rceil$ and $\text{tw}(K_n) = n - 1$ for any $n > 0$.

Any tree decomposition of width k is also a relaxed tree decomposition of width $k + 1$. This bound is tight, as exhibited by the following graph. Let $k \geq 2$ and let $U = \{u_1, \ldots, u_{k+1}\}$ and $W = \{w_1, \ldots, w_{k+1}\}$ be disjoint sets of $k + 1$ vertices each. Add edges to make U a clique and for each $w_i \in W$ add edges to all vertices $u_j \in U \backslash \{u_i\}$. That is, for each $i \in \{1, \ldots, k + 1\}$, (u_i, w_i) is a nonedge. Call the resulting graph G.

Clearly, $\text{tw}(G) = k$ and $\text{rtw}(G) \leq k + 1$. Suppose for sake of contradiction that G has a relaxed tree decomposition (T, X) of width at most k. Since $|U| = k + 1$, no bag can contain all vertices of U. As U is a clique, there must be bags X_t and $X_{t'}$ such that $(t, t') \in E(T)$ and $U \subseteq X_t \cup X_{t'}$. Moreover, there must be vertices $u_i, u_j \in U$ such that $u_i \in X_t - X_{t'}$ and $u_j \in X_{t'} - X_t$. This implies that $W - \{w_i, w_j\} \subseteq X_t \cup X_{t'}$. But then $|X_t \cup X_{t'}| \geq 2k$ and thus $X_t \cap X_{t'} = \emptyset$. Assume w.l.o.g. that $|X_t \cap U| \leq |X_{t'} \cap U|$. Then $w_j \in X_t \cup X_{t'}$ and thus $|X_t \cup X_{t'}| \geq 2k + 1$, a contradiction.

The same arguments hold for (relaxed) path decompositions. □

The definition of relaxed tree- and pathwidth is reminiscent of the definition of strong treewidth by Seese [239].

Definition 5.1.6 *A* strong tree decomposition (T, X) *is a relaxed tree decomposition where the bags are pairwise disjoint. The* strong treewidth $\text{stw}(G)$ *of a graph* G *is the minimum width over all strong tree decompositions of* G.

The definitions of the notions of *strong path decomposition* and *strong pathwidth* $\text{spw}(G)$ easily follow.

Proposition 5.1.7 *On any graph* G, $\text{rtw}(G) \leq \text{stw}(G)$, $\text{tw}(G) \leq 2 \cdot \text{stw}(G) - 1$, $\text{rpw}(G) \leq \text{spw}(G)$, *and* $\text{pw}(G) \leq 2 \cdot \text{spw}(G) - 1$.

Proof: The bound $\text{rtw}(G) \leq \text{stw}(G)$ is immediate from Definition 5.1.6. The bound $\text{tw}(G) \leq 2 \cdot \text{stw}(G) - 1$ follows from Proposition 5.1.5 and was also proved by Seese [239]. The inequalities for strong pathwidth follow similarly. □

It is not clear yet what relation exists between the strong treewidth of a graph and its (relaxed) treewidth, although a large gap might exist.

Proposition 5.1.8 *Let* G *be a wheel graph with* $k(k+2)$ *spokes for* $k \geq 2$. *Then* $\text{tw}(G) = \text{rtw}(G) = 3$, *but* $\text{stw}(G) = k + 1$.

Proof: One can easily show that $\text{tw}(G), \text{rtw}(G) \leq 3$. As G has a K_4-minor, $\text{tw}(G) \geq 3$. A cycle has relaxed treewidth 2 and from this $\text{rtw}(G) \geq 3$.

To prove that $\text{stw}(G) = k + 1$, consider any strong tree decomposition (T, X) of G. Let X_h be the bag containing the hub of G. Any other vertex of G must be either in X_h or a bag X_t with $(h, t) \in E(T)$. Consider any edge $(u, v) \in E(G)$ that is not a spoke and for which there is no bag containing both u and v. Then w.l.o.g. $u \in X_h$ and $v \in X_t$ for $t \neq h$. Call (u, v) a *crossing edge* of t. Clearly, any bag $t \neq h$ has at least two (unique) crossing edges. Moreover, any vertex in X_h can be incident with at most two crossing edges. As the hub is not incident with any crossing edge, $|V(T)| \leq |X_h|$. But then the width of (T, X) is greater than k by simple counting. A strong tree decomposition with k bags of $k+1$ vertices each and a central bag X_h with k vertices and the hub has width $k + 1$. Hence $\text{stw}(G) = k + 1$. □

Finally, we give an important property of relaxed tree decompositions, similar to a result by Kloks [168, p. 149] for tree decompositions.

Lemma 5.1.9 *Let* (T, X) *be a relaxed tree decomposition of width* w *of a graph* G. *Then there exists a relaxed tree decomposition* (T', X') *of width* w *of* G *such that any vertex in* T' *has degree at most* 3.

Proof: If $t \in V(T)$ has degree larger than three, split the set of neighbors of t into two sets (N_1 and N_2) of cardinality at least two. Remove t from T and replace it by two vertices t', t'', connected by an edge. Let $X_{t'} = X_{t''} = X_t$. Now connect all vertices in N_1 to t' and all vertices in N_2 to t''. Clearly the result is a relaxed tree decomposition of G. Iteratively apply the above splitting process to obtain the requested relaxed tree decomposition. □

5.2 Thickness

The reason for studying strong and relaxed path and tree decompositions is that there is a natural way to bound the width of such decompositions on unit disk graphs. We capture the width in a geometric notion, called thickness. Below we define the thickness of a unit disk graph and give a relation to strong and relaxed pathwidth.

Assume that we are given a unit disk graph G with a known 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 and $r_v = 1/2$ is its radius.

The thickness of a unit disk graph is determined by a slab decomposition of a representation of that graph. Given an angle α ($0 \leq \alpha < \pi$) and a point $p \in \mathbb{R}^2$, partition the plane using an infinite set of parallel lines (called *bars*) such that the distance between any two consecutive lines is precisely 1, the lines are parallel and intersect the x-axis at angle α, and (exactly) one line goes through p. The area within distance 1/2 of a bar is called a *slab*. A disk is said to be *in* a slab if its center is contained in the interior or lies on the left boundary of the slab. The parallel lines thus induce a partition of $\mathcal{D}$, which we call a *slab decomposition* (α, p) of $\mathcal{D}$. This in turn induces a decomposition of $V(G)$ into pairwise disjoint, but collectively exhaustive subsets $Y_1, \ldots, Y_k$ such that Y_j contains the vertices corresponding to the disks contained in the j-th nonempty slab from the left.

Definition 5.2.1 *The* thickness *of a slab decomposition* (α, p) *of* $\mathcal{D}$ *is the maximum number of disks of* $\mathcal{D}$ *in any slab of the decomposition, i.e. it is* $\max_{1 \leq j \leq k}\{|Y_j|\}$.

For any fixed angle α ($0 \leq \alpha < \pi$), the *min-thickness* $t^*_\alpha(\mathcal{D})$ is the minimum thickness of any slab decomposition (α, p) over all $p \in \mathbb{R}^2$. Similarly, the *max-thickness* $\bar{t}_\alpha(\mathcal{D})$ is the maximum thickness of any slab decomposition (α, p) over all $p \in \mathbb{R}^2$.

Definition 5.2.2 *The* thickness $t^*(\mathcal{D})$ *is the minimum min-thickness* $t^*_\alpha(\mathcal{D})$ *over all angles* α $(0 \leq \alpha < \pi)$. *The* minimax thickness $\bar{t}(\mathcal{D})$ *is the minimum max-thickness* $\bar{t}_\alpha(\mathcal{D})$ *over all angles* α $(0 \leq \alpha < \pi)$.

The thickness and the minimax thickness can both be computed in polynomial time by exhaustively enumerating all relevant angles and points [258].

There is an easy relation between the thickness of $\mathcal{D}$ and its minimax thickness, given by the following proposition.

Proposition 5.2.3 $t^*(\mathcal{D}) \leq \bar{t}(\mathcal{D}) \leq 2 \cdot t^*(\mathcal{D})$. *Moreover these bounds are tight.*

Proof: The inequality $t^*(\mathcal{D}) \leq \bar{t}(\mathcal{D})$ holds by definition. For any angle α, a slab of one slab decomposition with angle α can overlap at most two of any other. Hence $\bar{t}_\alpha(\mathcal{D}) \leq 2 \cdot t^*_\alpha(\mathcal{D})$ for any angle α ($0 \leq \alpha < \pi$) and thus $\bar{t}(\mathcal{D}) \leq 2 \cdot t^*(\mathcal{D})$. The tightness of these inequalities is demonstrated by respectively two disjoint disks and two disks with intersecting interiors. □

The definition of minimax thickness already hints at a relation to the pathwidth of a unit disk graph.

Theorem 5.2.4 *The pathwidth of a unit disk graph with representation $\mathcal{D}$ is at most $\bar{t}(\mathcal{D}) - 1$.*

Proof: Let α be an angle such that $\bar{t}_\alpha(\mathcal{D}) = \bar{t}(\mathcal{D})$. Let B be the set of all straight lines parallel to a line intersecting the x-axis at angle α. For any $b \in B$, define X_b as the set of disks intersecting b. Let B' be any minimal subset of B such that $\{X_b \mid b \in B'\} = \{X_b \mid b \in B\}$. By definition, $\max_{b \in B'}\{|X_b|\} \leq \bar{t}_\alpha(\mathcal{D})$. Moreover, by ordering the bars in B' from left to right, the sets X_b with $b \in B'$ induce a path decomposition. □

This implies that the relaxed pathwidth is at most $\bar{t}(\mathcal{D})$. One can however improve on this bound by considering the strong pathwidth of a unit disk graph instead of its pathwidth.

Theorem 5.2.5 *The strong pathwidth of a unit disk graph with representation $\mathcal{D}$ is at most $t^*(\mathcal{D})$.*

Proof: Consider any slab decomposition of thickness $t^*(\mathcal{D})$ and let $Y_1, \ldots, Y_k$ be the induced partition of $V(G)$. Since the slabs have width 1 and two unit disks intersect if and only if the distance between their centers is at most 1, $Y_1, \ldots, Y_k$ is a strong path decomposition of G of width $t^*(\mathcal{D})$. □

Corollary 5.2.6 *The relaxed pathwidth of a unit disk graph with representation $\mathcal{D}$ is at most $t^*(\mathcal{D})$.*

Observe that following Proposition 5.2.3 this bound is always better, possibly by as much as a factor of 2, then the bound that followed from Theorem 5.2.4. Hence it makes sense to consider algorithms for relaxed path decompositions, as we might attain better worst-case running times than when using the path decomposition given by Theorem 5.2.4.

5.3 Algorithms on Strong, Relaxed Tree Decompositions

We have shown how to transform an optimum slab decomposition into a strong and a relaxed path decomposition. However, as far as we know, no algorithms exist for problems like Maximum Independent Set and Minimum (Connected) Dominating Set that make use of such a decomposition. We develop these algorithms in this section and prove that we can match or improve on the running times attained for path decompositions. In fact, the given algorithms apply to relaxed tree decompositions and thus can also be applied in a more general setting.

5.3.1 Maximum Independent Set and Minimum Vertex Cover

The idea behind the algorithm is similar to the idea behind the algorithm on tree decompositions by Telle and Proskurowski [249]. Let G be any graph and (T, X) a relaxed tree decomposition of G. Fix a vertex $r \in V(T)$ and root the tree T at r. Define a function size as follows. If $u \in V(T)$ has no children, then for any independent set $A_u \subseteq X_u$,

$$\mathsf{size}_u(A_u) = 0.$$

For all $u \in V(T)$ with children $u_1, \ldots, u_\rho$, let for any independent set $A_u \subseteq X_u$,

$$\mathsf{size}_u(A_u) = \sum_{i=1}^{\rho} \max_{A_{u_i}} \Big\{ |A_{u_i}| + \mathsf{size}_{u_i}((A_u \cup A_{u_i}) \cap X_{u_i}) \Big\},$$

where the maximum is over all independent sets $A_{u_i} \subseteq X_{u_i} - X_u - N(A_u)$.

For any $u \in V(T)$, let T_u denote the subtree of T rooted at u.

Lemma 5.3.1 *The cardinality of a maximum independent set of graph G is* $\max_{A_r}\{\mathsf{size}_r(A_r) + |A_r|\}$. *The maximum is over all independent sets $A_r \subseteq X_r$.*

Proof: We claim that for any $u \in V(T)$ and for any $A_u \subseteq X_u$, $\mathsf{size}_u(A_u)$ is the maximum cardinality of any set $I \subseteq \bigcup_{t \in V(T_u) \setminus \{u\}} X_t - X_u$ for which $I \cup A_u$ is an independent set, or $-\infty$ if no such set exists.

We prove this by induction. If u has no children, then $\bigcup_{t \in V(T_u) \setminus \{u\}} X_t = \emptyset$ and the claim is immediate from the definition of size_u. So suppose that u has children $u_1, \ldots, u_\rho$ and that the claim holds for $u_1, \ldots, u_\rho$. For any $i \in \{1, \ldots, \rho\}$ and any independent set $A_u \subseteq X_u$, it follows by induction that

$$\max_{A_{u_i}} \Big\{ |A_{u_i}| + \mathsf{size}_{u_i}((A_u \cup A_{u_i}) \cap X_{u_i}) \Big\},$$

where the maximum is over all independent sets $A_{u_i} \subseteq X_{u_i} - X_u - N(A_u)$, is the maximum cardinality of any set $I_i \subseteq \bigcup_{t \in V(T_{u_i})} X_t - X_u$ for which $I_i \cup A_u$ is an independent set, or $-\infty$ is no such set exists. Furthermore, $\bigcup_{t \in V(T_{u_i})} X_t - X_u$ and $\bigcup_{t \in V(T_{u_j})} X_t - X_u$ for any $1 \leq i < j \leq \rho$ are not connected. Hence $\mathsf{size}_u(A_u)$ is indeed as claimed.

The lemma is immediate from the proof of the claim. □

It is also easy to compute (the cardinality of) a maximum independent set using size. Let w be the width of (T, X), let $M = \max_{(t,t') \in E(T)} |X_t \cup X_{t'}|$, and let $n = |V(G)|$.

Theorem 5.3.2 *One can compute (the cardinality of) a maximum independent set of G in* $\mathrm{O}(nw\,2^M)$ *time.*

Proof: This follows from the definition of size and from Lemma 5.3.1. □

If (T, X) is a strong or a relaxed tree decomposition, then M is at most twice the width of the decomposition. If (T, X) is an ordinary tree decomposition of width w, then we can assume that $M \leq w + 1$ [168, p. 149].

Corollary 5.3.3 *One can compute (the cardinality of) a maximum independent set of G in $\mathrm{O}(nw\, 2^{2\cdot \mathrm{stw}(G)})$, $\mathrm{O}(nw\, 2^{2\cdot \mathrm{rtw}(G)})$, or $\mathrm{O}(nw\, 2^{\mathrm{tw}(G)})$ time, assuming the appropriate graph decomposition is given.*

Recall that $\mathrm{rtw}(G) \leq \mathrm{tw}(G) \leq 2 \cdot \mathrm{rtw}(G) - 1$. For Maximum Independent Set, it is therefore more beneficial to have a tree decomposition.

The same holds for Minimum Vertex Cover. It is well-known that I is an independent set of a graph G if and only if $V(G) - I$ is a vertex cover of G.

Theorem 5.3.4 *One can compute (the cardinality of) a minimum vertex cover of G in $\mathrm{O}(nw\, 2^{2\cdot \mathrm{stw}(G)})$, $\mathrm{O}(nw\, 2^{2\cdot \mathrm{rtw}(G)})$, or $\mathrm{O}(nw\, 2^{\mathrm{tw}(G)})$ time, assuming the appropriate graph decomposition is given.*

5.3.2 Minimum Dominating Set

We give two algorithms for Minimum Dominating Set. First, we present a simple algorithm using a strong path decomposition. It has a running time of $\mathrm{O}(nw\, 2^{3\cdot \mathrm{spw}(G)})$ and is crucial to the approximation schemes of Chapter 6. We then consider a more complex algorithm, which can use any relaxed tree decomposition as input and has running time $\mathrm{O}(nw\, 2^{2\cdot \mathrm{rtw}(G)})$ or $\mathrm{O}(nw\, 3^{\mathrm{pw}(G)})$.

A Simple Algorithm

Let G be any graph and (T, X) a strong path decomposition of G. We can view the bags as being in sequence and number them accordingly $X_1, \ldots, X_p$. Observe that for any i, the vertices in X_i can only be dominated by vertices in X_{i-1} (if $i > 1$), X_i, or X_{i+1} (if $i < p$). This idea can be exploited as follows. Define for any $A_1 \subseteq X_1$ and $A_2 \subseteq X_2$,

$$\mathsf{size}_1(A_1, A_2) = \begin{cases} 0 & \text{if } A_1 \cup A_2 \text{ dominates } X_1 \\ \infty & \text{otherwise.} \end{cases}$$

Define for any $i = 2, \ldots, p-1$, any $A_i \subseteq X_i$, and any $A_{i+1} \subseteq X_{i+1}$,

$$\begin{aligned} \mathsf{size}_i(A_i, A_{i+1}) \quad = \quad & \min\{|A_{i-1}| + \mathsf{size}_{i-1}(A_{i-1}, A_i) \mid \\ & \qquad A_{i-1} \subseteq X_{i-1} \text{ and } A_{i-1} \cup A_i \cup A_{i+1} \text{ dominates } X_i\}. \end{aligned}$$

Moreover, for any $A_p \subseteq X_p$,

$$\begin{aligned} \mathsf{size}_p(A_p) \quad = \quad & \min\{|A_{p-1}| + \mathsf{size}_{p-1}(A_{p-1}, A_p) \mid \\ & \qquad A_{p-1} \subseteq X_{p-1} \text{ and } A_{p-1} \cup A_p \text{ dominates } X_p\}. \end{aligned}$$

Lemma 5.3.5 *The cardinality of a minimum dominating set of graph G is $\min_{A_p}\{\mathsf{size}_p(A_p) + |A_p|\}$, where the minimum is over all $A_p \subseteq X_p$.*

Proof: We claim that for any $i \in \{1, \ldots, p-1\}$, any $A_i \subseteq X_i$, and any $A_{i+1} \subseteq X_{i+1}$, $\mathsf{size}_i(A_i, A_{i+1})$ is the cardinality of a smallest set $R \subseteq \bigcup_{j=1}^{i-1} X_j$ such that $R \cup A_i \cup A_{i+1}$ dominates $\bigcup_{j=1}^{i} X_j$, or ∞ if no such set exists. Applying induction, this follows immediately from the definition of size. Hence for any $A_p \subseteq X_p$, $\mathsf{size}_p(A_p)$ is the cardinality of a smallest set $R \subseteq \bigcup_{j=1}^{p-1} X_j$ such that $R \cup A_p$ dominates $\bigcup_{j=1}^{p} X_j$, or ∞ if no such set exists. As G has a dominating set, $\min_{A_p}\{\mathsf{size}_p(A_p) + |A_p|\}$ is the cardinality of a minimum dominating set. □

Theorem 5.3.6 *One can compute (the cardinality of) a minimum dominating set of G in $\mathrm{O}(nw\, 2^{3w})$ time, given a strong path decomposition of width w.*

Clearly this algorithm extends to strong tree decompositions, although this increases the running time of the algorithm. It can even be extended to relaxed tree decompositions at a higher cost. Here we chose to keep the algorithm and its description simple. We will only use it on a strong path decomposition. Moreover, a much faster algorithm exists on relaxed tree decompositions.

A Better Algorithm

We improve on the running time of the above algorithm by extending it to relaxed tree decompositions. We apply an idea similar to the one used in the algorithm by Alber et al. [6, 8].

Let G be any graph and (T, X) a relaxed tree decomposition of G. By Lemma 5.1.9, we may assume that any vertex of T has degree at most three. Fix a vertex $r \in V(T)$ which has degree at most one and root the tree T at r. Define a function size as follows. If $u \in V(T)$ has no children, then for any $A_u \subseteq X_u$ and any $B_u \subseteq X_u - A_u$,

$$\mathsf{size}_u(A_u, B_u) = \begin{cases} 0 & \text{if } B_u \subseteq N(A_u) \\ \infty & \text{otherwise.} \end{cases}$$

If $u \in V(T)$ has children $u_1, \ldots, u_\rho$, then for any $A_u \subseteq X_u$ and $B_u \subseteq X_u - A_u$,

$$\begin{aligned} &\mathsf{size}_u^i(A_u, B_u) \\ &\quad = \min\Big\{|A_{u_i}| + \mathsf{size}_{u_i}\Big((A_u \cup A_{u_i}) \cap X_{u_i},\ X_{u_i} - X_u - N[A_u \cup A_{u_i}]\Big) \;\Big|\; \\ &\qquad\qquad A_{u_i} \subseteq X_{u_i} - X_u, B_u \subseteq N(A_{u_i} \cup A_u)\Big\} \end{aligned}$$

$$\mathsf{size}_u(A_u, B_u) = \min_{(B_1, \ldots, B_\rho) \text{ partitioning } B_u} \left\{ \sum_{i=1}^{\rho} \mathsf{size}_u^i(A_u, B_i) \right\}.$$

Lemma 5.3.7 *The cardinality of a minimum dominating set of graph G is $\min_{A_r \subseteq X_r}\{\mathsf{size}_r(A_r, X_r - A_r) + |A_r|\}$.*

Proof: It suffices to prove the following. We claim that for any $u \in V(T)$, any $A_u \subseteq X_u$, and any $B_u \subseteq X_u - A_u$, $\mathsf{size}_u(A_u, B_u)$ is the cardinality of a smallest set $R \subseteq \bigcup_{t \in V(T_u)} X_t - X_u$ such that $R \cup A_u$ dominates $B_u \cup \big(\bigcup_{t \in V(T_u)} X_t - X_u\big)$, or ∞ if no such set exists.

If $u \in V(T)$ has no children, then this is immediate from the definition of size. So suppose that $u \in V(T)$ has children $u_1, \ldots, u_\rho$ and the claim holds for each child. Let $A_u \subseteq X_u$, $B_u \subseteq X_u - A_u$. Then $\mathsf{size}^i_u(A_u, B_u)$ is the cardinality of a smallest set $R_i \subseteq \bigcup_{t \in V(T_{u_i})} X_t - X_u$ such that $R_i \cup A_u$ dominates $B_u \cup (\bigcup_{t \in V(T_{u_i})} X_t - X_u)$, or ∞ if no such set exists. Observe that $\bigcup_{t \in V(T_{u_i})} X_t - X_u$ and $\bigcup_{t \in V(T_{u_j})} X_t - X_u$ are disjoint for any $1 \leq i < j \leq \rho$. Hence using its definition, $\mathsf{size}_u(A_u, B_u)$ is indeed as claimed. □

A trivial way to compute the function size would be to follow its definition. However, the running time would be $\mathrm{O}(nw\,(2^w 3^w + 5^w))$, where w is the width of the decomposition. One can improve on this naive analysis to get a better worst-case running time. Let again $M = \max_{(t,t') \in E(T)} |X_t \cup X_{t'}|$.

Theorem 5.3.8 *One can compute (the cardinality of) a minimum dominating set of graph G in $\mathrm{O}(nw\,(2^M + 4^w))$ time.*

Proof: If $u \in V(T)$ has no children, then size_u is computable in $\mathrm{O}(w3^w)$ time.

Consider some $u \in V(T)$ with children $u_1, \ldots, u_\rho$. We compute size^i_u in two phases. First, enumerate all $A_u \subseteq X_u$ and $A_{u_i} \subseteq X_{u_i} - X_u$. Set $B_u = N(A_u \cup A_{u_i}) \cap X_u$. If this set B_u was not encountered before or if $|A_{u_i}| + \mathsf{size}_{u_i}(\ldots)$ is smaller than the previous value of $\mathsf{size}^i_u(A_u, B_u)$, then set

$$\mathsf{size}^i_u(A_u, B_u) = |A_{u_i}| + \mathsf{size}_{u_i}((A_u \cup A_{u_i}) \cap X_{u_i}, X_{u_i} - X_u - N[A_u \cup A_{u_i}]).$$

This takes $\mathrm{O}(w\rho\, 2^M)$ time.

There is no guarantee that we see all values $B_u \subseteq X_u - A_u$ for a given A_u. However, we do see all 'maximal' values. Note that the following inequality should hold: $\mathsf{size}^i_u(A_u, B'_u) \leq \mathsf{size}^i_u(A_u, B_u)$ for any $A_u \subseteq X_u$ and any $B'_u \subseteq B_u \subseteq X_u - A_u$. So enumerate all $A_u \subseteq X_u$ and $B_u \subseteq X_u - A_u$ with $B_u \neq \emptyset$ *in order*, meaning that if $B'_u \subseteq B''_u$, then B''_u is considered before B'_u. Now update size^i_u as follows. For any $x \in B_u$, let

$$\mathsf{size}^i_u(A_u, B_u \backslash \{x\}) = \min\{\mathsf{size}^i_u(A_u, B_u \backslash \{x\}), \mathsf{size}^i_u(A_u, B_u)\},$$

where we assume that $\mathsf{size}^i_u(A_u, B_u) = \infty$ or $\mathsf{size}^i_u(A_u, B_u \backslash \{x\}) = \infty$ if B_u or $B_u \backslash \{x\}$ respectively was not considered in the first step. This takes $\mathrm{O}(w\, 3^w)$ time. Moreover, we have now correctly computed size^i_u.

To compute size_u, recall that any vertex of T has degree at most three and that r is a leaf of T. Hence $\rho \leq 2$. Now by enumerating all $A_u \subseteq X_u$, $B_1 \subseteq X_u - A_u$, $B_2 \subseteq X_u - A_u - B_2$, and letting $B_u = B_1 \cup B_2$, it is easy to compute size_u. This takes $\mathrm{O}(4^w)$ time. The theorem follows. □

Observe that if a relaxed path decomposition of width w is given, then $\rho \leq 1$, and the described algorithm runs in $\mathrm{O}(nw\,(2^M + 3^w))$ time. If a strong path decomposition of width w is given, this implies a running time of $\mathrm{O}(nw\,2^{2w})$, which is a factor of 2^w faster than the previous algorithm.

Corollary 5.3.9 *One can compute (the cardinality of) a minimum dominating set of graph G in $\mathrm{O}(nw\,2^{2\cdot \mathrm{rtw}(G)})$ or $\mathrm{O}(nw\,3^{\mathrm{pw}(G)})$ time, assuming the appropriate graph decomposition is known.*

The running time for path decompositions follows from the fact that given a path decomposition of width w, one can assume that $M \leq w + 1$ [168].

Minimum Dominating Set is a clear case where having a relaxed tree decomposition is preferable to having an ordinary tree decomposition. As $\mathrm{rtw}(G) \leq \mathrm{tw}(G)$, and possibly $\mathrm{rtw}(G) = \frac{1}{2} \cdot \mathrm{tw}(G) + \frac{1}{2}$, the worst case running time of $\mathrm{O}(nw\,2^{2\cdot \mathrm{rtw}(G)})$ is preferable to the $\mathrm{O}(n\,2^{2\cdot \mathrm{tw}(G)})$ algorithm by Alber et al. [6, 8]. When considering path decompositions, we improve if $\mathrm{rpw}(G) \leq (\frac{1}{2}\log 3) \cdot \mathrm{pw}(G)$, where $\frac{1}{2}\log 3 \approx 0.792$.

5.3.3 Minimum Connected Dominating Set

We develop a technique for solving the minimum connected dominating set problem that augments both proposed algorithms for Minimum Dominating Set. Hence we effectively obtain two algorithms.

The technique is based on the following general ideas. Let G be a graph and (T, X) a relaxed tree decomposition of G. Root the tree T at some fixed vertex $r \in V(T)$. Given $X \subseteq V(G)$, a set $W \subseteq V(G)$ is *partially connected* (with respect to X) if each connected component of W intersects X. This definition is crucial because of the following easy fact. Let $u \in V(T)$ and let G_u be the graph induced by $\bigcup_{t \in V(T_u)} X_t$. Then for any connected dominating set $W \subseteq V(G)$, $W \cap V(G_u)$ is partially connected with respect to $W \cap X_u$ if $W \cap X_u \neq \emptyset$. This suggests that during the dynamic programming, we should construct partially connected dominating sets.

Suppose that $W \subseteq V(G_u)$ is partially connected with respect to $W \cap X_u$ for some $u \in V(T) \backslash \{r\}$. Then W induces an equivalence relation $\sim_{W,u}$ on the connected components of $W \cap X_u$, namely $C \sim_{W,u} C'$ for connected components C, C' of $W \cap X_u$ if and only if there is a C–C' path in $G[W]$. Let $(u, v) \in E(T)$ such that v is closer to r than u. Given a subset $A_u \subseteq X_u$, an equivalence relation $\sim$ on the connected components of A_u, and a subset $A_v \subseteq X_v - X_u$, we say that A_v is *compatible to* $(A_u, \sim)$ if for any equivalence class $\mathcal{C}$ of $\sim$, there is a connected component $C \in \mathcal{C}$ such that there is a C–A_v path in $G[C \cup A_v]$. Then given a set $W \subseteq V(G_u)$ that is partially connected with respect to $W \cap X_u$ and any $A_v \subseteq X_v - X_u$, $W \cup A_v$ is partially connected with respect to A_v if and only if A_v is compatible to $(W \cap X_u, \sim_{W,u})$. Note that given a compatible set A_v, $\sim_{W \cup A_v, v}$ is uniquely determined by $\sim_{W,u}$. We say that $\sim_{W,u}$ *determines* $\sim_{W \cup A_v, v}$.

We can now use the above ideas as follows. First, we adapt the simple algorithm given in Paragraph 5.3.2. Suppose that (T, X) is a strong path decomposition, i.e. the bags are numbered in sequence, $X_1, \dots, X_p$. Define a function size as follows. Let $\sim_{\text{id}}$ denote the identity equivalence relation, i.e. $C \sim_{\text{id}} C'$ if and only if $C = C'$, and let $\sim_{\text{all}}$ denote the total equivalence relation, i.e. $C \sim_{\text{all}} C'$ for any C, C'. Then for any $A_1 \subseteq X_1$, any $A_2 \subseteq X_2$, and any equivalence relation $\sim$ on the connected components of $G[A_1]$ such that A_2 is compatible to $(A_1, \sim)$,

$$\mathsf{size}_1(A_1, A_2, \sim) = \begin{cases} 0 & \text{if } A_1 \cup A_2 \text{ dominates } X_i \text{ and } \sim \equiv \sim_{\text{id}} \\ \infty & \text{otherwise.} \end{cases}$$

For any $i \in \{2, \dots, p-1\}$, any $A_i \subseteq X_i$, any $A_{i+1} \subseteq X_{i+1}$, and any equivalence relation $\sim$ on the connected components of $G[A_i]$ such that A_{i+1} is compatible to $(A_i, \sim)$,

$$\begin{aligned} \mathsf{size}_i(A_i, A_{i+1}, \sim) \quad = \quad \min \big\{ & |A_{i-1}| + \mathsf{size}_{i-1}(A_{i-1}, A_i, \sim') \mid \\ & A_{i-1} \subseteq X_{i-1}, A_i \text{ compatible to } (A_{i-1}, \sim'), \\ & A_{i-1} \cup A_i \cup A_{i+1} \text{ dominates } X_i, \\ & \text{and } \sim' \text{ determines } \sim \big\}. \end{aligned}$$

Finally, for any $A_p \subseteq X_p$,

$$\begin{aligned} \mathsf{size}_p(A_p) \quad = \quad \min \big\{ & |A_{p-1}| + \mathsf{size}_{p-1}(A_{p-1}, A_p, \sim') \mid \\ & A_{p-1} \subseteq X_{p-1}, A_p \text{ compatible to } (A_{p-1}, \sim'), \\ & A_{p-1} \cup A_p \text{ dominates } X_p, \\ & \text{and } \sim' \text{ determines } \sim, \text{ where } \sim \equiv \sim_{\text{all}} \big\}. \end{aligned}$$

For the analysis, we define $q(G, T, X)$ as the maximum number of equivalence relations one needs to consider during the algorithm for any subset $A_u \subseteq X_u$ for any $u \in V(T)$.

Lemma 5.3.10 *Given a set $\{c_0, \dots, c_k\}$, there is a data structure for storing equivalence relations of this set that uses space the number of stored relations times $\mathrm{O}(k \log k)$ and where insert, update, and find cost $\mathrm{O}(k \log k)$ time.*

Proof: Observe that for any equivalence relation we can always assume that c_h is in the j-th equivalence class for some $j \leq h + 1$. Now consider the following tree. Level h of the tree corresponds to c_h. A node on level h has $h+1$ children, where the j-th child is used to describe that c_{h+1} is in equivalence class j. A path in the tree (from the root to a leaf) thus provides a complete description of an equivalence relation. One can use level $k+1$ to store any value associated with this equivalence relation.

To quickly find the j-th child of a node, we use a balanced binary search tree to store the children. It follows immediately from this description that insert, update, and find each cost $\mathrm{O}(k \log k)$ time. The tree has $k + 1$ levels and on each level, one has a search tree of height $\mathrm{O}(\log k)$ to navigate.

Of course, one does not store the entire tree, but keeps only those paths that correspond to a stored relation. Each path costs $\mathrm{O}(k \log k)$ space to maintain. As each node uses a balanced binary search tree to store its children, insert, update, and find still take $\mathrm{O}(k \log k)$ time. The lemma follows. □

Theorem 5.3.11 *One can compute (the cardinality of) a minimum connected dominating set of graph G in $\mathrm{O}(n \cdot w \log w \cdot q(G,T,X) \cdot 2^{3w})$ time, given a strong path decomposition (T,X) of width w.*

Proof: It immediately follows from the proof of Lemma 5.3.5 that the set attaining $\min_{A_p \subseteq X_p}\{\mathsf{size}_p(A_p) + |A_p|\}$ is a dominating set of G. Moreover, the compatibility constraints in the definition of size ensure that this set is connected. Hence it is easily shown that $\min_{A_p \subseteq X_p}\{\mathsf{size}_p(A_p) + |A_p|\}$ is the cardinality of a minimum connected dominating set of G.

To compute size, we use the data structure described in the above lemma. We only maintain those equivalence relations $\sim$ which are actually realizable, that is, for which $\mathsf{size}_i(A_i, A_{i+1}, \sim) \neq \infty$. Hence for fixed $A_i \subseteq X_i$, $A_{i+1} \subseteq X_{i+1}$, we have a data structure requiring $\mathrm{O}(w \log w \cdot q(G,T,X))$ space with $\mathrm{O}(w \log w)$ time insert, update, and find operations.

It follows that size_1 can be computed in $\mathrm{O}(w \log w \cdot 2^{2w})$ time. In computing size_i for $i > 1$, we deviate slightly from the recursive formula. For any $A_{i-1} \subseteq X_{i-1}$, $A_i \subseteq X_i$, and $A_{i+1} \subseteq X_{i+1}$ such that $A_{i-1} \cup A_i \cup A_{i+1}$ dominates X_i, we consider all relations $\sim'$ for which $\mathsf{size}_{i-1}(A_{i-1}, A_i, \sim')$ is stored. Then we find the equivalence relation $\sim$ determined by $\sim'$ and check whether A_{i+1} is compatible to $(A_i, \sim)$. We then store $\mathsf{size}_i(A_i, A_{i+1}, \sim) = |A_{i-1}| + \mathsf{size}_{i-1}(A_{i-1}, A_i, \sim')$ if $\sim$ was not encountered before, or update this value if necessary. One computes size_p similarly.

Observe that during this process, it is not necessary to explicitly enumerate all equivalence relations. The relevant relations may be obtained from the tables. Thus in the analysis of the running time, it suffices to bound the maximum number of equivalence relations, which is $q(G,T,X)$. Hence for any $i > 1$, we spend $\mathrm{O}(w \log w \cdot q(G,T,X) \cdot 2^{3w})$ time. The theorem follows. □

In a similar manner, one can adapt the second algorithm of Section 5.3.2.

Theorem 5.3.12 *One can compute (the cardinality of) a minimum connected dominating set of graph G in $\mathrm{O}(n \cdot w \log w \cdot q(G,T,X) \cdot (2^M + 4^w))$ time, given a relaxed tree decomposition (T,X) of width w.*

It remains to bound $q(G,T,X)$. For an arbitrary graph G and an arbitrary relaxed tree decomposition (T,X) of width w, $q(G,T,X)$ is bounded by the number of different equivalence relations on a set of cardinality w. In other words, it is bounded by the number of distinct partitions of a w-element set.

Definition 5.3.13 *Given a set S, $S_1, \ldots, S_p$ is a* partition *of S if $S_i \neq \emptyset$ for any $1 \leq i \leq p$, $S = \bigcup_{i=1}^{p} S_i$, and for any $1 \leq i < j \leq p$, $S_i \cap S_j = \emptyset$.*

The number of partitions of an w-element set is equal to the w-th Bell number, ϖ_w, named for E.T. Bell [29, 30]. We give an upper bound.

Proposition 5.3.14 *$\varpi_w \leq (\frac{w}{\ln w})^w$ for $w \geq 2$.*

Proof: Since $x \leq e^{x-1}$ for any $x \in \mathbb{R}$, we have that $ex \leq e^x$. Then for any k,

$$\begin{aligned} e \cdot \frac{k \ln w}{w} &\leq e^{\frac{k \ln w}{w}} && \Rightarrow \\ \left(\frac{ke \ln w}{w}\right)^w &\leq e^{k \ln w} && \Rightarrow \\ k^w &\leq \left(\frac{w}{e \ln w}\right)^w w^k. \end{aligned}$$

Hence following Dobiński's formula [88] (see also Wilf [271]),

$$\varpi_w = \frac{1}{e} \sum_{k=0}^{\infty} \frac{k^w}{k!} \leq \left(\frac{w}{e \ln w}\right)^w \sum_{k=0}^{\infty} \frac{w^k}{k!} = \left(\frac{w}{e \ln w}\right)^w e^w = \left(\frac{w}{\ln w}\right)^w.$$

The bound follows. □

Corollary 5.3.15 *One can compute (the cardinality of) a minimum connected dominating set of graph G in $O(n \cdot w \log w \cdot (\frac{w}{\ln w})^w\, 2^{3w})$ time, given a strong path decomposition (T, X) of width w.*

Corollary 5.3.16 *One can compute (the cardinality of) a minimum connected dominating set of graph G in $O(n \cdot w \log w \cdot (\frac{w}{\ln w})^w\, (2^M + 4^w))$ time, given a relaxed tree decomposition (T, X) of width w.*

It can be easily verified that all algorithms in this section also apply to the weighted case of the problems.

5.4 Unit Disk Graphs of Bounded Thickness

Recall from Theorem 5.2.5 that the strong pathwidth of a unit disk graph with representation $\mathcal{D}$ is at most $t^*(\mathcal{D})$. The results in the previous section yield the following theorem.

Theorem 5.4.1 *Let G be a unit disk graph with representation $\mathcal{D}$ and let $t = t^*(\mathcal{D})$. Then Maximum Independent Set, Minimum Vertex Cover, and Minimum Dominating Set can be solved in $O(nt\, 2^{2t})$ time. Minimum Connected Dominating Set can be solved in $O(n \cdot t \log t \cdot (\frac{t}{\ln t})^t\, 2^{2t})$ time.*

The theorem implies that if the thickness is bounded, say by a constant, then these problems can be solved in polynomial time.

One can improve on the worst-case analysis of the algorithm for Minimum Connected Dominating Set by using that the given strong path decomposition comes from a strip decomposition. So let (T, X) be such a strong path decomposition. Instead of bounding $q(G, T, X)$ by ϖ_w, we will show that many relations are in fact the same. Hence we can use so-called noncrossing partitions to bound $q(G, T, X)$ instead of general partitions.

Definition 5.4.2 *Let S be a finite set and $\preceq$ a total order of its elements. Then $S_1, \ldots, S_p$ is a* noncrossing partition *of S if $S_1, \ldots, S_p$ is a partition of S and for any $1 \leq i, j \leq p$ with $i \neq j$, any $a, c \in S_i$, and any $b, d \in S_j$, $a \preceq b \preceq c \preceq d$ is false.*

Noncrossing partitions were first considered by Becker [27, 28]. Refer to Simion [240] for numerous applications of such partitions. By Becker [28] and Kreweras [183], the number of noncrossing partitions of a w-element set is C_w, the w-th *Catalan number*. It is well-known that $C_w \approx 4^w$. If the cardinality of a set can be bounded by a polynomial in C_w, it is said to have *Catalan structure*. Catalan structure was considered before in the context of Minimum Connected Dominating Set, for example on planar graphs [84]. These ideas were subsequently generalized by Dorn, Fomin, and Thilikos [90, 91] to more general graph classes (see also Section 6.5). This however is the first application to unit disk graphs.

Let G be a unit disk graph with representation $\mathcal{D}$ and (T, X) a strong path decomposition induced by a strip decomposition through Theorem 5.2.5. The bags are numbered in sequence $X_1, \ldots, X_p$. We exhibit the existence of Catalan structure. We start with an easy bound. Without loss of generality, assume that no two disk centers have the same y-coordinate. Then the following is a total order. For any two vertices $u, v \in V(G)$, let $u \preceq v$ if and only if $u = v$ or the disk center of $\mathcal{D}(u)$ lies below the disk center of $\mathcal{D}(v)$.

Lemma 5.4.3 *For any $1 \leq i \leq p$, let $W \subseteq \bigcup_{j=1}^{i} X_j$ be partially connected with respect to $A_i := W \cap X_i$ and for any $u, v \in A_i$, $u \sim v$ if and only if $u = v$ or there is a u–v path P_{uv} in W such that $P_{uv} \cap A_i = \{u, v\}$. If $a \sim c$ and $b \sim d$ for distinct $a, b, c, d \in A_i$ such that $a \preceq b \preceq c \preceq d$, then a, b, c, and d are connected in W.*

Proof: As $a \preceq b \preceq c \preceq d$, P_{ac} and P_{bd} must cross, meaning there are adjacent vertices $w, x \in P_{ac}$ and adjacent vertices $y, z \in P_{bd}$ such that the line segment $\overline{c_w c_x}$ intersects the line segment $\overline{c_y c_z}$. As G is a unit disk graph, both segments have length at most 1 and c_w is within distance 1 of c_y or c_z, or c_x is within distance 1 of c_y or c_z. Thus at least one of the edges (w, y), (w, z), (x, y), or (x, z) must be in $E(G)$. Hence a, b, c, and d are connected in W. □

Observe that a relation $\sim$ on the vertices of A_i as in the above lemma induces an equivalence relation $\sim_{W,i}$ on the connected components of $W \cap X_i$ as needed for the algorithm for Minimum Connected Dominating Set. In fact, one can construct a surjective map from equivalence relations $\sim$ to $\sim_{W,i}$. Hence if we bound the number of equivalence relations $\sim$, we have bounded the number of equivalence relations $\sim_{W,i}$.

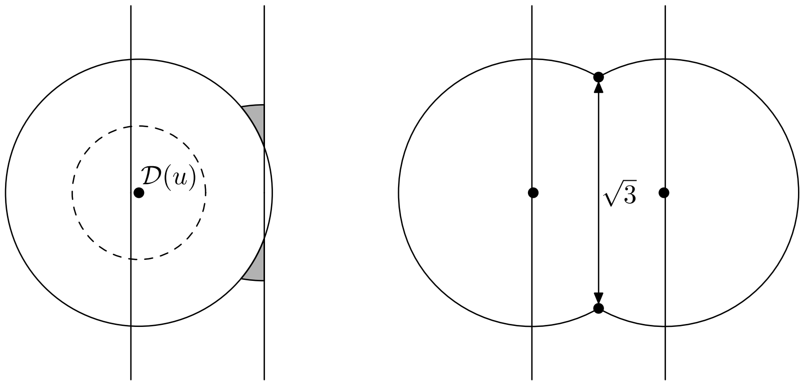

Figure 5.1: In the left figure, the disk $\mathcal{D}(u)$ is drawn dashed. The solid circle indicates the area in which the centers of $\mathcal{D}(v)$ and $\mathcal{D}(w)$ cannot lie, or they would intersect $\mathcal{D}(u)$. Then the centers of $\mathcal{D}(v)$ and $\mathcal{D}(w)$ have to be within the shaded areas. In the right figure, the area in which the centers of $\mathcal{D}(a)$ and $\mathcal{D}(b)$ cannot lie is indicated.

Theorem 5.4.4 *One can compute (the cardinality of) a minimum connected dominating set of a unit disk graph G in $\mathrm{O}(n \cdot t \log t \cdot 2^{5t})$ time, where $t = t^*(D)$ is the thickness of a representation $\mathcal{D}$ of G.*

Proof: From Lemma 5.4.3, it follows that $\sim$ induces a noncrossing partition on the vertices of A_i. Hence we can bound $q(G, T, X)$ by C_t. The theorem now follows from Theorem 5.3.11. □

Using Theorem 5.3.12, we can give the following slightly better result.

Theorem 5.4.5 *One can compute (the cardinality of) a minimum connected dominating set of a unit disk graph G in $\mathrm{O}(n \cdot t \log t \cdot 2^{4t})$ time, where $t = t^*(D)$ is the thickness of a representation $\mathcal{D}$ of G.*

It is possible to improve on this analysis if the number of connected components of A_i is smaller than $|A_i|$. Instead of applying noncrossing partitions to the vertices, we apply them to (parts of) connected components.

By rotating if necessary, we may assume that the slab boundaries of the underlying slab decomposition are parallel to the y-axis. Let $i \in V(T)$. Define the *y-range* $\bar{y}(S)$ of a connected subset $\emptyset \neq S \subseteq X_i$ as the range between the smallest and the largest y-coordinate of any disk center of S. The *contour* of S is the perimeter of the union of the disks of S. Consider the connected components of $G[A_i]$ for some $A_i \subseteq X_i$. Two connected components are said to *interact* if their y-ranges intersect.

Proposition 5.4.6 *No three connected components can pairwise interact.*

Proof: If three connected components C, C', and C'' pairwise interact, their y-ranges pairwise intersect and thus the y-ranges have a common y-coordinate. Without loss of generality, this common coordinate corresponds to the y-coordinate $y(u)$ of a disk center for some $u \in C$. Then $y(u) \in \bar{y}(C')$ and $y(u) \in \bar{y}(C'')$. Consider C' and let v be the vertex in C' such that $y(v)$ is minimal under the condition that $y(v) \geq y(u)$. Similarly, let $w \in V(C')$ such that $y(w)$ is maximal, but $y(w) \leq y(u)$. Define $a, b \in C''$ in a similar way. By symmetry, we may assume that the center of $\mathcal{D}(u)$ lies on or to the left of the middle of the slab. Then the centers of $\mathcal{D}(v)$ and $\mathcal{D}(w)$ lie in the shaded area of Figure 5.1. In the best case, u lies on the left boundary of the slab and $y(v) = y(u) + \epsilon$, $y(w) = y(u) - \epsilon$ for infinitesimally small $\epsilon > 0$. But then the centers of $\mathcal{D}(a)$ and $\mathcal{D}(b)$ must be at distance at least $\sqrt{3} > 1$, which contradicts that $\mathcal{D}(a)$ and $\mathcal{D}(b)$ intersect. □

Corollary 5.4.7 *If $G[A_i]$ has c connected components, then at most $c - 1$ pairs of connected components interact.*

Proof: Consider the (interval) graph H induced by the y-ranges of the connected components of $G[A_i]$. By Proposition 5.4.6, no three intervals of H pairwise intersect. Hence H is a forest, where each vertex corresponds to a connected component and each edge to an interaction. Since H is a forest, the number of edges of H, and thus the number of pairwise interacting connected components is at most $c - 1$. □

This corollary will prove to be useful in counting arguments later on.

Following the proof of Proposition 5.4.6 and as the contours of no two connected components intersect, we can say of two interacting connected components C and C' that C is *in front* of C', if (when their y-ranges overlap) the contour of C lies to the left of the contour of C'.

Suppose that two connected components C and C' interact and C is in front of C'. Now split the vertices of C' into three parts: those vertices u for which $y(u) > \bar{y}(C)$, for which $y(u) \in \bar{y}(C)$, and for which $y(u) < \bar{y}(C)$. Iteratively decompose the connected components this way. What remains are *blocks* of vertices. Let $\mathcal{B}$ denote the set of all blocks. By construction, for any two distinct blocks $B, B' \in \mathcal{B}$, either $\bar{y}(B) \cap \bar{y}(B') = \emptyset$, $\bar{y}(B) \subseteq \bar{y}(B')$, or $\bar{y}(B) \supseteq \bar{y}(B')$. Note that the vertices of a block are connected. If $\bar{y}(B) \subseteq \bar{y}(B')$ for two blocks B and B', then B is *occluded* by B'. We may also call B' the *occluder* of B.

Consider the following total order $\preceq$ on the blocks. For $B, B' \in \mathcal{B}$, let $B \preceq B'$ if and only if $B = B'$, $\bar{y}(B) < \bar{y}(B')$, or $\bar{y}(B) \subseteq \bar{y}(B')$. Let $W \subseteq \bigcup_{j=1}^{i} X_j$ be such that W is partially connected with respect to $A_i := W \cap X_i$. Define $\sim$ on $\mathcal{B}$ such that $B \sim B'$ if and only if $B = B'$ or there is a B–B' path in $(W - A_i) \cup B \cup B'$. We claim that $\sim$ induces a noncrossing partition on $\mathcal{B}$ with respect to $\preceq$.

Lemma 5.4.8 *Consider distinct $B_a, B_b, B_c, B_d \in \mathcal{B}$ such that $B_a \sim B_c$, $B_b \sim B_d$, and $B_a \preceq B_b \preceq B_c \preceq B_d$. Then $B_a \sim B_b$.*

Proof: We aim to show that $\bar{y}(B_a) < \bar{y}(B_b) < \bar{y}(B_c) < \bar{y}(B_d)$. The lemma then follows from Lemma 5.4.3. Let (B, B') be an arbitrary pair of (B_a, B_b), (B_b, B_c), (B_c, B_d). By assumption, $B \preceq B'$ and $B \not\sim B'$. We show that $\bar{y}(B) < \bar{y}(B')$. Using the definition of $\preceq$ and since $B \neq B'$, it suffices to prove that assuming $\bar{y}(B) \subseteq \bar{y}(B')$ leads to a contradiction.

So suppose that $\bar{y}(B) \subseteq \bar{y}(B')$. By the construction of the blocks, this implies that B' occludes B. But then $N(B) \cap X_{i-1} \subseteq N(B') \cap X_{i-1}$. Hence any path in W to or from B must contain a vertex not in B neighboring a vertex in B'. As there is such a path, $B \sim B'$, a contradiction.

It follows that $\bar{y}(B_a) < \bar{y}(B_b) < \bar{y}(B_c) < \bar{y}(B_d)$ and thus that $B_a \sim B_b$ by using Lemma 5.4.3. □

Observe that one can construct a surjective map from relations $\sim$ on the blocks to relations $\sim_{W,i}$ on connected components of A_i. Hence it suffices to bound the number of relations $\sim$ on the blocks.

Theorem 5.4.9 *One can compute (the cardinality of) a minimum connected dominating set of a unit disk graph G with representation $\mathcal{D}$ in $\mathrm{O}(n \cdot c \log c \cdot 2^{3c} 2^{3t})$ time, where $t = t^*(\mathcal{D})$ and c is the largest number of connected components in any $A_i \subseteq X_i$ in the strong path decomposition (T, X) corresponding to the slab decomposition supporting $t^*(\mathcal{D})$.*

Proof: It follows from Corollary 5.4.7 that the number of blocks is at most $3(c-1)$. Lemma 5.4.8 showed that $\sim$ induces a noncrossing partition of $\mathcal{B}$ with respect to $\preceq$. Hence the number of relations $\sim$ is at most $C_{|\mathcal{B}|} \leq C_{3(c-1)}$ and we can bound $q(G, T, X)$ by $C_{3(c-1)}$. The theorem immediately follows from Theorem 5.3.11. □

Because $c \leq t$, the running time of the above algorithm is $\mathrm{O}(n \cdot t \log t \cdot 2^{6t})$ in the worst case, which is worse than the running time guaranteed by Theorem 5.4.4. We will see in the next chapter however that there are cases where $c \ll t$.

The results developed above also apply to the algorithm of Theorem 5.3.12 and we thus improve on the running time of Theorem 5.4.9 as follows.

Theorem 5.4.10 *One can compute (the cardinality of) a minimum connected dominating set of a unit disk graph G with representation $\mathcal{D}$ in $\mathrm{O}(n \cdot c \log c \cdot 2^{3c} \cdot 2^{2t})$ time, where $t = t^*(\mathcal{D})$ and c is the largest number of connected components in any $A_i \subseteq X_i$ in the strong path decomposition (T, X) corresponding to the slab decomposition supporting $t^*(\mathcal{D})$.*

This gives a worst-case running time of $\mathrm{O}(n \cdot t \log t \cdot 2^{5t})$.

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 $O(dh^2/\log h)$ 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. □

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 $(2 + 4/\pi)\sqrt{dn} + o(\sqrt{dn})$ 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 $O(n\sqrt{dn}\, 2^{6\sqrt{dn}})$ time, Minimum Dominating Set in $O(n\sqrt{dn}\, 3^{6\sqrt{dn}})$ time, and Minimum Connected Dominating Set in $O(dn^2\, 2^{22.8\sqrt{nd\log n}})$ 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 $O^*(2^{O(\sqrt{n})})$ 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 $O^*(2^{O(\sqrt{n})})$ 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 $n^{O(\sqrt{k})}$-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 $2^{O(\sqrt{k})}$ 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. □

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 $\mathcal{D}_a^b$,

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

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

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

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

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

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

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

it follows that

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

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

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 $k \geq 1$) 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 $k \geq 3$ 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 $bk + a \leq i \leq (b+1)k + a + 1$, 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 $0 \leq a \leq k-1$ 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. □

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

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. □

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 $P = p_1 \ldots p_m$ 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 $vp_3 \ldots p_m$ is a shorter Y–Z path than P, a contradiction. Hence $C^A \cup C^B$ is connected. □

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 $\mathcal{D}^b_a$ to be the set of disks in the b-th strip induced by shifting parameter a. Let $S^b_a = (\mathcal{D}^{b-1}_a \cap \mathcal{D}^b_a) \cup (\mathcal{D}^b_a \cap \mathcal{D}^{b+1}_a)$ and let N^b_a 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^b_a dominating $\mathcal{D}^b_a - N^b_a$ such that $C^b_a \cap Z$ is connected for each connected component Z of $\mathcal{D}^b_a$ 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 $0 \leq a \leq k-1$, $C_a = \bigcup_{b \in \mathbb{Z}} C^b_a$ is a connected dominating set, since $\{\mathcal{D}^b_a, \mathcal{D} - (\mathcal{D}^b_a - S^b_a)\}$ 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 $b \in \mathbb{Z}$, $|C^b_a| \leq |\mathcal{C} \cap \mathcal{D}^b_a| + 2|\mathcal{C} \cap N^b_a|$. As $C_a = \bigcup_{b \in \mathbb{Z}} C^b_a$,

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

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

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

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

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. □

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

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 $O(\text{poly}(n, 1/\epsilon)\, d^{O(\epsilon^{-l-1})})$ 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 $O(n^{o(1)})$. 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 $O(\text{poly}(n, 1/\epsilon)\, 2^{O(d/\epsilon)})$ 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 $o(\log n)$. 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 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 $d = d(n) = o(\log n)$.

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 $0 \le \alpha \le 1$ 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. □

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 $\mathrm{tw}(L_a^b) = O_H(k)$ *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 $\mathrm{tw}(\hat{L}_a^b) = O_H(k)$. As L_a^b is a minor of $\hat{L}_a^b$, $\mathrm{tw}(L_a^b) \leq \mathrm{tw}(\hat{L}_a^b)$ and thus $\mathrm{tw}(L_a^b) = O_H(k)$. □

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 $\mathrm{bw}(G) \geq w$ 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. □

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. □

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.

Chapter 7

Better Approximation Schemes on Disk Graphs

In the previous chapter, we considered unit disk graphs of bounded density, leading to new approximation schemes for several optimization problems. Here we extend these ideas to disk graphs and introduce the notion of bounded level density. We give an eptas for Maximum Independent Set on disk graphs of bounded level density, which is also a ptas on arbitrary disk graphs. Furthermore, we show that there is an eptas for Minimum Vertex Cover on arbitrary disk graphs, improving results of Erlebach, Jansen, and Seidel [103]. The given description of these schemes also establishes a general framework, making it easier to obtain efficient approximation schemes for other problems. We will in fact see further applications of this framework in later chapters.

These results all form a geometric generalization to the schemes for planar graphs obtained by Baker [22], because each planar graph is a disk graph of ply 1 [169, 210], and thus a disk graph of bounded level density as well.

7.1 The Ply of Disk Graphs

Let $\mathcal{D} = \{D_i \mid i = 1, \ldots, n\}$ be a set of disks in the plane and $G = (V, E)$ the corresponding disk graph. Scale the disks by a factor 2^w for some integer w, such that each disk has radius at least $\frac{1}{2}$. In the following, we will not distinguish between the disks in $\mathcal{D}$ and the vertices of the graph they induce.

Previously, we showed that an eptas exists for Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set on unit disk graphs of bounded density. The density of a unit disk graph is (informally) the maximum number of disk centers in any 1×1 box. A careful examination of the proof of these schemes showed that they can be extended to disk graphs of bounded density and constant maximum radius, but do not generalize to disk graphs of arbitrary density and radius. Hence another approach is needed.

The *ply* of a point p in the plane with respect to $\mathcal{D}$ is the number of disks of $\mathcal{D}$ strictly containing p (i.e. having p strictly inside the disk). Then the *ply* of $\mathcal{D}$ is the maximum ply of any point in the plane [210]. Observe that disk graphs of bounded ply are more general than disk graphs of bounded density and bounded maximum radius. Hence an eptas on disk graphs of bounded ply would generalize previous results. Below we give such an approximation scheme for Minimum Vertex Cover. The analysis relies heavily on the following properties of disk graphs of bounded ply.

Lemma 7.1.1 *Given a set $\mathcal{D}$ of disks of ply γ, the number of disks of radius at least r intersecting*

- *a line of length k is at most $\frac{4}{r\pi}(k+4r)\gamma$,*
- *the boundary of a $k \times k$ square ($k \geq 4r$) is at most $\frac{16}{r\pi}k\gamma$,*
- *a $k \times k$ square is at most $\frac{(k+4r)^2}{r^2\pi}\gamma$,*
- *two perpendicular, intersecting lines of length k is at most $\frac{8}{r\pi}(k+2r)\gamma$.*

Proof (Sketch): Consider a line of length k. Replace each disk D of radius at least r intersecting the line by a canonical disk D' of radius precisely r, such that D' intersects the line and $D' \subseteq D$. Any such canonical disk is contained in a size $(4r) \times (k+4r)$ rectangle centered over the line. As the canonical disks have ply at most γ and each has area $r^2\pi$, one can readily see that at most $\frac{4}{r\pi}(k+4r)\gamma$ disks intersect the line. The other bounds follow similarly. □

Lemma 7.1.2 *A set $\mathcal{D}$ of disks of ply γ, radius at least r and at most r', and intersecting a $k \times k$ square, has a path decomposition of width at most $\frac{4}{r\pi}(k+2r'+4r)\gamma - 1$ and consisting of at most $\frac{(k+4r)^2}{r^2\pi}\gamma$ bags.*

Proof (Sketch): Sweep a vertical line of length $k+2r'$ through the square from left to right. At any position of the line, place the disks intersecting the line in a bag. This yields a valid path decomposition. Moreover, one can find such a decomposition in $O(|\mathcal{D}| \log |\mathcal{D}|)$ time [258]. The bounds follow from the previous lemma. □

7.2 Approximating Minimum Vertex Cover

To approximate the minimum vertex cover problem, we use the *(geometric) shifting technique* introduced by Hochbaum and Maass [150]. To apply this technique, a decomposition of the minimum vertex cover problem into smaller subproblems is needed. Here we use a decomposition of the disks similar to the ones proposed by Hochbaum and Maass [150], Erlebach, Jansen, and Seidel [103], and Chan [57]. Combining the shifting technique with this decomposition yields the desired approximation factor (see Section 7.2.4).

First partition the disks into levels. A disk has *level* $j \in \mathbb{Z}_{\geq 0}$ if its radius r satisfies $2^{j-1} \leq r < 2^j$. Since all disks have radius at least $\frac{1}{2}$, each disk is indeed assigned a level. The level of the largest disk is denoted by l. For a set of disks $\mathcal{D}$, let $\mathcal{D}_{=j}$ denote the set of disks in $\mathcal{D}$ of level j. Similarly, we define $\mathcal{D}_{\geq j}$ as the set of disks of level at least j, and so on. Finally, $\mathcal{D}_{>j,<j'}$ is the set of disks of level greater than j, but less than j'.

Now let $k \geq 5$ be an odd positive integer (whose precise value is determined later). For each level j, we decompose the plane into squares of size $k2^j \times k2^j$ such that these squares induce a quadtree. Formally, for each level j, we consider the horizontal lines $y = hk2^j$ and vertical lines $x = vk2^j$ $(h, v \in \mathbb{Z})$. The squares induced by these lines are called *level j squares*, or put simply, *j-squares*.

Note that each j-square is completely contained in some $(j+1)$-square. Conversely, each $(j+1)$-square S contains exactly four j-squares, denoted by S_1 through S_4. The squares $S_1, \ldots, S_4$ are *siblings* of each other. We let $\mathcal{D}^S$ denote the set of disks intersecting S and $\mathcal{D}^{\mathrm{b}(S)}$ denotes the set of disks which intersect the boundary of S. Furthermore, we define $\mathcal{D}^{\mathrm{i}(S)} = \mathcal{D}^S - \mathcal{D}^{\mathrm{b}(S)}$ (i.e. the set of disks fully contained in the interior of S) and let $\mathcal{D}^{+(S)} = \mathcal{D}^{\mathrm{i}(S)} - \bigcup_{i=1}^4 \mathcal{D}^{\mathrm{i}(S_i)} = \bigcup_{i=1}^4 \mathcal{D}^{\mathrm{b}(S_i)} - \mathcal{D}^{\mathrm{b}(S)}$ (i.e. the set of disks intersecting the boundary of at least one of the four children of S, but not the boundary of S itself). The meaning of combinations like $\mathcal{D}^{\mathrm{b}(S)}_{\leq j}$ should be self-explaining. We use $j(S)$ to denote the level of a square S.

7.2.1 A Close to Optimal Vertex Cover

We prove the following theorem, which will be auxiliary to the main theorem.

Theorem 7.2.1 *Let $\mathcal{D}$ be a set of n disks of ply γ and $k \geq 5$ an odd positive integer. Then in time $\mathrm{O}(k^2 n^2 \gamma^{64k/\pi})$, one can find a vertex cover VC of $\mathcal{D}$ such that $|VC| \leq \sum_S \left| OPT^S_{=j(S)} \right|$, where the sum is over all squares S and OPT is any minimum vertex cover of $\mathcal{D}$.*

We can obtain a vertex cover of the required cardinality by applying bottom-up dynamic programming to the j-squares. Roughly speaking, for each j-square S, we consider all subsets of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ (the disks of level greater than j intersecting the boundary of S). For each such subset, we compute a close to optimal vertex cover for $\mathcal{D}^S$ containing this subset. Formally, we define for each j-square S and each $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$ a function $\mathsf{size}(S, W)$. The function is defined recursively on j.

$$\mathsf{size}(S, W) = \begin{cases} \min\left\{ |T| \,\middle|\, T \subseteq \mathcal{D}^S_{=j} \cup \mathcal{D}^{\mathrm{i}(S)}_{>j};\ T \cup W \text{ covers } \mathcal{D}^S \right\} & \text{if } j = 0; \\ \min\limits_{U \subseteq \mathcal{D}^{+(S)}_{\geq j} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j}} \left\{ |U| + \sum_{i=1}^{4} \mathsf{size}\big(S_i, (U \cup W)^{\mathrm{b}(S_i)}\big) \right\} & \text{if } j > 0. \end{cases}$$

Here we define the minimum over an empty set to be ∞. Observe that W must be a vertex cover of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ and U must be a vertex cover of $\mathcal{D}^{+(S)}_{\geq j} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j}$. Let $\mathsf{sol}(S, W)$ be the subfamily of $\mathcal{D}$ attaining $\mathsf{size}(S, W)$, or $\emptyset$ if $\mathsf{size}(S, W)$ is ∞.

7.2.2 Properties of the size- and sol-Functions

We first show that the sum of $\mathsf{size}(S, \emptyset)$ over all level l squares S attains the value stated in Theorem 7.2.1. In fact, we prove a slightly more general result.

Let $\mathcal{C}$ be any vertex cover for $\mathcal{D}$.

Lemma 7.2.2 $\sum_{S;\,j(S)=l} \mathsf{size}(S,\emptyset) \le \sum_S \left|\mathcal{C}^S_{=j(S)}\right|$.

Proof: Apply induction on j. We prove that the following invariant holds:

$$\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}\right) \le \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \sum_{S' \subseteq S} \left|\mathcal{C}^{S'}_{=j(S')}\right|.$$

Here S is some j-square. For $j = 0$, the correctness of the invariant follows from the definition of size. So assume that $j > 0$ and that the invariant holds for all j'-squares with $j' < j$. Note that

$$\sum_{i=1}^{4} \left|\mathcal{C}^{\mathrm{i}(S_i)}_{>j-1}\right| = \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \left|\mathcal{C}^{\mathrm{i}(S)}_{=j}\right| - \left|\mathcal{C}^{+(S)}_{\ge j}\right|.$$

Then from the description of size and by applying induction,

$$\begin{aligned}
&\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}\right)\\
\le\ & \left|\mathcal{C}^{+(S)}_{\ge j}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \left(\mathcal{C}^{\mathrm{b}(S)}_{\ge j} \cup \mathcal{C}^{+(S)}_{\ge j}\right)^{\mathrm{b}(S_i)}\right)\\
=\ & \left|\mathcal{C}^{+(S)}_{\ge j}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \mathcal{C}^{\mathrm{b}(S_i)}_{>j-1}\right)\\
\le\ & \left|\mathcal{C}^{+(S)}_{\ge j}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \left|\mathcal{C}^{\mathrm{i}(S_i)}_{>j-1}\right| + \sum_{i=1}^{4} \sum_{S'_i \subseteq S_i} \left|\mathcal{C}^{S'_i}_{=j(S'_i)}\right|\\
=\ & \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \left|\mathcal{C}^{S}_{=j}\right| + \sum_{i=1}^{4} \sum_{S'_i \subseteq S_i} \left|\mathcal{C}^{S'_i}_{=j(S'_i)}\right|\\
=\ & \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \sum_{S' \subseteq S} \left|\mathcal{C}^{S'}_{=j(S')}\right|.
\end{aligned}$$

Since l is the level of the largest disk, $\mathcal{C}^{\mathrm{i}(S)}_{>j} = \emptyset$ and $\mathcal{C}^{\mathrm{b}(S)}_{>j} = \emptyset$ for all j-squares S with $j \ge l$. Hence

$$\sum_{S;\,j(S)=l} \mathsf{size}(S,\emptyset) \le \sum_{S;\,j(S)=l} \sum_{S' \subseteq S} \left|\mathcal{C}^{S'}_{=j(S')}\right| = \sum_S \left|\mathcal{C}^S_{=j(S)}\right|.$$

This proves the lemma. □

The lemma implies that $\sum_{S;\,j(S)=l} \mathsf{size}(S,\emptyset) \le \sum_S \left|OPT^S_{=j(S)}\right|$, where OPT is a minimum vertex cover of $\mathcal{D}$. We now prove that the union of $\mathsf{sol}(S,\emptyset)$ over all level l squares S is a vertex cover of $\mathcal{D}$.

Lemma 7.2.3 $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S, \emptyset)$ *is a vertex cover of* $\mathcal{D}$.

Proof: For any level j_0 and any collection of sets $\{W_S \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j_0} \mid j(S) = j_0\}$, we prove the following claim:

$$\bigcup_{S;j(S)=j_0} \mathsf{sol}(S, W_S) \cup W_S \text{ covers } \bigcup_{S;j(S)=j_0} \mathcal{D}^S \text{ if } \sum_{S;j(S)=j_0} \mathsf{size}(S, W_S) \neq \infty.$$

Apply induction on j_0. For $j_0 = 0$, this follows trivially from the definition of size and sol. So assume that $j_0 > 0$ and that the claim holds for all $j_0' < j_0$.

Suppose $\sum_{S;j(S)=j_0} \mathsf{size}(S, W_S) \neq \infty$ is true for some collection of sets $\{W_S \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j_0} \mid j(S) = j_0\}$. For any j_0-square S, let

$$U_S^* = \arg \min_{U \subseteq \mathcal{D}^{+(S)}_{\geq j_0} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j_0}} \left\{ |U| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, (U \cup W_S)^{\mathrm{b}(S_i)}\right) \right\}.$$

As $\mathsf{size}(S, W_S) \neq \infty$, it must be that $\mathsf{size}(S_i, (U_S^* \cup W_S)^{\mathrm{b}(S_i)}) \neq \infty$ for $i = 1, \ldots, 4$ as well. For any S' where $S' = S_i$ for some j_0-square S and $i \in \{1, \ldots, 4\}$, let $W_{S'} = (U_S^* \cup W_S)^{\mathrm{b}(S')}$. It follows that

$$\sum_{S';j(S')=j_0-1} \mathsf{size}(S', W_{S'}) \neq \infty.$$

Then by induction, $\bigcup_{S';j(S')=j_0-1} \mathsf{sol}(S', W_{S'}) \cup W_{S'}$ covers $\bigcup_{S';j(S')=j_0-1} \mathcal{D}^{S'}$. Observe that for any j_0-square S

$$\begin{aligned} W_S \cup \mathsf{sol}(S, W_S) &= W_S \cup U_S^* \cup \bigcup_{i=1}^{4} \mathsf{sol}(S_i, (U_S^* \cup W_S)^{\mathrm{b}(S_i)}) \\ &= \bigcup_{i=1}^{4} \left((U_{S_i}^* \cup W_{S_i})^{\mathrm{b}(S_i)} \cup \mathsf{sol}\left(S_i, (U_{S_i}^* \cup W_{S_i})^{\mathrm{b}(S_i)}\right)\right) \\ &= \bigcup_{\substack{S'=S_i \\ i=1,\ldots,4}} W_{S'} \cup \mathsf{sol}(S', W_{S'}). \end{aligned}$$

As $\bigcup_{S';j(S')=j_0-1} \mathcal{D}^{S'} = \bigcup_{S;j(S)=j_0} \mathcal{D}^S$, we have $\bigcup_{S;j(S)=j_0} \mathsf{sol}(S, W_S) \cup W_S$ covers $\bigcup_{S;j(S)=j_0} \mathcal{D}^S$.

From the previous lemma, we know that $\sum_{S;j(S)=l} \mathsf{size}(S, \emptyset) \neq \infty$. Because each edge is induced by $\mathcal{D}^S$ for some l-square S, $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S, \emptyset)$ is a vertex cover of $\mathcal{D}$. □

7.2.3 Computing the size- and sol-Functions

We show that it is sufficient to compute size and sol for a limited number of j-squares. This can be done in the time stated in Theorem 7.2.1.

Call a j-square *nonempty* if it is intersected by a level j disk and *empty* otherwise. A j-square S is *relevant* if one of its three siblings is nonempty or there is a nonempty square S' containing S, such that S' has level at most $j + \lceil \log k \rceil$ (so each nonempty j-square is relevant). Note that this definition induces $O(k^2 n)$ relevant squares. A relevant square S is said to be a *relevant child* of another relevant square S' if $S \subset S'$ and there is no third relevant square S'', such that $S \subset S'' \subset S'$. Conversely, if S is a relevant child of S', S' is a *relevant parent* of S.

Lemma 7.2.4 *For each relevant* 0*-square* S*, all* size*- and* sol*-values for* S *can be computed in* $O\left(nk^3\gamma\,\gamma^{(24k+8)/\pi}\right)$ *time.*

Proof: From Lemma 7.1.1, $\left|\mathcal{D}^{\mathrm{b}(S)}_{>0}\right|$ is bounded by $16k\gamma/\pi$. As an independent set has $\gamma = 1$, all independent sets and hence all vertex covers of $\mathcal{D}^{\mathrm{b}(S)}_{>0}$ can be enumerated in $O(k\,\gamma^{16k/\pi})$ time using Lemma 6.3.2. For a fixed set W, $\mathsf{size}(S, W)$ is defined as the cardinality of a minimum subset of $\mathcal{D}^{S}_{=0} \cup \mathcal{D}^{\mathrm{i}(S)}_{>0}$, such that this subset and W cover $\mathcal{D}^S$. We may assume that W covers $\mathcal{D}^{\mathrm{b}(S)}_{>0}$, otherwise such a subset does not exist and $\mathsf{size}(S, W)$ is ∞. Then the requested subset is a minimum vertex cover for $\mathcal{D}^S - W$. Similar to Lemma 7.1.2, one can show that $\mathcal{D}^S$ has a path decomposition of width at most $\frac{8}{\pi}(k + 4)\gamma$ and $O(|\mathcal{D}^S|)$ bags. Moreover, these path decompositions can be precomputed for all level 0 squares in $O(n \log n)$ time. Adapting the algorithm of Lemma 5.3.4 and using Lemma 6.3.2, the cover can be computed in $O(|\mathcal{D}^S|k^2\gamma\,\gamma^{8(k+4)/\pi})$ time. Therefore one can compute all size- and sol-values for S in $O\left(nk^3\gamma\,\gamma^{(24k+8)/\pi}\right)$ time. □

Assume that the size- and sol-values of all relevant children of S are known.

Lemma 7.2.5 *For each relevant* j*-square* S $(j > 0)$ *with relevant* $(j-1)$*-square children, all* size*- and* sol*-values for* S *can be computed in* $O(k\,\gamma^{64k/\pi})$ *time.*

Proof: If one of the children $S_1, \ldots, S_4$ of S is relevant, then, by the definition of relevant, all children of S must be relevant. Following the definition of size, we enumerate all vertex covers W of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ and for each such W all vertex covers U of $\mathcal{D}^{+(S)}_{\geq j} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j}$. Using the ideas of Lemma 7.1.1, we can show that $\left|\mathcal{D}^{\mathrm{b}(S)}_{>j}\right| \leq 16k\gamma/\pi$, $\left|\mathcal{D}^{+(S)}_{\geq j} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j}\right| \leq 48k\gamma/\pi$, and $\left|\mathcal{D}^{\mathrm{b}(S)}_{\geq j} \cup \mathcal{D}^{+(S)}_{\geq j}\right| \leq 48k\gamma/\pi$. Then all independent sets, and hence all vertex covers, of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ and of $\mathcal{D}^{+(S)}_{\geq j} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j}$ can be enumerated in $O(k\,\gamma^{64k/\pi})$ time by applying Lemma 6.3.2. Since size and sol of all relevant children of S are known and assuming that for a given W and U we can compute $|U| + \sum_{i=1}^{4} \mathsf{size}(S_i, (U \cup W)^{\mathrm{b}(S_i)})$ in constant time, the running time of $O(k\,\gamma^{64k/\pi})$ follows immediately. □

Lemma 7.2.6 *For each relevant* j*-square* S $(j > 0)$ *with no relevant children of level* $j-1$*, all* size*- and* sol*-values for* S *can be computed in* $O(n\,\gamma^{32/\pi})$ *time.*

Proof: We start with two simple observations. The first is that S must be empty, because S has no relevant children of level $j-1$. Secondly, one notes that by the definition of relevant, the nearest nonempty square containing S (if it exists) has level at least $j + \lceil \log k \rceil$. Hence $\mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}$.

Now consider any j'-square $S' \subseteq S$ for which there is no relevant square S'' such that $S' \subseteq S'' \subset S$. Then the nearest nonempty square containing S' (if it exists) has level at least $j + \lceil \log k \rceil$. Hence any disk of level at least j' intersecting S' has level at least $j + \lceil \log k \rceil$. This implies that $\mathcal{D}^{\mathrm{b}(S'_i)}_{>j'-1} = \mathcal{D}^{\mathrm{b}(S')}_{>j'}$ for any $i = 1, \ldots, 4$ and that $\mathcal{D}^{\mathrm{b}(S')}_{=j'} = \emptyset$. Since S' is empty, it also follows that $\mathcal{D}^{\mathrm{i}(S')}_{\geq j'} = \emptyset$ and, if $j' > 0$, $\mathcal{D}^{+(S')}_{\geq j'} = \emptyset$ as well.

Using these observations, we can simplify the definition of size considerably for such S'. For any set $W' \subseteq \mathcal{D}^{\mathrm{b}(S')}_{>j'}$,

$$\mathsf{size}(S', W') = \begin{cases} 0 & \text{if } j' = 0,\ W' \text{ covers } \mathcal{D}^{\mathrm{b}(S')}_{>j'}; \\ \infty & \text{if } j' = 0,\ W' \text{ not covers } \mathcal{D}^{\mathrm{b}(S')}_{>j'} \\ \sum_{i=1}^{4} \mathsf{size}\left(S'_i, W'^{\mathrm{b}(S'_i)}\right) & \text{if } j' > 0. \end{cases}$$

Applying this simplification repeatedly, it can be seen that for any $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$,

$$\mathsf{size}(S, W) = \begin{cases} 0 & \text{if } S \text{ has no relevant children and} \\ & \qquad W \text{ covers } \mathcal{D}^{\mathrm{b}(S)}_{>j}; \\ \infty & \text{if } S \text{ has no relevant children and} \\ & \qquad W \text{ doesn't cover } \mathcal{D}^{\mathrm{b}(S)}_{>j}; \\ \sum_{S''} \mathsf{size}\left(S'', W^{\mathrm{b}(S'')}\right) & \text{otherwise,} \end{cases}$$

where the sum is over all relevant children S'' of S.

Any relevant child of S is either nonempty, or the sibling of a nonempty square. As the number of nonempty squares is $\mathrm{O}(n)$ and a square has three siblings, the number of relevant children of S is $\mathrm{O}(n)$. So for fixed W, it takes $\mathrm{O}(n)$ time to compute $\mathsf{size}(S, W)$. As $\mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}$, we know from Lemma 7.1.1 and Lemma 6.3.2 that all vertex covers $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}$ can be enumerated in $\mathrm{O}(\gamma^{32/\pi})$ time. □

Lemma 7.2.7 $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset)$ *can be computed in* $\mathrm{O}(k^2 n^2 \gamma^{64k/\pi})$ *time.*

Proof: Recall that there are $\mathrm{O}(k^2 n)$ relevant squares. Let S be a relevant j-square without a relevant parent. Following Lemmas 7.2.4, 7.2.5, and 7.2.6, we can compute $\mathsf{size}(S, \emptyset)$ for all such squares S in $\mathrm{O}(k^2 n^2 \gamma^{64k/\pi})$ time.

Now consider any level l square S. If S is relevant, then it cannot have a relevant parent. Hence by the preceding argument, $\mathsf{size}(S, \emptyset)$ is known. If S is not relevant, then we can use the same arguments as in Lemma 7.2.6 to show that $\mathsf{size}(S, \emptyset) = \sum_{S''} \mathsf{size}(S'', \emptyset)$, where the sum is over all relevant j''-squares $S'' \subset S$ without a relevant parent. It follows that $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset)$ can be computed in $\mathrm{O}(k^2 n^2 \gamma^{64k/\pi})$ time. □

Proof of Theorem 7.2.1: Follows by Lemmas 7.2.2, 7.2.3, and 7.2.7. □

7.2.4 An eptas for Minimum Vertex Cover

We now apply the shifting technique to obtain a $(1+\epsilon)$ approximation of the optimum. For some integer a $(0 \le a \le k-1)$, define the decomposition as follows. We call a line of level j *active* if it is of the form $y = (hk + a2^{l-j})2^j$ or $x = (vk + a2^{l-j})2^j$ $(h, v \in \mathbb{Z})$. The active lines partition the plane into j-squares as before, except that they are now shifted by the shifting parameter a. The structure however remains the same, and thus we can apply Theorem 7.2.1 to compute a close to optimal vertex cover.

Let VC_a denote the set returned by the algorithm for some value of a $(0 \le a \le k-1)$ and let $VC_{\min}$ be a smallest such set.

Lemma 7.2.8 $|VC_{\min}| \le (1 + \frac{12}{k})\,|OPT|$.

Proof: We claim a line of level j (i.e. of the form $y = h'2^j$ or $x = v'2^j$) is active for precisely one value of a. A horizontal line $y = h'2^j$ is active if $h' = hk + a2^{l-j}$ for some h and a, i.e. if $h' \equiv a2^{l-j} \bmod k$. As $\gcd(k, 2^{l-j}) = 1$, such a value of a exists. Hence the line is active for at least one value of a.

Suppose that a horizontal line of level j is active for two values of a. Then $hk + a2^{l-j} = h'k + a'2^{l-j}$ for some choice of h, h', a, and a'. Simplifying gives $(h - h')k = (a' - a)2^{l-j}$, or $k|(a' - a)2^{l-j}$. Since k is odd, $k|(a' - a)$, which is impossible as $1 \le |a' - a| \le k - 1$. Hence each horizontal line of level j is active for precisely one value of a. The same arguments hold for vertical lines of level j.

Define $\mathcal{D}^b_a$ as the set of disks intersecting the boundary of a j-square S at their level, i.e. $\mathcal{D}^b_a = \bigcup_S \mathcal{D}^{\mathrm{b}(S)}_{=j(S)}$. A level j disk is in $\mathcal{D}^b_a$ if and only if it intersects an active line of level j. It can be in $\mathcal{D}^b_a$ for at most four different values of a, intersecting both a horizontal and a vertical active line at most twice, because a line of level j is active for exactly one value of a, the distance between consecutive lines is 2^j, and disks of level j have radius less than 2^j. Hence there is a value of a (say a^*) for which $\left|OPT \cap \mathcal{D}^b_{a^*}\right| \le \frac{4}{k}\,|OPT|$.

From Lemma 7.2.2, we know that $|VC_{a^*}| \le \sum_S |OPT^S_{=j(S)}|$. Observe that for a fixed value of a, any disk can intersect at most four squares at its level. Then

$$\begin{aligned} |VC_{a^*}| &\le \sum_S \left|OPT^S_{=j(S)}\right| \\ &= \sum_S \left|OPT^S_{=j(S)} - OPT^{\mathrm{b}(S)}_{=j(S)}\right| + \sum_S \left|OPT^{\mathrm{b}(S)}_{=j(S)}\right| \end{aligned}$$

$$\begin{aligned} &\leq |OPT| - \left|OPT \cap \mathcal{D}^b_{a^*}\right| + 4\left|OPT \cap \mathcal{D}^b_{a^*}\right| \\ &\leq |OPT| + \tfrac{12}{k}\,|OPT|\,. \end{aligned}$$

Hence $|VC_{\min}| \leq |VC_{a^*}| \leq (1 + \frac{12}{k})\,|OPT|$ and the lemma follows. □

Combining Theorem 7.2.1 and Lemma 7.2.8, we obtain the following result.

Theorem 7.2.9 *There is an eptas for Minimum Vertex Cover on disk graphs.*

Proof: The idea is similar to Lemma 6.3.14 and Theorem 6.3.15, except we have no polynomial-time algorithm to find a maximum clique in a disk graph. However, it suffices to reduce the ply. Consider a point p in the plane of ply more than $\frac{1}{\epsilon}$. Note that the set of disks $\mathcal{D}_p$ containing p form a clique. Marx [202] observed that $\mathcal{D}_p$ is actually a $(1+\epsilon)$-approximation of a minimum vertex cover for $\mathcal{D}_p$. Hence we remove $\mathcal{D}_p$ from $\mathcal{D}$ and repeat until the ply is bounded by $\frac{1}{\epsilon}$. Using the algorithm by Eppstein, Miller, and Teng [100] to determine the ply of a set of disks, this can be done in $O(n^3 \log n)$ time.

Let $\mathcal{D}_0$ denote the remaining set of disks. Choose k as the smallest odd integer larger than $\frac{12}{\epsilon}$. Compute and output $VC_{\min}$ in $O(k^3 n^2\, k^{64k/\pi})$ time using Theorem 7.2.1. Following Lemma 7.2.8 and the choice of k, this results in a $(1+\epsilon)$-approximation of a minimum vertex cover of $\mathcal{D}_0$. Combining the different approximations gives a $(1+\epsilon)$-approximation of a minimum vertex cover of $\mathcal{D}$. This gives the eptas. □

This result improves the $n^{O(\epsilon^{-2})}$-time ptas for Minimum Vertex Cover on disk graphs by Erlebach, Jansen, and Seidel [103].

7.3 Approximating Maximum Independent Set

The maximum independent set problem can also be approximated well using the ideas of the previous sections. We can show that it has an eptas on disk graphs of bounded ply. In fact, we can prove a more general result, namely that Maximum Independent Set has an eptas on disk graphs of bounded level density. This notion is defined as follows. Partition the disks into levels as before (i.e. a disk has level j if its radius is in $[2^{j-1}, 2^j)$). For each level j, let d_j denote the maximum number of level j disks in any $2^j \times 2^j$ box. Then the *level density*, denoted by d, is the maximum d_j over all levels j. Scaling a set of disks by a constant factor can reduce the level density by a factor of 2, but this is of little consequence to the analysis of the algorithm below.

Disk graphs of bounded level density are more general than disk graphs of bounded ply, as a disk graph of ply γ has level density at most 4γ. However, a disk graph of bounded level density can contain overlapping disks from an arbitrary number of levels, giving it arbitrarily large ply.

Consider a set of disks $\mathcal{D}$ of level density d. Let $k \geq 5$ be an odd positive integer to be determined later and let the plane be partitioned into j-squares as before. We prove the following auxiliary theorem.

Theorem 7.3.1 *Let $\mathcal{D}$ be a set of disks of level density d and $k \geq 5$ an odd positive integer. Then one can find in $\mathrm{O}(k^3 n^9 (2ed)^{32k/\pi})$ time an independent set IS of $\bigcup_S \mathcal{D}^{\mathrm{i}(S)}_{=j(S)}$ such that $|IS| \geq \sum_S \left| OPT^{\mathrm{i}(S)}_{=j(S)} \right|$, where sum and union are over all squares S and OPT is any maximum independent set of $\mathcal{D}$.*

We employ a similar approach as with Minimum Vertex Cover. For any j-square S and any independent set $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$, we compute (the cardinality of) a close to maximum independent set of $\mathcal{D}^{\mathrm{i}(S)}_{>j} \cup \bigcup_{S' \subseteq S} \mathcal{D}^{\mathrm{i}(S')}_{=j(S')}$ that is independent of W. For each j-square S and each independent set $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$,

$$\mathsf{size}(S, W) = \begin{cases} \max\left\{ |T| \;\middle|\; T \subseteq \mathcal{D}^{\mathrm{i}(S)};\ T \cup W \text{ independent} \right\} & \text{if } j = 0; \\ \max\limits_{U \subseteq \mathcal{D}^{+(S)}_{\geq j}} \left\{ |U| + \sum\limits_{i=1}^{4} \mathsf{size}\Big(S_i, (U \cup W)^{\mathrm{b}(S_i)}\Big) \right\} & \text{if } j > 0. \end{cases}$$

Let $\mathsf{sol}(S, W)$ be the subset of $\mathcal{D}$ attaining $\mathsf{size}(S, W)$.

Lemma 7.3.2 $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset) \geq \sum_S \left| \mathcal{I}^{\mathrm{i}(S)}_{=j(S)} \right|$ *for any independent set $\mathcal{I}$.*

Proof: Define

$$up(S) = \bigcup_{S' \supset S} \mathcal{I}^{\mathrm{i}(S')}_{=j(S')}.$$

We use induction on j to prove the following invariant for any j-square S:

$$\mathsf{size}\left(S, (up(S))^{\mathrm{b}(S)}\right) \geq \left| \mathcal{I}^{\mathrm{i}(S)}_{>j} \right| + \sum_{S' \subseteq S} \left| \mathcal{I}^{\mathrm{i}(S')}_{=j(S')} \right|.$$

It follows immediately from the definition of size that the invariant is true for $j = 0$. So consider a $j > 0$ and assume the invariant holds for all $j' < j$. Then

$$\begin{aligned} &\mathsf{size}\left(S, (up(S))^{\mathrm{b}(S)}\right) \\ &\quad \geq \left| \mathcal{I}^{+(S)}_{\geq j} \right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \left(\mathcal{I}^{+(S)}_{\geq j} \cup (up(S))^{\mathrm{b}(S)}\right)^{\mathrm{b}(S_i)}\right) \\ &\quad = \left| \mathcal{I}^{+(S)}_{\geq j} \right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, (up(S_i))^{\mathrm{b}(S_i)}\right) \\ &\quad \geq \left| \mathcal{I}^{+(S)}_{\geq j} \right| + \sum_{i=1}^{4} \left| \mathcal{I}^{\mathrm{i}(S_i)}_{>j-1} \right| + \sum_{i=1}^{4} \sum_{S'_i \subseteq S_i} \left| \mathcal{I}^{\mathrm{i}(S'_i)}_{=j(S'_i)} \right| \\ &\quad = \left| \mathcal{I}^{\mathrm{i}(S)}_{>j} \right| + \sum_{S' \subseteq S} \left| \mathcal{I}^{\mathrm{i}(S')}_{=j(S')} \right|, \end{aligned} \tag{7.1}$$

where the equality

$$\left(\mathcal{I}^{+(S)}_{\geq j} \cup (up(S))^{\mathrm{b}(S)}\right)^{\mathrm{b}(S_i)} = (up(S_i))^{\mathrm{b}(S_i)}$$

holds, because by definition

$$\left(\mathcal{I}^{+(S)}_{\geq j}\right)^{\mathrm{b}(S_i)} = \left(\mathcal{I}^{\mathrm{i}(S)}_{\geq j}\right)^{\mathrm{b}(S_i)} = \left(\left(\bigcup_{S' \supseteq S} \mathcal{I}^{\mathrm{i}(S')}_{=j(S')}\right)^{\mathrm{i}(S)}\right)^{\mathrm{b}(S_i)}$$

and $\left(\mathcal{I}^{\mathrm{i}(S)}_{=j(S)}\right)^{\mathrm{b}(S)} = \emptyset$. Then

$$\begin{aligned}
&\left(\mathcal{I}^{+(S)}_{\geq j} \cup \left(\bigcup_{S' \supset S} \mathcal{I}^{\mathrm{i}(S')}_{=j(S')}\right)^{\mathrm{b}(S)}\right)^{\mathrm{b}(S_i)} \\
&\quad = \left(\left(\bigcup_{S' \supseteq S} \mathcal{I}^{\mathrm{i}(S')}_{=j(S')}\right)^{\mathrm{i}(S)} \cup \left(\bigcup_{S' \supseteq S} \mathcal{I}^{\mathrm{i}(S')}_{=j(S')}\right)^{\mathrm{b}(S)}\right)^{\mathrm{b}(S_i)} \\
&\quad = \left(\bigcup_{S' \supseteq S} \mathcal{I}^{\mathrm{i}(S')}_{=j(S')}\right)^{\mathrm{b}(S_i)} \\
&\quad = (up(S_i))^{\mathrm{b}(S_i)}.
\end{aligned}$$

Returning to Equation 7.1, as l is the level of the largest disk, $up(S) = \emptyset$ and $\mathcal{I}^{\mathrm{i}(S)}_{>j} = \emptyset$ for any square S of level at least l. Then

$$\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset) \geq \sum_{S;\, j(S)=l} \sum_{S' \subseteq S} \left|\mathcal{I}^{\mathrm{i}(S')}_{=j}\right| = \sum_{S} \left|\mathcal{I}^{\mathrm{i}(S)}_{=j(S)}\right|.$$

The lemma follows. □

Clearly, the lemma implies that $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset) \geq \sum_S \left|OPT^{\mathrm{i}(S)}_{=j(S)}\right|$, where OPT is a maximum independent set.

Lemma 7.3.3 $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S, \emptyset)$ *is an independent set of* $\bigcup_S \mathcal{D}^{\mathrm{i}(S)}_{=j(S)}$.

This lemma follows straightforwardly from the definitions of size and sol in a similar way as in Lemma 7.2.3.

To compute $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset)$, it is again sufficient to consider only relevant j-squares, where the definition of relevant is the same as before. As was observed earlier, we need only to consider independent sets W, T, and U in the definition of size, as size will be $-\infty$ otherwise. Crucial in the analysis of the algorithm will therefore be bounds on the maximum cardinality of certain independent sets. In particular, we apply the following theorem.

Theorem 7.3.4 *The maximum number of disjoint disks of radius* r *intersecting a square of size* $2r \times 2r$ *is* 7.

The (lengthy) proof is detailed in Section 7.5.

Lemma 7.3.5 *For each relevant* 0*-square* S*, all* size- *and* sol*-values can be computed in* $\mathrm{O}(k^2 n^8 d\,(2ed)^{24k/\pi})$ *time.*

Proof: As disks in $\mathcal{D}^{\mathrm{b}(S)}_{\geq \lceil \log k \rceil}$ have radius at least $\frac{1}{2}k$, we can use Theorem 7.3.4 to bound the maximum cardinality of any independent set in $\mathcal{D}^{\mathrm{b}(S)}_{\geq \lceil \log k \rceil}$ by 7. Hence all independent sets in $\mathcal{D}^{\mathrm{b}(S)}_{\geq \lceil \log k \rceil}$ can be enumerated in $\mathrm{O}(n^7)$ time.

To enumerate all independent subsets of $\mathcal{D}^{\mathrm{b}(S)}_{>0}$, we should consider independent subsets of $\mathcal{D}^{\mathrm{b}(S)}_{>0,<\lceil \log k \rceil}$ as well. For some j' with $0 < j' < \lceil \log k \rceil$, we can use an area bound to show that $\left|\mathcal{D}^{\mathrm{b}(S)}_{=j'}\right| \leq 4kd2^{j-j'}$. Then

$$\left|\mathcal{D}^{\mathrm{b}(S)}_{>0,<\lceil \log k \rceil}\right| \;\leq\; \sum_{j'=1}^{\lceil \log k \rceil - 1} \left(4kd2^{j-j'}\right) \;\leq\; 4kd.$$

As an independent set of disks has ply 1, it follows from Lemma 7.1.1 that any independent subset of $\mathcal{D}^{\mathrm{b}(S)}_{>0,<\lceil \log k \rceil}$ has cardinality at most $16k/\pi$. Then, following Lemma 6.3.2, all independent sets of disks in $\mathcal{D}^{\mathrm{b}(S)}_{>0,<\lceil \log k \rceil}$ can be enumerated in $\mathrm{O}(k\,(ed)^{16k/\pi})$ time. Hence all independent sets $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>0}$ can be enumerated in $\mathrm{O}(kn^7\,(ed)^{16k/\pi})$ time.

For fixed $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>0}$, it remains to compute a maximum $T \subseteq \mathcal{D}^{\mathrm{i}(S)}_{\geq 0}$ such that $T \cup W$ is an independent set. That is, to compute a maximum independent set of $\mathcal{D}^{\mathrm{i}(S)}_{\geq 0} - N[W]$, where $N[W]$ is the closed neighborhood of W. We use a path decomposition to find this set. First, observe that $\mathcal{D}^{\mathrm{i}(S)}_{\geq \lceil \log k \rceil} = \emptyset$. For any j' with $0 \leq j' < \lceil \log k \rceil$, the number of disks of $\mathcal{D}^{\mathrm{i}(S)}_{=j'}$ intersecting a vertical line of length k in S is bounded by $2\lceil (k - 2^{j'})/2^{j'} \rceil d$. Hence the number of disks of $\mathcal{D}^{\mathrm{i}(S)}_{\geq 0}$ intersecting such a line is at most

$$\sum_{j'=0}^{\lceil \log k \rceil - 1} \left(2\lceil 2^{-j'}k - 1\rceil d\right) \;\leq\; 4kd.$$

It follows that $\mathcal{D}^{\mathrm{i}(S)}_{\geq 0} - N[W]$ has a path decomposition of width at most $4kd$ and $\mathrm{O}(n)$ bags. Following Lemma 7.1.1, any independent set of disks in $\mathcal{D}^{\mathrm{i}(S)}_{\geq 0} - N[W]$ intersecting a vertical line in S has cardinality at most $8k/\pi$. Adapting the algorithm of Theorem 5.3.2 to only consider independent sets, the maximum T can be found in $\mathrm{O}(k^2 nd\,(2ed)^{8k/\pi})$ time. □

Lemma 7.3.6 *For each relevant* j*-square* S *with relevant* $(j-1)$*-square children, all* size- *and* sol*-values can be computed in* $\mathrm{O}(k^2 n^7\,(2ed)^{32k/\pi})$ *time.*

Proof: Using the same arguments as in the previous lemma, all independent subsets W of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ can be enumerated in $\mathrm{O}(kn^7\,(ed)^{16k/\pi})$ time. For any j' with $j-1 < j' < j + \lceil \log k \rceil$, $\left|\mathcal{D}^{+(S)}_{=j'}\right| \le (4\lceil 2^{j-j'}k - 1\rceil - 4)d$. Hence

$$\left|\mathcal{D}^{+(S)}_{>j-1}\right| \;\le\; \sum_{j'=j}^{j+\lceil \log k \rceil - 1} (4\lceil 2^{j-j'}k - 1\rceil - 4)d \;\le\; 8kd.$$

Following Lemma 7.1.1, any independent set of disks in $\mathcal{D}^{+(S)}_{\ge j}$ has cardinality at most $16k/\pi$. Then, for any fixed $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$, all independent sets $U \subseteq \mathcal{D}^{+(S)}_{\ge j} - N[W]$ can be enumerated in $\mathrm{O}(k\,(2ed)^{16k/\pi})$ time. □

Lemma 7.3.7 *For each relevant j-square S with no relevant $(j-1)$-square children, all* size *and* sol*-values can be computed in* $\mathrm{O}(n^8)$ *time.*

Proof: Using the same arguments as in Lemma 7.2.6, for any $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$,

$$\mathsf{size}(S, W) = \begin{cases} 0 & \text{if } S \text{ has no relevant children and } W \text{ is an independent set;} \\ -\infty & \text{if } S \text{ has no relevant children and } W \text{ is not an independent set;} \\ \sum_{S''} \mathsf{size}\left(S'', W^{\mathrm{b}(S'')}\right) & \text{otherwise,} \end{cases}$$

where the sum is over all relevant children S'' of S. Since the number of relevant children of S is $\mathrm{O}(n)$, for fixed W, it takes $\mathrm{O}(n)$ time to compute $\mathsf{size}(S, W)$. As $\mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\ge j + \lceil \log k \rceil}$, we know from previous lemmas that all $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\ge j + \lceil \log k \rceil}$ can be enumerated in $\mathrm{O}(n^7)$ time. □

Proof of Theorem 7.3.1: Applying similar ideas as in Lemma 7.2.7, this follows immediately from Lemmas 7.3.2, 7.3.3, 7.3.5, 7.3.6, and 7.3.7. □

Let a $(0 \le a \le k-1)$ be an integer. Shift the decomposition as before. Let IS_a be the independent set returned by the algorithm for some value of a $(0 \le a \le k-1)$ and let $IS_{\max}$ be a largest such set. Using similar ideas as in Lemma 7.2.8, we obtain the following.

Lemma 7.3.8 $|IS_{\max}| \ge (1 - \frac{4}{k})\,|OPT|$.

Proof: Define $\mathcal{D}^b_a$ again as the set of disks intersecting the boundary of a j-square S at their level, i.e. $\mathcal{D}^b_a = \bigcup_S \mathcal{D}^{\mathrm{b}(S)}_{=j(S)}$. Following Lemma 7.2.8, a disk of level j is in $\mathcal{D}^b_a$ for at most 4 different values of a. Hence there is a value of a (say a^*) for which $\left|OPT \cap \mathcal{D}^b_{a^*}\right| \le \frac{4}{k}\,|OPT|$.

From Theorem 7.3.1, we know that $|IS_{a^*}| \ge \sum_S \left|OPT^{\mathrm{i}(S)}_{=j(S)}\right|$.

Then

$$\begin{aligned} |IS_{a^*}| &\geq \textstyle\sum_S \left| OPT^{\mathrm{i}(S)}_{=j(S)} \right| \\ &= |OPT| - \left|(OPT \cap \mathcal{D}^b_{a^*})\right| \\ &\geq |OPT| - \tfrac{4}{k}\,|OPT|\,. \end{aligned}$$

Hence $|IS_{\max}| \geq |IS_{a^*}| \geq (1 - \frac{4}{k})\,|OPT|$ and the lemma follows. □

We can now prove the following.

Theorem 7.3.9 *There is an eptas for Maximum Independent Set on disk graphs of bounded level density, i.e. of level density* $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$.

Proof: Consider any $\epsilon > 0$. Choose k as the largest odd integer such that $(32k/\pi) \cdot \log(2ed) \leq \log n$. If $k < 5$, output any single vertex. Otherwise, using Theorem 7.3.1 and the choice of k, compute and output $IS_{\max}$ in $\mathrm{O}(n^{10} \log^4 n)$ time. Furthermore, if $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$, there is a c_ϵ such that $k \geq 4/\epsilon$ and $k \geq 5$ for all $n \geq c_\epsilon$. Therefore, if $n \geq c_\epsilon$, it follows from Lemma 7.3.8 and the choice of k that $IS_{\max}$ 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 level bounded density, i.e. of level density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$. Because the existence of a fiptas$^\omega$ implies the existence of an eptas (see Theorem 2.2.4), the theorem follows. □

Now observe that d is bounded by n. Hence the worst case running time of the scheme is $\mathrm{O}(k^4 n^9 (2en)^{\frac{32}{\pi}k})$.

Theorem 7.3.10 *The above algorithm is a ptas for Maximum Independent Set on disk graphs.*

The ptas given here improves on the $n^{\mathrm{O}(k^2)}$-time ptas by Erlebach, Jansen, and Seidel [103] and matches the $n^{\mathrm{O}(k)}$-time ptas by Chan [57].

7.4 Further Improvements

We gave an eptas for Minimum Vertex Cover on general disk graphs and an eptas for Maximum Independent Set on disk graphs of bounded level density. The latter scheme is also a ptas on general disk graphs. These algorithms extend to any constant dimension. Furthermore, they can be extended to intersection graphs of more general objects than disks, such as squares, triangles, etc., as long as the objects are sufficiently 'disk-like'. In other words, the objects should be *fat*. Many formal definitions of 'fat' exist, but as an example, it is easy to see the algorithms work for α-fat objects (a convex subset s of $\mathbb{R}^2$ is *α-fat* for some $\alpha \geq 1$ if the ratio between the radii of the smallest disk enclosing s and the largest disk inscribed in s is at most α [97]).

We cannot hope for a ptas on intersection graphs of nonfat objects in three dimensions, even if they have ply 1. Theorem 3.3.1 showed that any graph is an intersection graph of a set of three-dimensional convex polytopes of ply 1. Hence Maximum Independent Set and Minimum Vertex Cover are as hard on such intersection graphs as on general graphs.

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 level density is bounded. These schemes are a ptas on disk graphs of arbitrary density. Moreover, they extend to fat objects and to any constant dimension. Unfortunately, the idea of Theorem 7.2.9 that reduces the ply of the disk graph for the minimum vertex cover problem does not seem to carry over to Minimum-Weight Vertex Cover. Therefore the question of the existence of an eptas in the weighted case on disks of arbitrary size remains open.

Beyond these generalizations, an important question is whether one can improve on the algorithms given in this chapter? Here we refer to the results of Section 6.4. Recall that Maximum Independent Set on unit disk graphs of density d cannot have a ptas with running time $2^{\mathrm{O}(\mathrm{poly}(1/\epsilon))} d^{\mathrm{o}(1/\epsilon)} n^{\mathrm{O}(1)}$. Furthermore, there is a constant d_0 such that Minimum Vertex Cover on unit disk graphs of density at most d_0 has no $2^{\mathrm{O}(\mathrm{poly}(1/\epsilon))} n^{\mathrm{O}(1)}$ time eptas. Both results are under the condition that the exponential time hypothesis is true. Because the notions of ply, level density, and density are essentially the same for unit disk graphs, the following result immediately follows from Theorem 6.4.3 and Theorem 6.4.7.

Theorem 7.4.1 *If there exist constants $\delta \geq 1$, $0 < \beta < 1$ such that Maximum Independent Set on disk graphs of level 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. If there is a constant $0 < \beta < 1$ such that for any constant γ_0 Minimum Vertex Cover on disk graphs of ply at most γ_0 has a $2^{\mathrm{O}(1/\epsilon)^{1-\beta}} n^{\mathrm{O}(1)}$-time eptas, then the exponential time hypothesis is false.*

The approximation schemes for Maximum Independent Set on disk graphs of level density d and for Minimum Vertex Cover on disk graphs of ply γ described in this chapter are clearly optimal with respect to the above theorem.

Theorem 7.4.2 *Maximum Independent Set on n-vertex disk graphs of level density $d = d(n) = \Omega(n^{\alpha})$ for some constant $0 < \alpha \leq 1$ cannot have an eptas, unless FPT=W[1].*

The bound in this theorem matches the bound in Theorem 7.3.9, where we showed that Maximum Independent Set has an eptas on n-vertex disk graphs of level density $d = d(n) = \mathrm{O}(n^{\mathrm{o}(1)})$.

These results make it very unlikely that one can (significantly) improve on the schemes in this section.

7.5 The Maximum Number of Disjoint Unit Disks Intersecting a Unit Square is 7

We prove Theorem 7.3.4, which basically asks the following. Consider a unit square (i.e. a 1×1 square) and unit disks (i.e. disks of radius $\frac{1}{2}$). Determine the maximum number of nonintersecting unit disks intersecting the unit square. Here touching disks are assumed to intersect.

A trivial lower bound is 7. Placing a central disk in the center of the unit square and 6 disks around it gives a set of 7 nonintersecting disks.

A trivial upper bound is 9. All disks intersecting the unit square are completely contained in a 3×3 square around the unit square. De Groot, Peikert, and Würtz [81, 223] have shown that in the densest packing of 10 nonoverlapping (but possibly touching) disks in a 3×3 square, the disks have radius ≈ 0.444612. Hence a packing with 10 radius $\frac{1}{2}$ disks cannot exist. The upper bound of 9 follows.

We now aim to lower the upper bound. We first prove an upper bound of 8 and then further reduce it to 7, matching the lower bound.

In the following, we assume without loss of generality that the unit square is axis-aligned and that its center lies on the origin. We will not directly prove upper bounds on the number of nonintersecting unit disks intersecting the unit square, but instead focus on the bounding the number of nonintersecting unit disks intersecting the unit square, *but not intersecting the origin*. It can be readily seen that an upper bound of x on the latter number implies an upper bound of $x + 1$ on the former number.

Consider Figure 7.1. The unit square is drawn dashed. The rounded rectangle R around the unit square contains all points at distance exactly $\frac{1}{2}$ from the unit square. Now the center of any unit disk intersecting the unit square, but not intersecting the origin, must lie on or within R, but outside of the unit disk C centered on the origin. This 'allowed' area is shaded in the figure and is denoted by A.

Let c be the center of an arbitrary unit disk with $c \in A$. Consider the line segment from the origin to c and extend this segment until it intersects R (see Figure 7.2). Call this intersection point c^p. We use superscript 'p' to indicate that c^p is the projection of c onto R.

Now let $\mathsf{disk}(x, r)$ denote the disk of radius r centered on point x. If c and c' are the centers of two nonintersecting unit disks, then obviously $\mathsf{disk}(c, \frac{1}{2}) \cap \mathsf{disk}(c', \frac{1}{2}) = \emptyset$. Equivalently, it must be that $c' \notin \mathsf{disk}(c, 1)$ and $c \notin \mathsf{disk}(c', 1)$. We will combine this observation with the following lemma.

Lemma 7.5.1 *Let c and x be arbitrary points in A. If $x \notin \mathsf{disk}(c, 1)$, then $x \notin \mathsf{disk}(c^p, 1)$.*

Proof: It is sufficient to prove that the two intersection points of the boundaries of $\mathsf{disk}(c, 1)$ and $\mathsf{disk}(c^p, 1)$ are not in A. Furthermore, we only need to consider points c on C. Because if c is on C, then the two intersection points

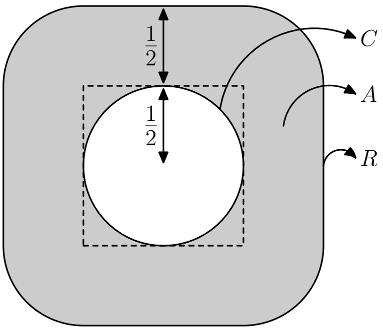

Figure 7.1: The unit disk C, rounded rectangle R at distance $\frac{1}{2}$ from the unit square, and the allowed area A.

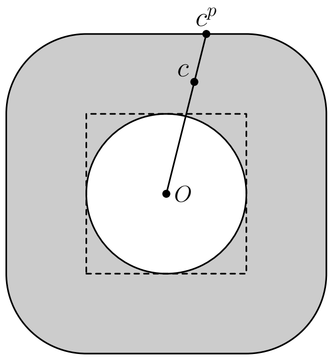

Figure 7.2: The projection point c^p of c on R.

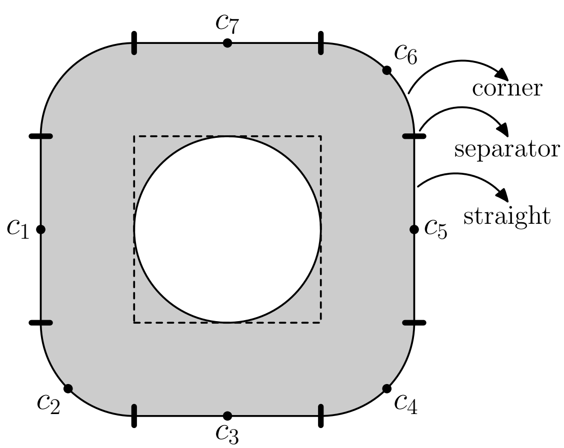

Figure 7.3: Corners, separators, and straights.

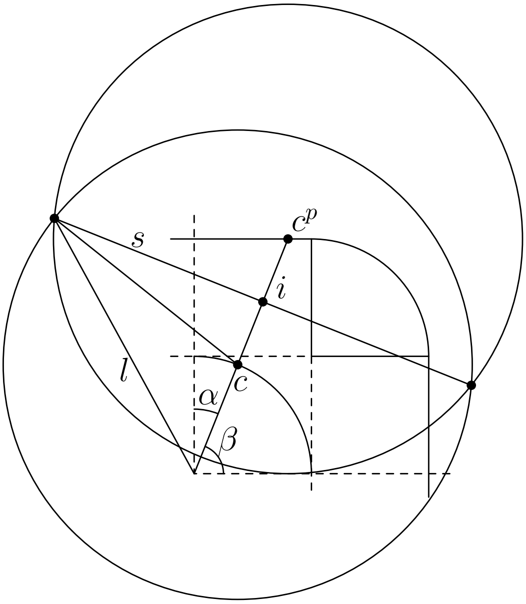

Figure 7.4: The situation if $\frac{1}{2}\pi - \tan^{-1}(\frac{1}{2}) \leq \beta \leq \frac{1}{2}\pi$.

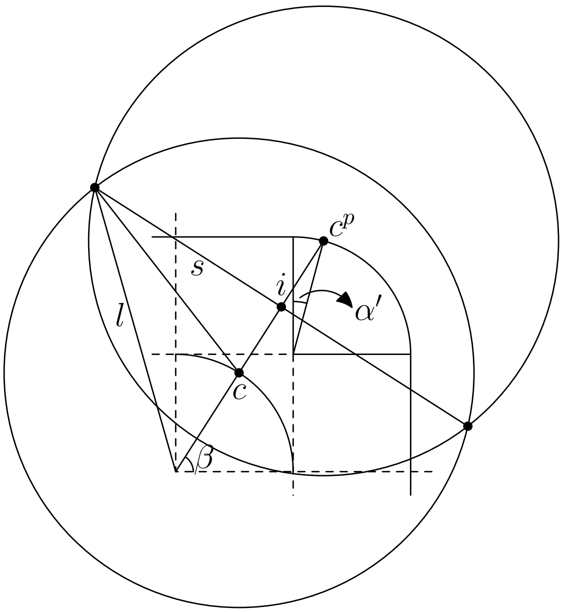

Figure 7.5: The situation if $\frac{1}{4}\pi \leq \beta \leq \frac{1}{2}\pi - \tan^{-1}(\frac{1}{2})$.

of the boundaries of $\mathsf{disk}(c,1)$ and $\mathsf{disk}(c^p,1)$ are closer to A than the two intersection points of the boundaries of $\mathsf{disk}(c',1)$ and $\mathsf{disk}(c'^p,1)$ for some c' on $\overline{cc^p}$. Due to symmetry, it is sufficient to consider angles β with $\frac{1}{4}\pi \le \beta \le \frac{1}{2}\pi$ (see Figure 7.4 and 7.5).

In Figure 7.4 and 7.5, i is the point in the middle between c and c^p. Then

$$\|i\| = \tfrac{1}{2}\left(\|c^p\| - \tfrac{1}{2}\right) + \tfrac{1}{2} = \tfrac{1}{2}\left(\|c^p\| + \tfrac{1}{2}\right)$$

and

$$\|s\| = \sqrt{1 - \|i-c\|^2} = \sqrt{1 - \left(\tfrac{1}{2}\left(\|c^p\| - \tfrac{1}{2}\right)\right)^2}.$$

This implies that

$$\begin{aligned} \|l\| &= \sqrt{\|i\|^2 + \|s\|^2} \\ &= \sqrt{\tfrac{1}{4}\|c^p\|^2 + \tfrac{1}{4}\|c^p\| + \tfrac{1}{16} + 1 - \tfrac{1}{4}\|c^p\|^2 + \tfrac{1}{4}\|c^p\| - \tfrac{1}{16}} \\ &= \sqrt{1 + \tfrac{1}{2}\|c^p\|}. \end{aligned}$$

It remains to determine $\|c^p\|$. We distinguish two cases.

If $\frac{1}{2}\pi - \tan^{-1}(\frac{1}{2}) \le \beta \le \frac{1}{2}\pi$, then α is between 0 and $\tan^{-1}(\frac{1}{2})$. But then we can easily see that $\|c^p\| = \frac{1}{\cos\alpha}$, and thus

$$\|l\| = \sqrt{1 + \frac{1}{2\cos\alpha}}$$

We next compute the derivative

$$\frac{\mathrm{d}\,\|l\|}{\mathrm{d}\alpha} = \frac{1}{\sqrt{1 + \frac{1}{2\cos\alpha}}} \cdot \frac{1}{4\cos^2\alpha} \cdot \sin\alpha$$

For $0 \le \alpha \le \tan^{-1}(\frac{1}{2})$, $\frac{\mathrm{d}\|l\|}{\mathrm{d}\alpha}$ is strictly positive. Hence for $\frac{1}{2}\pi - \tan^{-1}(\frac{1}{2}) \le \beta \le \frac{1}{2}\pi$, $\|l\|$ is at least $\sqrt{1 + \frac{1}{2\cos 0}} = \sqrt{3/2} \approx 1.225$.

If $\frac{1}{4}\pi \le \beta \le \frac{1}{2}\pi - \tan^{-1}(\frac{1}{2})$, then $0 \le \alpha' \le \frac{1}{4}\pi$. Using the Cosine Law,

$$\|c^p\| = \sqrt{\tfrac{1}{2} + \tfrac{1}{4} - 2\tfrac{1}{\sqrt{2}}\tfrac{1}{\sqrt{4}} \cdot \cos\left(\alpha' + \tfrac{3}{4}\pi\right)} = \sqrt{\tfrac{3}{4} - \tfrac{1}{2}\sqrt{2} \cdot \cos\left(\alpha' + \tfrac{3}{4}\pi\right)}.$$

Then

$$\|l\| = \sqrt{1 + \tfrac{1}{2}\sqrt{\tfrac{3}{4} - \tfrac{1}{2}\sqrt{2} \cdot \cos\left(\alpha' + \tfrac{3}{4}\pi\right)}}.$$

We again look at the derivative

$$\tfrac{\mathrm{d}\|l\|}{\mathrm{d}\alpha'} = \tfrac{1}{16}\sqrt{2} \cdot \tfrac{1}{\|l\|}\tfrac{1}{\|c^p\|} \cdot \sin\left(\alpha' + \tfrac{3}{4}\pi\right).$$

If $0 \le \alpha' \le \frac{1}{4}\pi$, then $\frac{\mathrm{d}\|l\|}{\mathrm{d}\alpha'}$ is nonnegative. Hence for $\frac{1}{4}\pi \le \beta \le \frac{1}{2}\pi - \tan^{-1}(\frac{1}{2})$, $\|l\|$ is at least $\sqrt{1 + \frac{1}{2}\sqrt{5/4}} \approx 1.249$.

Observe that the radius of the smallest circle enclosing R is $\frac{1}{2} + \frac{1}{2}\sqrt{2} \approx 1.207$. Since $\|l\| \geq \sqrt{3/2} \approx 1.225$, $\|l\| > \frac{1}{2} + \frac{1}{2}\sqrt{2}$. Hence the two intersection points of the boundaries of $\mathsf{disk}(c, 1)$ and $\mathsf{disk}(c^p, 1)$ are not in A. $\square$

Given any set of nonintersecting disks intersecting the unit square, but not intersecting the origin, with centers $c_1, \ldots, c_k$, we can thus find an equivalent set of disks with centers $c_1^p, \ldots, c_k^p$, which are also nonintersecting and intersect the unit square, but not the origin. So we may assume that all disk centers of disks not intersecting the origin are on R.

Theorem 7.5.2 *The maximum number of nonintersecting unit disks, intersecting the unit square, but not intersecting the origin, is at most* 7.

Proof: Because the centers of any such a set of unit disks lie on R, the distance between any two centers on R must be at least 1 as well (follows from the triangle inequality). Because R has length $4+\pi \approx 7.142$, there can be at most 7 such centers on R. $\square$

Corollary 7.5.3 *The number of nonintersecting unit disks intersecting the unit square is at most* 8.

In the proof of the above theorem, we used that the distance on R between any two disk centers must be at least 1. By considering this distance more closely, we can improve the bound.

Theorem 7.5.4 *The number of nonintersecting unit disks, intersecting the unit square, but not intersecting the origin, is at most* 6.

Proof: For sake of contradiction, assume $c_1, \ldots, c_7$ are the centers of 7 such unit disks. Partition R into *corners* and *straights* as shown in Figure 7.3. A point on a separator is assumed to belong to the adjacent corner. Then the 4 corners and the 4 straights partition R. Furthermore, a corner or a straight can contain at most one disk center c_i ($1 \leq i \leq 7$). Then there is either a corner or a straight which does not contain a disk center c_i.

Suppose there is a corner which does not contain a disk center. Consider a corner which *does* contain a disk center (see Figure 7.6). We determine the length on R of edges $\overline{c_i c_{i+1}}$ and $\overline{c_i c_{i-1}}$, depending on α. So let $l = \|\overline{c_i c_{i+1}}\|_R + \|\overline{c_i c_{i-1}}\|_R$. We know that $\|\overline{c_i c_{i+1}}\| \geq 1$. Using that triangle T is an isosceles triangle, $\beta = \frac{\pi-\alpha}{2}$. Then basic trigonometry gives that the part of the straight covered by $\overline{c_i c_{i+1}}$ has length at least

$$\sqrt{1 - \left(\tfrac{1}{2} - \tfrac{1}{2}\cos\alpha\right)^2} - \tfrac{1}{2}\sin\alpha.$$

But then

$$\|\overline{c_i c_{i+1}}\|_R \geq \sqrt{1 - \left(\tfrac{1}{2} - \tfrac{1}{2}\cos\alpha\right)^2} - \tfrac{1}{2}\sin\alpha + \tfrac{1}{2}\alpha.$$

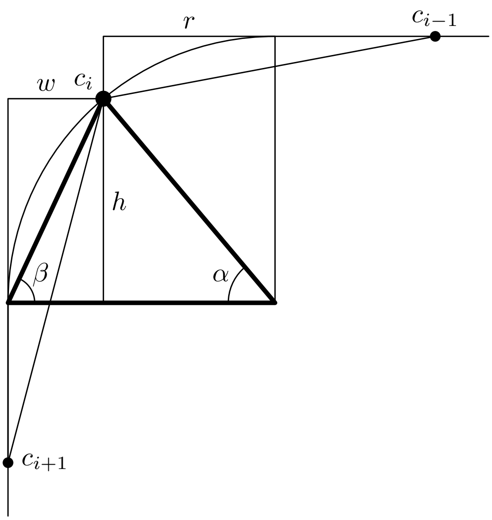

Figure 7.6: The triangle T is formed by the three thick lines. As T is an isosceles triangle, $\beta = \frac{\pi-\alpha}{2}$. From the figure, we derive that $\|h\| = \frac{1}{2}\sin\alpha$, $\|r\| = \frac{1}{2}\cos\alpha$, and $\|w\| = \frac{1}{2} - \frac{1}{2}\cos\alpha$.

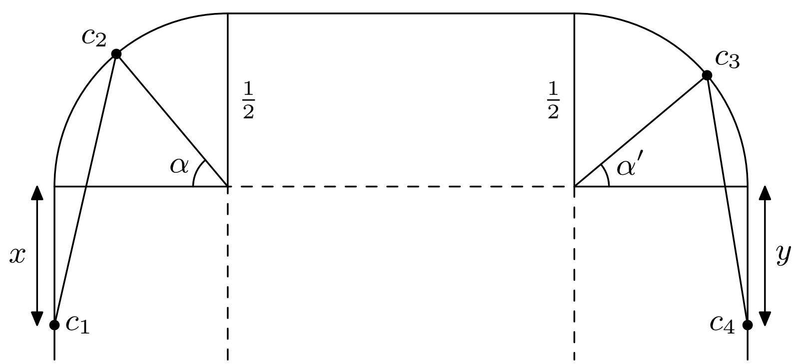

Figure 7.7: The top straight contains no disk center. Clearly, $\|x\|+\|y\|$ is minimal if $\alpha = \alpha' = \frac{1}{2}$.

Similarly,

$$\|\overline{c_i c_{i-1}}\|_R \geq \sqrt{1 - \left(\tfrac{1}{2} - \tfrac{1}{2}\sin\alpha\right)^2} - \tfrac{1}{2}\cos\alpha + \tfrac{1}{4}\pi - \tfrac{1}{2}\alpha.$$

Hence

$$l \geq \sqrt{1 - \left(\tfrac{1}{2} - \tfrac{1}{2}\cos\alpha\right)^2} + \sqrt{1 - \left(\tfrac{1}{2} - \tfrac{1}{2}\sin\alpha\right)^2} - \tfrac{1}{2}\sin\alpha - \tfrac{1}{2}\cos\alpha + \tfrac{1}{4}\pi.$$

Then the derivative is

$$\begin{aligned}\frac{\mathrm{d}l}{\mathrm{d}\alpha} \quad = \quad & -\tfrac{1}{2}\frac{1}{\sqrt{1-\left(\frac{1}{2}-\frac{1}{2}\cos\alpha\right)^2}}\left(\tfrac{1}{2} - \tfrac{1}{2}\cos\alpha\right)\sin\alpha \\ & +\tfrac{1}{2}\frac{1}{\sqrt{1-\left(\frac{1}{2}-\frac{1}{2}\sin\alpha\right)^2}}\left(\tfrac{1}{2} - \tfrac{1}{2}\sin\alpha\right)\cos\alpha \\ & -\tfrac{1}{2}\cos\alpha + \tfrac{1}{2}\sin\alpha.\end{aligned}$$

If $0 \leq \alpha < \frac{1}{4}\pi$, then $\frac{\mathrm{d}l}{\mathrm{d}\alpha} < 0$. If $\frac{1}{4}\pi < \alpha \leq \frac{1}{2}\pi$, $\frac{\mathrm{d}l}{\mathrm{d}\alpha} > 0$. If $\alpha = \frac{1}{4}\pi$, $\frac{\mathrm{d}l}{\mathrm{d}\alpha} = 0$. So $l = \|\overline{c_i c_{i+1}}\|_R + \|\overline{c_i c_{i-1}}\|_R$ is minimal if $\alpha = \frac{1}{4}\pi$ and has value $l_{\min} \approx 2.057$. Because the disk centers can be numbered arbitrarily, we may assume that c_2, c_4, and c_6 are on corners. Using symmetry, we may also assume that they are on counter-clockwise consecutive corners, as shown in Figure 7.3. Then

$$\begin{aligned}&(\|\overline{c_1 c_2}\|_R + \|\overline{c_2 c_3}\|_R) + (\|\overline{c_3 c_4}\|_R + \|\overline{c_4 c_5}\|_R) + (\|\overline{c_5 c_6}\|_R + \|\overline{c_6 c_7}\|_R) + \|\overline{c_7 c_1}\|_R \\ &\geq 3l_{\min} + 1 \approx 7.17\end{aligned}$$

This is larger than the length of R, which is a contradiction.

So there must be a straight which does not contain a disk center. Using symmetry, we may assume that this is the top straight (see Figure 7.7). Note that $0 \leq \alpha, \alpha' \leq \frac{1}{2}\pi$. We minimize $\|x\| + \|y\|$. Trivially, this minimum is attained if $\alpha = \alpha' = \frac{1}{2}\pi$. In this case,

$$\|x\| + \|y\| = 2\left(\sqrt{1 - \left(\tfrac{1}{2}\right)^2} - \tfrac{1}{2}\right) = \sqrt{3} - 1.$$

Hence

$$\begin{aligned}&(\|\overline{c_1 c_2}\|_R + \|\overline{c_2 c_3}\|_R + \|\overline{c_3 c_4}\|_R) + \|\overline{c_4 c_5}\|_R + \|\overline{c_5 c_6}\|_R + \|\overline{c_6 c_7}\|_R + \|\overline{c_7 c_1}\|_R \\ &\geq (\sqrt{3} - 1 + \tfrac{1}{2}\pi + 1) + 4 \approx 7.30\end{aligned}$$

This is larger than the length of R, which is a contradiction. Therefore there can be at most 6 nonintersecting unit disks intersecting the unit square, but not intersecting the origin. □

Using this upper bound and the lower bound given before, we have proved Theorem 7.3.4.

Chapter 8

Domination on Geometric Intersection Graphs

This chapter only treats the minimum dominating set problem on geometric intersection graphs. Although on general graphs the approximability of Minimum Dominating Set has been settled [156, 197, 66, 108], the problem is still open on numerous graph classes, including several classes of geometric intersection graphs.

In studying approximation algorithms for fundamental graph optimization problems on geometric intersection graphs, we demonstrated the power of the geometric shifting technique to approximate these problems. In particular, we were able to obtain better polynomial-time approximation schemes for Maximum Independent Set and Minimum Vertex Cover on unit disk graphs (Chapter 6) and on general disk graphs (Chapter 7). Moreover, we found a better ptas for Minimum (Connected) Dominating Set on unit disk graphs (Chapter 6), again using the shifting technique. These algorithms extend to intersection graphs of (unit) fat objects in any constant dimension and (at least partially) to the weighted case (see Section 6.3.5 and 7.4).

Interestingly, as pointed out by Erlebach, Jansen, and Seidel [103], these techniques do not seem sufficient for handling Minimum Dominating Set on intersection graphs of objects of different sizes. As far as we know, there are no results on intersection graphs of arbitrary disks, squares, etc., beyond the $(1 + \ln n)$-approximation ratio of the greedy algorithm [156, 197, 66]. In particular, we know of no constant-factor approximation algorithm or approximation hardness results. In this chapter, we address this open problem by studying Minimum Dominating Set on intersection graphs of different types of fat objects and providing new insights into its approximability.

In Section 8.2, we present a new general approach to deriving approximation algorithms for Minimum Dominating Set on geometric intersection graphs. We apply it to obtain the first constant-factor approximation algorithms for Minimum Dominating Set on intersection graphs of pairwise homothetic polygons with a constant number of corners and on intersection graphs of rectangles of bounded aspect-ratio.

We also obtain a constant-factor approximation algorithm for Minimum Dominating Set on disk graphs of constant ply (see Section 8.4). A surprising corollary of this is a constant integrality gap of the standard linear program (LP) for Minimum Dominating Set on planar graphs. For disk graphs of bounded ply, we can improve this result to a $(3+\epsilon)$-approximation algorithm by using a new variant of the shifting technique. This algorithm extends to intersection graphs of fat objects of bounded ply and constant dimension.

The type of fat objects one considers has a strong impact on the approximability of Minimum Dominating Set, as shown in Section 8.5. We prove that on intersection graphs of n convex fat objects, approximation ratio $(1-\epsilon)\ln n$ is not achievable in polynomial time for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{O(\log \log n)}$). This also holds on intersection graphs of pairwise homothetic objects. Finally, we solve an open problem of Chlebík and Chlebíková [65], who asked whether their APX-hardness results for Minimum Dominating Set on intersection graphs of d-dimensional axis-parallel boxes if $d \geq 3$ extend to the case where $d = 2$. We affirm this by showing that Minimum Dominating Set is APX-hard on rectangle intersection graphs.

8.1 Small ϵ-Nets

The core of the algorithmic results of Section 8.2 relies on the availability of small ϵ-nets. Given a universe $\mathbb{U}$, a family $\mathcal{S}$ of subsets of $\mathbb{U}$ (called *objects*), and a (positive) weight function w over $\mathcal{S}$, we say that $\mathcal{R} \subseteq \mathcal{S}$ is an *ϵ-net* for $\mathcal{S}$ if any element $u \in \mathbb{U}$ for which $\sum_{s \in \mathcal{S}: u \in s} w(s) > \epsilon W$ is covered by $\mathcal{R}$ (i.e. $u \in \bigcup \mathcal{R}$), where $W = \sum_{s \in \mathcal{S}} w(s)$. In the classic definition of an ϵ-net, it assumed that all weights are equal to 1. That is, $\mathcal{R} \subseteq \mathcal{S}$ is a *binary ϵ-net* for $\mathcal{S}$ if any element $u \in \mathbb{U}$ covered by more than $\epsilon|\mathcal{S}|$ sets in $\mathcal{S}$ is also covered by $\mathcal{R}$. The *size* of a (binary) ϵ-net is the cardinality of $\mathcal{R}$.

We should note that in a way there are two definitions of an ϵ-net, that are essentially dual to each other [143, 68]. In the covering version of ϵ-nets, described above, we aim to select objects to cover elements that are covered by a lot of objects. In the dual definition, the hitting version, we need to select elements to hit all objects containing a large number of elements. Here we only need the covering variant and thus disregard the hitting version.

There have been several results on ϵ-nets in the past (e.g. [143, 35, 170, 205, 62, 44, 68, 188, 226]). The most general result is the following. Given a (finite) universe $\mathbb{U}$ and a family $\mathcal{S}$ of subsets of $\mathbb{U}$, let $S(u) = \{s \in S \mid u \in s\}$ for any $S \subseteq \mathcal{S}$. Then the *dual Vapnik-Chervonenkis dimension* or *dual VC-dimension* of $(\mathbb{U}, \mathcal{S})$ is equal to the cardinality of a largest set $S \subseteq \mathcal{S}$ for which $\{S(u) \mid u \in \mathbb{U}\}$ equals the power set of S [143].

Theorem 8.1.1 ([170]) *Suppose that* $(\mathbb{U}, \mathcal{S})$ *has dual VC-dimension* d. *Then for any* $\epsilon > 0$ *that is sufficiently small with respect to* d *there is a binary* ϵ*-net for* $\mathcal{S}$ *of size at most* $(d/\epsilon) \cdot (\log(1/\epsilon) + 2 \log\log(1/\epsilon) + 3)$.

There are many examples of set systems with constant dual VC-dimension. For instance, recall from Chapter 3 the representation of an arbitrary graph as an intersection graph. Given a graph G, let $\mathbb{U} = E(G)$ and $\mathcal{S} = \{S_v \mid v \in V(G)\}$, where $S_v = \{(u,v) \in E(G) \mid u \in V(G)\}$ for any $v \in V(G)$. This set system can easily be shown to have dual VC-dimension at most 2. Hence, by Theorem 8.1.1, it has an ϵ-net of size $\mathrm{O}(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$. One can however improve on this bound.

Theorem 8.1.2 *Let $(\mathbb{U}, \mathcal{S})$ be induced by a graph G (as described above) and let w be a positive weight function over $\mathcal{S}$. Then one can find an ϵ-net of $\mathcal{S}$ of size at most $2/\epsilon$ in linear time.*

Proof: We need to cover all elements of $\mathbb{U}$ covered by sets of $\mathcal{S}$ with total weight exceeding ϵW. Any $u \in \mathbb{U}$ is in at most 2 sets of $\mathcal{S}$, say s_u^1 and s_u^2. If $w(s_u^1) + w(s_u^2) > \epsilon W$, then $\max\{w(s_u^1), w(s_u^2)\} > \epsilon W/2$. Hence $\mathcal{R} = \{s \in \mathcal{S} \mid w(s) > \epsilon W/2\}$ is an ϵ-net. Moreover, $|\mathcal{R}| < 2/\epsilon$. □

The bound of Theorem 8.1.2 is essentially tight. For $m > 0$, let $G = K_{2m}$ and let $(\mathbb{U}, \mathcal{S})$ be the set system induced by G. Set $w(s) = 1$ for each $s \in \mathcal{S}$ and set $\epsilon = 1/(m + \frac{1}{4})$. An ϵ-net for $(\mathbb{U}, \mathcal{S})$ is equal to a vertex cover of all edges $(u,v) \in E(G)$ for which $w(u) + w(v) > \epsilon W$. Clearly, $w(u) + w(v) = 2 > \epsilon \cdot 2m$ for each $(u,v) \in E(G)$. But each vertex cover of G needs at least $2m - 1$ vertices, while $2/\epsilon < 2m + 1$. As m tends to infinity, this is tight.

For geometric intersection graphs one can prove similar bounds. A family $\mathcal{S}$ of subsets of $\mathbb{U} = \mathbb{R}^2$ is a family of *pseudo-disks* if the sets in $\mathcal{S}$ are bounded by simple closed Jordan curves, such that each pair of curves intersects at most twice. Examples are families of disks, squares, or homothetic polygons. Given such $\mathbb{U}$ and $\mathcal{S}$, the next theorem follows from results of Chazelle and Friedman [62], Clarkson and Varadarajan [68], and Kedem et al. [161].

Theorem 8.1.3 *For any $\epsilon > 0$, there is a binary ϵ-net for $\mathcal{S}$ of size* $\mathrm{O}(1/\epsilon)$.

Such a net can be found by a randomized algorithm with polynomial expected running time [62, 68]. By derandomizing the algorithm using the method of conditional expectations, we can prove that a binary ϵ-net as in Theorem 8.1.3 can be found in time polynomial in $|\mathcal{S}|$ and $1/\epsilon$ [62, 256].

The above results are actually corollaries of more general theorems that relate the size of the ϵ-net to the union complexity of the set $\mathcal{S}$. An extensive treatment may be found in [62, 68, 256].

Pyrga and Ray [226] recently improved on Theorem 8.1.3 and the associated algorithms. The ϵ-nets following from their results also have size $\mathrm{O}(1/\epsilon)$, but with a much better hidden constant. Moreover, both the analysis and the algorithm needed to compute the net are easier.

Theorem 8.1.4 *For any $\epsilon > 0$, one can obtain a binary ϵ-net for $\mathcal{S}$ of size $\mathrm{O}(1/\epsilon)$ in time polynomial in $|\mathcal{S}|$ and $1/\epsilon$.*

Linear-sized ϵ-nets also exist for three-dimensional objects. Clarkson and Varadarajan [68] showed that an ϵ-net exists for unit cubes. This result was subsequently generalized by Laue [188].

Theorem 8.1.5 ([188]) *For any $\epsilon > 0$, one can obtain a binary ϵ-net of size $O(1/\epsilon)$ for a set $\mathcal{S}$ of translates of a fixed three-dimensional polytope in time polynomial in $|\mathcal{S}|$ and $1/\epsilon$.*

Note that the above algorithms find binary ϵ-nets. One can transform them into algorithms to find a (weighted) ϵ-net at relatively small cost.

Definition 8.1.6 *Algorithm A is a* net finder with size-function g *for $(\mathbb{U}, \mathcal{S})$ if for any $\epsilon > 0$ and any (positive) weight function w over $\mathcal{S}$, A gives an ϵ-net for $(\mathbb{U}, \mathcal{S})$ of size at most $g(1/\epsilon)$ in time polynomial in $|\mathcal{S}|$, $1/\epsilon$, and the size of a representation of w.*

The definition of a *binary net finder with size-function g* is similar. We will always assume the size-function g to be nondecreasing.

Proposition 8.1.7 ([45]) *If A is a binary net finder with size-function g for some $(\mathbb{U}, \mathcal{S})$, then there is a net finder A' with size-function $g'(1/\epsilon) = g(2/\epsilon)$ for $(\mathbb{U}, \mathcal{S})$.*

Proof: Let some $\epsilon > 0$ and some (positive) weight function w over $\mathcal{S}$ be given. Scale the weights to w' such that $W' = \sum_{s \in \mathcal{S}} w'(s) = |\mathcal{S}|$. Take $\lceil w'(s) \rceil$ copies of each $s \in \mathcal{S}$ and denote the resulting set of objects by $\mathcal{S}'$. Then

$$|\mathcal{S}'| = \sum_{s \in \mathcal{S}} \lceil w'(s) \rceil < \sum_{s \in \mathcal{S}} (1 + w'(s)) = W' + |\mathcal{S}| = 2|\mathcal{S}| = 2W'.$$

Choose $\epsilon' = \epsilon/2$ and apply A to $\mathcal{S}'$ and ϵ'. This gives an ϵ'-net for $\mathcal{S}'$ of size $g(2/\epsilon)$. Since $\epsilon'|\mathcal{S}'| < \epsilon W'$, it induces an ϵ-net of $\mathcal{S}$ with respect to w', and hence with respect to w as well. Observe that the above algorithm takes time polynomial in $|\mathcal{S}|$, $1/\epsilon$, and the size of the representation of w. □

8.2 Generic Domination

We give a generic approach to approximating Minimum Dominating Set, particularly on geometric intersection graphs. To this end, we introduce the novel notion of $\preceq$-dominating sets, which we then use in combination with ϵ-nets to approximate Minimum Dominating Set.

Let $\preceq$ be a binary reflexive relation on the vertices of a graph G. For example, if G is some geometric intersection graph with representation $\mathcal{S}$, $u \preceq v$ if the size of $\mathcal{S}(u)$ is at most the size of $\mathcal{S}(v)$. We say that $v \in V(G)$ is *$\preceq$-larger* than $u \in V(G)$ if $u \preceq v$. Denote by $N_{\preceq}(u) = \{v \in V(G) \mid (u, v) \in E(G), u \preceq v\}$ the set of $\preceq$-larger neighbors of some $u \in V(G)$ and let $N_{\preceq}[u] = N_{\preceq}(u) \cup \{u\}$ denote u's closed $\preceq$-larger neighborhood. Similarly, we define $N_{\succeq}(u) = \{v \in V(G) \mid (u, v) \in E(G), v \preceq u\}$ and $N_{\succeq}[u] = N_{\succeq}(u) \cup \{u\}$.

Definition 8.2.1 *Given a graph G and a binary reflexive relation $\preceq$ on the vertices of G, $C \subseteq V(G)$ is a $\preceq$-*dominating set *of G if for any $u \in V(G)$, $u \in C$ or there is a $\preceq$-larger neighbor of u in C.*

Alternatively, $C \subseteq V(G)$ is a *$\preceq$-dominating set* of G if $C \cap N_{\preceq}[u] \neq \emptyset$ for all $u \in V(G)$. Observe that $\preceq$-dominating sets are a proper generalization of ordinary dominating sets. Simply take $\preceq$ to be the complete relation, i.e. $u \preceq v$ for all $u, v \in V(G)$. Moreover, the definition of $\preceq$-dominating set is sound, as $V(G)$ is a $\preceq$-dominating set of G, regardless of the definition of $\preceq$.

For a given relation $\preceq$, one can try to find a relation between the cardinality of a smallest dominating and of a smallest $\preceq$-dominating set.

Definition 8.2.2 *Given a graph G and a binary reflexive relation $\preceq$ on $V(G)$, the* $\preceq$-factor *is the cardinality of a minimum $\preceq$-dominating set divided by the cardinality of a minimum dominating set.*

Clearly, the $\preceq$-factor is at least 1 for any relation $\preceq$. Knowing an upper bound on the $\preceq$-factor is more interesting however, as this leads to one of the main theorems of this chapter.

Theorem 8.2.3 *Let $(\mathbb{U}, \mathcal{S})$ be a set system for which a net finder with size-function g exists and let $\preceq$ be a binary reflexive relation on the vertices of $G = G[\mathcal{S}]$ with $\preceq$-factor at most c_1 such that for any $u \in V(G)$ there exist at most c_2 elements of $\mathbb{U}$ in $\mathcal{S}(u)$ jointly hitting all $\mathcal{S}(v)$ with $v \in N_{\preceq}(u)$. If the cardinality of a minimum dominating set of G is k, then one can find a dominating set of G of cardinality at most $g(c_1c_2k)$ in time polynomial in $|\mathcal{S}|$.*

Proof: Consider the standard integer LP of the minimum $\preceq$-dominating set problem:

$$\begin{aligned} z_I^* \quad = \quad \min \quad & \sum_{u \in V(G)} x_u \\ \text{s.t.} \quad & \sum_{v \in N_{\preceq}[u]} x_v \geq 1 \quad && \forall u \in V(G) \\ & x_u \in \{0, 1\} && \forall u \in V(G) \end{aligned}$$

Observe that $z_I^* \leq c_1 k$. Relax the above integer LP by replacing its last constraint by $x_u \geq 0$ $\forall u \in V(G)$. Let x^* be a vector attaining the optimum fractional value z^*. Because for any $u \in V(G)$, all $\mathcal{S}(v)$ with $v \in N_{\preceq}(u)$ can be jointly hit by c_2 elements in $\mathcal{S}(u)$, each $\mathcal{S}(u)$ contains an element p such that $\sum_{v : p \in \mathcal{S}(v)} x_v^* \geq 1/c_2$. Call such an element *heavily covered.*

Now define a weight function w by $w(\mathcal{S}(u)) := x_u^* |\mathcal{S}|/z^*$. Let $W = \sum_{u \in V(G)} w(\mathcal{S}(u))$ and $\epsilon = 1/(c_2 z^*)$. Following the previous observation, this implies that any object $s \in \mathcal{S}$ contains an element p such that

$$\sum_{v : p \in \mathcal{S}(v)} w(\mathcal{S}(v)) = \sum_{v : p \in \mathcal{S}(v)} x_v^* |\mathcal{S}|/z^* = (|\mathcal{S}|/z^*) \cdot \sum_{v : p \in \mathcal{S}(v)} x_v^* \geq |\mathcal{S}|/(c_2 z^*) = \epsilon W.$$

Hence an ϵ-net $\mathcal{R} \subseteq \mathcal{S}$ for this choice of ϵ will cover all heavily covered elements. But then $\mathcal{R}$ induces a dominating set of G. Moreover,

$$|\mathcal{R}| \leq g(c_2 z^*) \leq g(c_2 z_I^*) \leq g(c_1 c_2 k).$$

Finally note that $\mathcal{R}$ can be found in time polynomial in $|\mathcal{S}|$. The optimum solution to the linear program can be found in polynomial time [163, 159]. Hence the weights of the weight function can be represented using a polynomial number of bits and therefore the ϵ-net can be found in polynomial time. $\square$

Observe that if instead of a (weighted) net finder we only have a binary net finder with size-function g, then the above algorithm yields a dominating set of cardinality at most $g(2c_1c_2k)$ by Proposition 8.1.7.

The running time of the algorithm described in Theorem 8.2.3 is determined by the time it takes to find the ϵ-net and to solve the linear program. The latter takes $O(n^{3.5}\log^2 n)$ time [159], where we ignore some sublogarithmic terms. Young [275] showed that a $(1+\delta)$-approximate solution to the linear program can be found much quicker, in $O(n^2 \log n/\delta^2)$ time. If we use such a solution in Theorem 8.2.3, the dominating set has cardinality $g((1+\delta)c_1c_2k)$.

The proof of Theorem 8.2.3 solves a linear program and finds an ϵ-net once, following a technique of Even, Rawitz, and Shahar [107]. Alternatively, one could use the iterative reweighting technique proposed by Brönnimann and Goodrich [45], where an ϵ-net is constructed in every iteration. In this chapter, finding the ϵ-net is usually quite expensive and hence we prefer the technique of Even, Rawitz, and Shahar. Moreover, it makes for an easier proof.

Another consequence of the proof of Theorem 8.2.3 is a bound on the integrality gap of the standard LP of Minimum Dominating Set. The *integrality gap* of an LP is the ratio of its optimum integral value and its optimum fractional value. For this bound, we need a fractional equivalent of the $\preceq$-factor.

Definition 8.2.4 *Given a graph G and a binary reflexive relation $\preceq$ on $V(G)$, the* fractional $\preceq$-factor *is the ratio of the optimum fractional value of the standard LP for Minimum $\preceq$-Dominating Set and the optimum fractional value of the standard LP for Minimum Dominating Set.*

For all relations $\preceq$ described in this chapter, we can find the same bound on the $\preceq$-factor as on the fractional $\preceq$-factor. It is not clear whether this is a coincidence.

We can now prove a fractional equivalent of Theorem 8.2.3.

Theorem 8.2.5 *Let $(\mathbb{U}, \mathcal{S})$ be a set system for which a net finder with size-function g exists and let $\preceq$ be a binary reflexive relation on the vertices of $G = G[\mathcal{S}]$ with fractional $\preceq$-factor at most c_3 such that for any $u \in V(G)$ there exist at most c_2 elements of $\mathbb{U}$ in $\mathcal{S}(u)$ jointly hitting all $\mathcal{S}(v)$ with $v \in N_{\preceq}(u)$. If the optimum fractional value of the standard LP for Minimum Dominating Set is z^*, then one can find a dominating set of G of cardinality at most $g(c_2c_3z^*)$ in time polynomial in $|\mathcal{S}|$.*

Proof: Let $z^*_{\preceq}$ denote the optimum fractional value of the standard LP for Minimum $\preceq$-Dominating Set on G. Following the proof of Theorem 8.2.3, one can find a dominating set of cardinality at most $g(c_2z^*_{\preceq})$ in polynomial time. As $z^*_{\preceq} \le c_3z^*$, this dominating set has cardinality at most $g(c_2c_3z^*)$. $\square$

In other words, the integrality gap is at most $g(c_2c_3z^*)/z^*$.

Again if only a binary net finder with size-function g exists, then one can find a dominating set of G of cardinality at most $g(2c_2c_3z^*)$ in polynomial time. This implies that the integrality gap is at most $g(2c_2c_3z^*)/z^*$.

As an example and to demonstrate the generality of Theorem 8.2.3 and Theorem 8.2.5, we apply them to general graphs. For a graph G, let $\Delta(G)$ denote the maximum degree of a vertex of $V(G)$.

Theorem 8.2.6 *Minimum Dominating Set has a polynomial-time* $2\Delta(G)$*-approximation algorithm on any graph* G*. Moreover, the integrality gap is at most* $2\Delta(G)$*.*

Proof: Let G be any graph and $(\mathbb{U}, \mathcal{S})$ a representation of G, i.e. $\mathbb{U} = E(G)$ and $\mathcal{S} = \{S_v \mid v \in V(G)\}$, where $S_v = \{(u, v) \in E(G) \mid u \in V(G)\}$ for any $v \in V(G)$. Define a binary relation $\preceq$ such that $u \preceq v$ for all $u, v \in V(G)$. Observe that the (fractional) $\preceq$-factor is 1. For any $u \in V(G)$, $N_{\preceq}(u) = N(u)$, and thus there exist (at most) $\Delta(G)$ elements of $\mathbb{U}$ in $\mathcal{S}(u)$ that jointly hit all $\mathcal{S}(v)$ with $v \in N_{\preceq}(u)$. Simply take all edges incident to u. Theorem 8.1.2 showed that any graph G with representation $(\mathbb{U}, \mathcal{S})$ has an ϵ-net of size $2/\epsilon$, which can be found in polynomial time. The theorem statement follows from Theorem 8.2.3 and Theorem 8.2.5. □

Note that the above algorithm can only guarantee an approximation ratio of $2\Delta(G)$, whereas a greedy algorithm giving ratio $1 + \ln \Delta(G)$ exists [156, 197, 66, 149]. Theorem 8.2.6 merely serves as an indication of the power of Theorem 8.2.3 and Theorem 8.2.5. The real challenges and offered improvements lie with geometric intersection graphs.

8.3 Dominating Set on Geometric Intersection Graphs

The main result of this section is a constant-factor approximation algorithm for Minimum Dominating Set on intersection graphs of homothetic convex polygons. The constant depends on the number of corners (i.e. the complexity) of the base polygon. We also show that on intersection graphs of regular polygons the dependence on the complexity of the base polygon can be reduced. Although homotheticity is crucial in the analysis of these results, we show that on intersection graphs of axis-parallel rectangles that are not necessarily homothetic, but have constant aspect-ratio, one can obtain a constant-factor approximation algorithm as well. A discussion of disk graphs is deferred to Section 8.4.

8.3.1 Homothetic Convex Polygons

We show that Minimum Dominating Set on intersection graphs of homothetic convex polygons with r corners has a polynomial-time $O(r^4)$-approximation algorithm. We require two auxiliary results before we are ready to prove this.

First we need a way to bound the (fractional) $\preceq$-factor of a relation $\preceq$. The next two lemmas hold for arbitrary graphs.

Lemma 8.3.1 *Let $\preceq$ be a binary reflexive relation on the vertices of a graph G such that for any $u \in V(G)$ a minimum $\preceq$-dominating set for $U_u = \{v \mid v \not\preceq u, v \in N(u)\}$ has cardinality at most c. Then the $\preceq$-factor is at most $c+1$.*

Proof: Consider some dominating set C of G and for each $u \in C$, let C_u be a minimum $\preceq$-dominating set of U_u. We claim that $C' = C \cup \bigcup_{u \in C} C_u$ is a $\preceq$-dominating set of G. For suppose that there is some $v \in V(G) - C'$ that is not $\preceq$-dominated by a vertex in C'. Because C is a dominating set of G and $C \subseteq C'$, $v \in U_u$ for some $u \in C$. But then v is $\preceq$-dominated by C_u and hence by C', a contradiction. Finally, note that $|C'| \leq (c+1) \cdot |C|$. Therefore the $\preceq$-factor is at most $c+1$. □

Observe that one only needs an upper bound on $|C_u|$ for vertices u appearing in the dominating set C.

Lemma 8.3.2 *Let $\preceq$ be a binary reflexive relation on the vertices of some graph G such that for any $u \in V(G)$ a minimum $\preceq$-dominating set for $U_u = \{v \mid v \not\preceq u, v \in N(u)\}$ has cardinality at most c. Then the fractional $\preceq$-factor is at most $c+1$.*

Proof: Let x^* be an optimal fractional solution to the standard LP for Minimum Dominating Set, with value z^*. For any $u \in V(G)$, let C_u be a minimum $\preceq$-dominating for U_u. Set x'_v to x^*_v for each $v \in V(G)$ and then add x^*_u to x'_v for each $u \in V(G)$ with $v \in C_u$. Then for any $u \in V(G)$,

$$\sum_{v \in N_{\preceq}[u]} x'_v \geq \sum_{v \in N_{\preceq}[u]} x^*_v + \sum_{v \in N[u]-N_{\preceq}[u]} x^*_v = \sum_{v \in N[u]} x^*_v \geq 1.$$

Hence x' is a solution to the standard LP for Minimum $\preceq$-Dominating Set. It has value

$$\sum_{u \in V(G)} x'_u \leq \sum_{u \in V(G)} (c+1)x^*_u = (c+1) \cdot z^*.$$

Therefore the fractional $\preceq$-factor is at most $c+1$. □

We are now ready to present the relation used in the approximation algorithm. Call the straight line segment between two corners of a convex polygon a *chord*. Observe that some chords correspond to sides of the polygon and that each chord is contained in the polygon. Let $G = G[\mathcal{S}]$ be the intersection graph of a set $\mathcal{S}$ of homothetic convex polygons. Define a relation $\preceq_{1/3}$ as follows. For any two vertices $u, v \in V(G)$, let $v \preceq_{1/3} u$ if $\mathcal{S}(u)$ contains a corner of $\mathcal{S}(v)$ or $\mathcal{S}(u)$ covers at least one third of a chord of $\mathcal{S}(v)$.

The next lemma is crucial to the analysis of the approximation algorithm. For its proof, recall the following definitions of points and lines of a triangle.

An *altitude* of a corner is the straight line through this corner, perpendicular to the side opposite the corner. A *median* of a corner is the straight line through this corner and the midpoint of the opposite side. The intersection point of the medians of a triangle is its *centroid* or *barycenter*.

Lemma 8.3.3 *Let $G = G[\mathcal{S}]$ be the intersection graph of a collection $\mathcal{S}$ of homothetic convex polygons with r corners for some $r \geq 3$. Then the $\preceq_{1/3}$-factor is at most $2r(r-2)+1$.*

Proof: Consider a dominating set C of G such that for each $u \in C$ there is no $v \in V(G)$ for which $\mathcal{S}(v)$ strictly contains $\mathcal{S}(u)$. Clearly, G always has a dominating set with this property that is also a minimum dominating set. Let $u \in C$ and consider the set $U = \{v \mid v \not\preceq_{1/3} u, v \in N(u)\}$. Following Lemma 8.3.1, it suffices to bound the cardinality of a minimum $\preceq_{1/3}$-dominating set of U by $2r(r-2)$ to prove the lemma.

Observe that for any $v \in U$, $\mathcal{S}(u)$ does not contain a corner of $\mathcal{S}(v)$. As the polygons are convex and homothetic, each $\mathcal{S}(v)$ with $v \in U$ must contain a corner of $\mathcal{S}(u)$. Consider a corner p of $\mathcal{S}(u)$ and let $U_p = \{v \mid v \in U, p \in \mathcal{S}(v)\}$ be the set of vertices $v \in U$ for which $p \in \mathcal{S}(v)$. Because $\mathcal{S}(v)$ has no corner in $\mathcal{S}(u)$ for each $v \in U_p$, there must be precisely one side of $\mathcal{S}(v)$ that intersects $\mathcal{S}(u)$. This side is not incident with the corner of $\mathcal{S}(v)$ corresponding to p. Let $U_{p,s}$ be the set of vertices $v \in U_p$ for which side s of $\mathcal{S}(v)$ intersects $\mathcal{S}(u)$.

For any p, s, reduce $\mathcal{S}(u)$ and each $\mathcal{S}(v)$ with $v \in U_{p,s}$ to the triangle induced by the corner corresponding to p and the side corresponding to s. This gives a collection $\mathcal{S}'$ of homothetic triangles all containing p, but no triangle $\mathcal{S}'(v)$ with $v \in U_{p,s}$ contains $\mathcal{S}'(u)$ or has a corner in $\mathcal{S}'(u)$. Moreover, the sides of the triangles correspond to chords of the original polygons.

Assume without loss of generality that one side of the triangles of $\mathcal{S}'$ is parallel to the x-axis and that p corresponds to the left corner of $\mathcal{S}'(u)$. Now let $v_t \in U_{p,s}$ be a vertex such that the top corner of $\mathcal{S}'(v_t)$ has the largest distance to the altitude of the left corner of $\mathcal{S}'(u)$ among all top corners of triangles in $U_{p,s}$. Similarly, let v_r be a vertex such that the right corner $\mathcal{S}'(v_r)$ has the largest distance to the altitude of the left corner of $\mathcal{S}'(u)$. We claim that v_t and v_r form a $\preceq_{1/3}$-dominating set of $U_{p,s}$.

Let w be a vertex in $U_{p,s}$ (see Figure 8.1). We may assume that $\mathcal{S}'(w)$ has no corner in $\mathcal{S}'(v_t)$, $\mathcal{S}'(v_r)$, or $\mathcal{S}'(u)$. Then $\mathcal{S}'(w)$ contains a corner of $\mathcal{S}'(v_t)$, $\mathcal{S}'(v_r)$, and $\mathcal{S}'(u)$. Furthermore, by the choice of v_t and v_r, $\mathcal{S}'(w)$ cannot strictly contain either $\mathcal{S}'(v_t)$ or $\mathcal{S}'(v_r)$, as the top or right corner of $\mathcal{S}'(w)$ would be further from the altitude than the top or right corner of $\mathcal{S}'(v_t)$ or $\mathcal{S}'(v_r)$ respectively.

Observe that there must be a side of $\mathcal{S}'(w)$ such that p is at least as far from this side as the centroid of $\mathcal{S}'(w)$. Suppose w.l.o.g. that $\mathcal{S}'(v_r)$ protrudes this side of $\mathcal{S}'(w)$. Then the corner of $\mathcal{S}'(v_r)$ in $\mathcal{S}'(w)$ is at least as far from this side as p, and thus at least as far from the side as the centroid of $\mathcal{S}'(w)$. An easy calculation shows that $\mathcal{S}'(v_r)$ covers at least one third of the side of $\mathcal{S}'(w)$. But

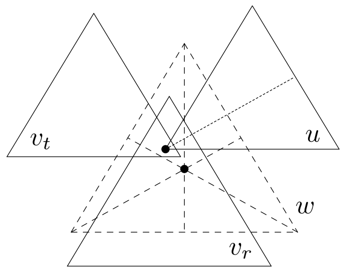

Figure 8.1: Triangles $\mathcal{S}'(u)$, $\mathcal{S}'(v_t)$, $\mathcal{S}'(v_r)$, and $\mathcal{S}'(w)$ of the proof of Lemma 8.3.3. The two dots represent p and the centroid of w. The dotted line inside $\mathcal{S}'(u)$ is the altitude of p.

then $\mathcal{S}(v_r)$ covers at least one third of a chord of $\mathcal{S}(w)$ and hence $w \preceq_{1/3} v_r$. Therefore v_t and v_r are a $\preceq_{1/3}$-dominating set of $U_{p,s}$.

As each of the r corners of the base polygon has $r-2$ sides not incident with it, U has a $\preceq_{1/3}$-dominating set of cardinality at most $2r(r-2)$. Following Lemma 8.3.1, the $\preceq_{1/3}$-factor is at most $2r(r-2)+1$. □

Combining Lemma 8.3.3 with Theorem 8.2.3, we obtain the following result.

Theorem 8.3.4 *Let $r \geq 3$ be an integer. There is a polynomial-time* $\mathrm{O}(r^4)$*-approximation algorithm for Minimum Dominating Set on intersection graphs of homothetic convex r-polygons.*

Proof: Let $G = G[\mathcal{S}]$ be the intersection graph of a collection $\mathcal{S}$ of homothetic convex r-polygons for some $r \geq 3$. Use the relation $\preceq_{1/3}$. Lemma 8.3.3 showed that the $\preceq_{1/3}$-factor is at most $2r(r-2)+1$. To hit all $\preceq_{1/3}$-larger neighbors of a vertex, place a point on each corner of the corresponding polygon and two on all chords, such that each chord is divided into three equal parts. This gives a total of $r + 2\binom{r}{2} = r^2$ points. Observe that homothetic convex polygons form a set of pseudo-disks. The theorem statement now follows from Theorem 8.1.4 and Theorem 8.2.3. □

This also implies an $\mathrm{O}(r^4)$-approximation algorithm for Minimum Connected Dominating Set on intersection graphs of homothetic convex r-polygons for $r \geq 3$ by using Proposition 6.3.24.

Another consequence of Theorem 8.3.4 is a constant-factor approximation algorithm for Minimum Dominating Set on max-tolerance (interval) graphs, because Kaufmann et al. [160] proved that max-tolerance graphs are intersection graphs of isosceles right triangles.

Using a similar proof as for Lemma 8.3.3, we can show that the fractional $\preceq_{1/3}$-factor is at most $2r(r-2)+1$. Then the following may be easily proved from Theorem 8.2.5.

Theorem 8.3.5 *Let $r \geq 3$ be an integer. The integrality gap of the standard LP for Minimum Dominating Set on intersection graphs of homothetic convex r-polygons is* $\mathrm{O}(r^4)$.

8.3.2 Regular Polygons

If the given polygons are homothetic regular polygons, then we can improve on the analysis of the previous section. We distinguish between regular polygons with an odd and with an even number of corners. Let $G = G[\mathcal{S}]$ be the intersection graph of a set $\mathcal{S}$ of homothetic odd regular polygons. Define a relation $\preceq_{1/2}$ such that for any $u, v \in V(G)$, $u \preceq_{1/2} v$ if $\mathcal{S}(v)$ contains a corner of $\mathcal{S}(u)$ or $\mathcal{S}(v)$ covers at least half of a side of $\mathcal{S}(u)$.

Lemma 8.3.6 *Let $G = G[\mathcal{S}]$ be the intersection graph of a set $\mathcal{S}$ of homothetic odd regular polygons with r corners for some odd integer $r \geq 5$. Then the $\preceq_{1/2}$-factor is at most $2r + 1$.*

Proof: Let C be a dominating set such that for each $u \in C$ there is no $v \in V(G)$ for which $\mathcal{S}(u) \subset \mathcal{S}(v)$. Consider the set $U = \{v \mid v \not\preceq_{1/2} u, v \in N(u)\}$ for some $u \in V(G)$. For each corner p of $\mathcal{S}(u)$, let $U_p = \{v \mid v \in U, p \in \mathcal{S}(v)\}$ be the set of vertices in U for which the corresponding polygon contains p. Because $\mathcal{S}(u)$ does not contain a corner of $\mathcal{S}(v)$ for any $v \in U_p$ and the polygons are odd and regular, $\mathcal{S}(u)$ protrudes the same side of each $\mathcal{S}(v)$ with $v \in U_p$.

Similar to Lemma 8.3.3, let v_t and v_b be two vertices for which this side of the corresponding polygons extends furthest in either direction. Then any $\mathcal{S}(w)$ with $w \in U$ is at most twice as large as $\mathcal{S}(v_t)$ or $\mathcal{S}(v_b)$, or this would contradict the choice of v_t or v_b. We may assume that $\mathcal{S}(v_t)$ and $\mathcal{S}(v_b)$ contain no corner of $\mathcal{S}(w)$, otherwise $w \preceq_{1/2} v_t$ or $w \preceq_{1/2} v_b$. Since $\mathcal{S}(w)$ intersects $\mathcal{S}(v_t)$ and $\mathcal{S}(v_b)$, the largest of $\mathcal{S}(v_t)$ and $\mathcal{S}(v_b)$ covers at least half of a side of $\mathcal{S}(w)$. Hence $w \preceq_{1/2} v_t$ or $w \preceq_{1/2} v_b$. But then v_t and v_b form a $\preceq_{1/2}$-dominating set of U. It follows immediately from Lemma 8.3.1 that the $\preceq_{1/2}$-factor is at most $2r + 1$. □

Theorem 8.3.7 *Let $r \geq 3$ be an odd integer. There is a polynomial-time* $\mathrm{O}(r^2)$*-approximation algorithm for Minimum Dominating Set on intersection graphs of homothetic regular r-polygons.*

Proof: The case when $r = 3$ follows from Theorem 8.3.4. So let $G = G[\mathcal{S}]$ be the intersection graph of a set $\mathcal{S}$ of homothetic regular r-polygons for some odd integer $r \geq 5$. Observe that all $\preceq_{1/2}$-larger neighbors of a $u \in V(G)$ can be hit by the corners of $\mathcal{S}(u)$ and the midpoint of each side. Then Theorem 8.1.4 and Theorem 8.2.3 immediately give the theorem. □

Furthermore, we can adapt Lemma 8.3.6 to bound the fractional $\preceq_{1/2}$-factor. Therefore the integrality gap of the standard LP of Minimum Dominating Set on intersection graphs of homothetic regular r-polygons for odd integers $r \geq 3$ is $\mathrm{O}(r^2)$ as well by Theorem 8.2.5.

For homothetic even regular polygons, we use a completely different relation to improve on the approximation ratio attained by the algorithm of Theorem 8.3.4. We require the following consequence of Lemma 8.3.1. A binary relation $\preceq$ is a *preorder* if it is both reflexive and transitive. It is *total* if $u \preceq v$ or $v \preceq u$ for any pair u, v.

Lemma 8.3.8 *Let $\preceq$ be a total preorder on the vertices of a graph G such that for any $u \in V(G)$ the cardinality of any independent set of $N_{\preceq}(u)$ is bounded by c. Then the $\preceq$-factor is at most $c+1$.*

Proof: Find a $\preceq$-dominating set of $N_{\preceq}(u)$ as follows. Since $\preceq$ is a total preorder, there is a $v \in N_{\preceq}(u)$ that is maximum, i.e. $w \preceq v$ for each $w \in N_{\preceq}(u)$. Observe that v $\preceq$-dominates $N(v) \cap N_{\preceq}(u)$. Now remove $N[v]$ from $N_{\preceq}(u)$ and iterate. This yields a $\preceq$-dominating set of $N_{\preceq}(u)$ that is also an independent set. Hence it has cardinality at most c. It follows from Lemma 8.3.1 that the $\preceq$-factor is at most $c+1$. $\square$

A similar lemma can be proved for the fractional $\preceq$-factor.

Now let $G = G[\mathcal{S}]$ be the intersection graph of a collection $\mathcal{S}$ of homothetic regular r-polygons for some even integer $r \geq 2$. Define a total preorder $\preceq_{\mathsf{Leb}}$ on $V(G)$ such that $u \preceq_{\mathsf{Leb}} v$ for $u, v \in V(G)$ if the Lebesgue measure of $\mathcal{S}(u)$ is at most the Lebesgue measure of $\mathcal{S}(v)$.

Lemma 8.3.9 *Let $G = G[\mathcal{S}]$ be the intersection graph of a collection $\mathcal{S}$ of homothetic convex compact sets in $\mathbb{R}^2$. Then the $\preceq_{\mathsf{Leb}}$-factor is at most 5 if $\mathcal{S}$ is a collection of homothetic parallelograms and at most 6 otherwise.*

Proof: Following Lemma 8.3.8, it suffices to bound the cardinality of any independent set of $N_{\preceq_{\mathsf{Leb}}}(u)$ for each $u \in V(G)$ by 4 and 5 respectively. So for some $u \in V(G)$, define a set $\mathcal{S}' = \{\mathcal{S}'(v) \mid v \in N_{\preceq_{\mathsf{Leb}}}[u]\}$ of translated copies of $\mathcal{S}(u)$ such that $\mathcal{S}'(v) \subseteq \mathcal{S}(v)$ and $\mathcal{S}'(v) \cap \mathcal{S}'(u) \neq \emptyset$ for each $v \in N_{\preceq_{\mathsf{Leb}}}[u]$. An independent set of $N_{\preceq_{\mathsf{Leb}}}(u)$ corresponds to one of $G[\mathcal{S}']$, and vice versa.

We now apply a result of Kim, Kostochka, and Nakprasit [167], who showed that if H is the intersection graph of a set of translated copies of a fixed convex compact set in the plane with $\omega(H) \geq 2$, then the maximum degree of H is at most $4\omega(H) - 4$ if this fixed set is a parallelogram and at most $6\omega(H) - 7$ otherwise, where $\omega(H)$ is the cardinality of a maximum clique of H. Let H' be the subgraph of $G[S']$ induced by u and any independent set of $G[\mathcal{S}']$ (i.e. of $N_{\preceq_{\mathsf{Leb}}}(u)$). Then $\omega(H') = 2$ and thus the degree of u in H' is bounded by 4 and 5 respectively. The lemma follows. $\square$

Note that the bounds of this lemma are tight, as demonstrated by a suitable representation of $K_{1,5}$ and $K_{1,6}$ respectively.

Theorem 8.3.10 *Let $r \geq 2$ be an even integer. There is a polynomial-time $\mathrm{O}(r)$-approximation algorithm for Minimum Dominating Set on intersection graphs of homothetic regular r-polygons.*

Proof: Use the relation $\preceq_{\mathsf{Leb}}$. Lemma 8.3.9 proved that the $\preceq_{\mathsf{Leb}}$-factor is at most 6. All $\preceq_{\mathsf{Leb}}$-larger neighbors of a vertex can be hit by placing a point on each corner of the corresponding polygon. The theorem statement then follows from Theorem 8.1.4 and Theorem 8.2.3. □

It follows from Theorem 8.2.5 that the integrality gap of the standard LP of Minimum Dominating Set is $\mathrm{O}(r)$ on intersection graphs of homothetic regular r-polygons for even integers $r \geq 2$.

Note that although the algorithm of Theorem 8.3.10 also applies to Minimum Dominating Set on intersection graphs of homothetic regular 2-polygons (i.e. interval graphs), a linear-time exact algorithm exists in this case [61] and the integrality gap of the standard LP is 1 [47].

8.3.3 More General Objects

Observe that the proof of Theorem 8.3.10 goes through for arbitrary homothetic parallelograms. In fact, we can extend Theorem 8.3.7 and Theorem 8.3.10 to the following theorem. An *affine regular polygon* is any polygon that can be obtained from a regular polygon by an invertible affine transformation.

Theorem 8.3.11 *For any integer $r \geq 2$, there is a polynomial-time approximation algorithm for Minimum Dominating Set on intersection graphs of homothetic affine regular r-polygons, attaining approximation ratio $\mathrm{O}(r)$ if r is even and $\mathrm{O}(r^2)$ otherwise.*

Proof: Let $\mathcal{S}$ be a collection of homothetic affine regular r-polygons for some $r \geq 2$. Apply the inverse affine transformation to transform $\mathcal{S}$ into a collection $\mathcal{S}'$ of homothetic regular r-polygons and note that $G[\mathcal{S}] = G[\mathcal{S}']$. The theorem statement is now immediate from Theorem 8.3.7 and Theorem 8.3.10. □

A consequence of this result is a constant-factor approximation algorithm for intersection graphs of homothetic rectangles. By placing a mild restriction on the type of rectangles, we can drop the homotheticity constraint.

We consider intersection graphs of axis-parallel rectangles whose aspect-ratio is constant. The *aspect-ratio* of a rectangle is the length of its longest side divided by the length of its shortest side.

Lemma 8.3.12 *Let $\mathcal{S}$ be a collection of axis-parallel rectangles with aspect-ratio at most c for some $c \geq 1$. Then for any $\epsilon > 0$, one can obtain a binary ϵ-net of size $\mathrm{O}(c/\epsilon)$ in time polynomial in $|\mathcal{S}|$ and c/ϵ.*

Proof: Construct a set of homothetic squares $\mathcal{S}'$ by replacing each rectangle $s \in \mathcal{S}$ by at most c axis-parallel squares, the union of which is precisely s. Now use Theorem 8.1.4 to find an ϵ'-net for $\mathcal{S}'$, where $\epsilon' = \epsilon/c$. □

Theorem 8.3.13 *For any integer $c \geq 1$, there is a polynomial-time $\mathrm{O}(c^3)$-approximation algorithm for Minimum Dominating Set on intersection graphs of axis-parallel rectangles with aspect-ratio at most c.*

Proof: Let $G = G[\mathcal{S}]$ be the intersection graph of a collection $\mathcal{S}$ of axis-parallel rectangles with aspect-ratio at most c, for some integer $c \geq 1$. Consider a $\preceq_{\mathsf{Leb}}$-larger neighbor v of some vertex u. Without loss of generality, $\mathcal{S}(u)$ is a $1 \times c$ rectangle. Then all sides of $\mathcal{S}(v)$ have length at least 1 and $\mathcal{S}(v)$ contains a corner of $\mathcal{S}(u)$ or covers at least a $1/c$-fraction of a long side of $\mathcal{S}(u)$. Hence $2c + 2$ points in $\mathcal{S}(u)$ suffice to hit all $\preceq_{\mathsf{Leb}}$-larger neighbors. But then any independent set of $N_{\preceq_{\mathsf{Leb}}}(u)$ has cardinality at most $2c+2$ and the $\preceq_{\mathsf{Leb}}$-factor is at most $2c+3$ by Lemma 8.3.8. The theorem now follows from Lemma 8.3.12 and Theorem 8.2.3. □

The result of Theorem 8.3.13 does not seem to extend to similarly defined variants of regular pentagons, regular hexagons, or other regular polygons.

To show that Theorem 8.2.3 may also be applied beyond two dimensions, we prove the following theorem about Minimum Dominating Set on intersection graphs of translated copies of an affine three-dimensional box. We should note that the results of Section 6.3.3 imply the existence of a ptas for this case.

Theorem 8.3.14 *There exists a constant-factor approximation algorithm for Minimum Dominating Set on intersection graphs of translated copies of an affine three-dimensional box.*

Proof: Using the idea of the proof of Theorem 8.3.11, we may assume that we are given the intersection graph $G = G[\mathcal{S}]$ of a set $\mathcal{S}$ of translated copies of a three-dimensional box. It is easy to see that the $\preceq_{\mathsf{Leb}}$-factor is at most 9 by Lemma 8.3.8 and that any $\preceq_{\mathsf{Leb}}$-larger neighborhood can be hit by 8 points. Hence, following a result by Laue [188] (see Theorem 8.1.5), we may apply Theorem 8.2.3 with a linear function g and the theorem follows. □

Since Theorem 8.1.5 applies to translated copies of any fixed three-dimensional polytope, it seems likely that the above theorem could be extended to more general or more complex three-dimensional objects.

8.4 Disk Graphs of Bounded Ply

The obvious class of intersection graphs missing in the above discussion is the class of disk graphs. We proved in Chapter 6 that Minimum Dominating Set has a ptas on unit disk graphs, but this scheme does not carry over to general disk graphs. The ideas developed in Chapter 6 also seem to be insufficient to handle this problem. Finally, even though the $\preceq_{\mathsf{Leb}}$-factor is at most 6 for disk graphs, we do not know how to apply Theorem 8.2.3. The problem (when using $\preceq_{\mathsf{Leb}}$) is that we cannot choose a constant number of points inside a disk to hit all $\preceq_{\mathsf{Leb}}$-larger neighbors. All $\preceq_{\mathsf{Leb}}$-larger neighbors of a disk can be hit by a constant number of points, but some would have to lie outside the disk. Unfortunately, Theorem 8.2.3 does not seem to extend to this case.

If we know however that the ply of the set of disks representing the disk graph is bounded, then the above techniques do work and we obtain a constant-factor approximation algorithm. We give these algorithms below, in order of descending approximation ratio. Recall that the *ply* of a set of objects is the maximum over all points p of the number of objects strictly containing p.

8.4.1 Ply-Dependent Approximation Ratio

The approximation ratio of the first approximation algorithms we present depend (linearly) on the ply of the set of disks representing the disk graph.

Lemma 8.4.1 *Given a set of disks of ply γ, the cardinality of the closed $\preceq_{\mathsf{Leb}}$-larger neighborhood of any disk is at most 9γ.*

The proof uses an area bound in a manner similar to Lemma 7.1.1 (see also Miller et al. [210]). We can now immediately prove the following result.

Theorem 8.4.2 *There is a polynomial-time $\mathrm{O}(\gamma)$-approximation algorithm for Minimum Dominating Set on disk graphs of ply γ.*

Proof: By Lemma 8.4.1, all $\preceq_{\mathsf{Leb}}$-larger neighbors of a disk can be hit by at most 9γ points. Lemma 8.3.9 shows that the $\preceq_{\mathsf{Leb}}$-factor is at most 6. The theorem now follows from Theorem 8.1.4 and Theorem 8.2.3. □

A different technique improves on Theorem 8.4.2. We essentially give a second general approach to approximate Minimum Dominating Set using $\preceq$-dominating sets, but this time without using ϵ-nets.

Theorem 8.4.3 *Let G be a graph and let $\preceq$ be a binary reflexive relation on the vertices of G with fractional $\preceq$-factor at most c_3. Suppose that the maximum cardinality of the $\preceq$-larger closed neighborhood of any $u \in V(G)$ is at most c_2. Then the integrality gap of the standard LP for Minimum Dominating Set on G is at most c_2c_3.*

Proof: From Definition 8.2.4, the integrality gap of (the standard LP for) Minimum $\preceq$-Dominating Set on G multiplied by the fractional $\preceq$-factor is an upper bound to the integrality gap of (the standard LP for) Minimum Dominating Set on G. By assumption, the fractional $\preceq$-factor is at most c_3. Hence it suffices to bound the integrality gap of Minimum $\preceq$-Dominating Set on G.

We transform the minimum $\preceq$-dominating set problem on G to an instance of Minimum Set Cover. Let $\mathbb{U} = V(G)$ and $\mathcal{S} = \{\mathcal{S}(v) \mid v \in V(G)\}$ where $\mathcal{S}(v) = \{u \mid v \in N_{\preceq}[u]\}$. Hochbaum [148] showed that the integrality gap of Minimum Set Cover is bounded by the element frequency. The element frequency of $(\mathbb{U}, \mathcal{S})$ is at most the maximum cardinality of any $\preceq$-larger closed neighborhood of G, which is at most c_2 by assumption.

Observe that a (fractional) $\preceq$-dominating set of G corresponds directly to a (fractional) set cover of $(\mathbb{U}, \mathcal{S})$ and vice versa. Hence the integrality gap of Minimum $\preceq$-Dominating Set on G is at most c_2. This gives a bound on the integrality gap of Minimum Dominating Set on G of c_2c_3. $\square$

Theorem 8.4.4 *The integrality gap of the standard LP for Minimum Dominating Set on disk graphs of ply γ is at most 54γ. If $\gamma = 1$, then the gap is at most 42. Hence the gap on planar graphs is at most 42.*

Proof: By Lemma 8.3.9, the fractional $\preceq_{\mathsf{Leb}}$-factor is at most 6. The maximum cardinality of any $\preceq_{\mathsf{Leb}}$-larger closed neighborhood of G is at most 9γ by Lemma 8.4.1. Hence the gap is at most 54γ by Theorem 8.4.3. If $\gamma = 1$, then the maximum cardinality of any $\preceq_{\mathsf{Leb}}$-larger closed neighborhood of $\mathcal{S}$ is at most 7, yielding the bound of 42 on the gap. As planar graphs are disks graphs of ply 1 [169, 210], the gap on planar graphs is at most 42. $\square$

Although a ptas for Minimum Dominating Set on planar graphs is known [22], we are not aware of any previous results on the integrality gap of the standard LP for Minimum Dominating Set on this class of graphs.

The reduction from Minimum $\preceq$-Dominating Set to Minimum Set Cover given in the proof of Theorem 8.4.3 can be exploited algorithmically.

Theorem 8.4.5 *Let G be a graph and let $\preceq$ be a binary reflexive relation on the vertices of G with $\preceq$-factor at most c_1. Suppose that the maximum cardinality of the $\preceq$-larger closed neighborhood of any $u \in V(G)$ is at most c_2. Then there is a linear-time c_1c_2-approximation algorithm for Minimum Dominating Set on G.*

Proof: Transform the minimum $\preceq$-dominating set instance on G to an instance of Minimum Set Cover, as in Theorem 8.4.3. Bar-Yehuda and Even [24] proved that Minimum Set Cover has a linear-time approximation algorithm with approximation ratio at most the maximum element frequency. Following the proof of Theorem 8.4.3, the maximum element frequency is at most c_2. As the $\preceq$-factor is at most c_1, the theorem follows. $\square$

Using the proof of Theorem 8.4.4, we can then show the following.

Theorem 8.4.6 *There exists a linear-time (54γ)-approximation algorithm for Minimum Dominating Set on disk graphs of ply γ.*

Note that the approximation ratio improves to 42 on disk graphs of ply 1, i.e. on planar graphs.

Theorem 8.4.3 and 8.4.5 also have implications for Minimum Dominating Set on general graphs. Following Lemma 8.3.1 and 8.3.2, the (fractional) $\preceq$-factor of any relation $\preceq$ is at most the maximum cardinality of any $\preceq$-larger closed neighborhood of G.

Corollary 8.4.7 *Let G be a graph and let $\preceq$ be a binary reflexive relation on the vertices of G. Suppose that the maximum cardinality of the $\preceq$-larger closed neighborhood of any $u \in V(G)$ is at most c. Then the integrality gap of the standard LP for Minimum Dominating Set on G is at most c^2. Moreover, there is a linear-time c^2-approximation algorithm for Minimum Dominating Set on G.*

Clearly, $c \leq \Delta(G)$ for any relation $\preceq$, yielding an integrality gap of $\Delta^2(G)$ and a linear-time $\Delta^2(G)$-approximation algorithm for Minimum Dominating Set on any graph G. This is far worse than the $(1+\ln \Delta(G))$-approximation algorithm for Minimum Dominating Set known in the literature [156, 197, 66, 149]. One could however imagine that a relation $\preceq$ for which c is minimum over all relations $\preceq$ beats this bound.

Theorem 8.4.8 *Let G be a graph. We can find in polynomial time a binary reflexive relation $\preceq$ such that the maximum cardinality of any $\preceq$-larger closed neighborhood of G is minimized.*

Proof: First note that there is an asymmetric binary reflexive relation $\preceq$ attaining the minimum. Now observe that an asymmetric binary reflexive relation $\preceq$ on G corresponds to an orientation $\vec{G}$ of G and vice versa. Simply let $u \preceq v$ if and only if there is a directed edge from u to v in $\vec{G}$. Hence it suffices to find an orientation $\vec{G}$ of G minimizing the maximum out-degree of any vertex. Using a result of Frank and Gyárfás [111], such an orientation can be found in polynomial time. □

If an upper bound to the maximum cardinality of any $\preceq$-larger closed neighborhood of G is known for some relation $\preceq$, then we can bound the approximation ratio of the algorithm of Corollary 8.4.7.

Theorem 8.4.9 *There exists a linear-time $(9\gamma)^2$-approximation algorithm for Minimum Dominating Set on disk graphs of ply γ, even if no representation of the graph is given.*

Proof: By Lemma 8.4.1, a disk graph G of ply γ has a binary reflexive relation $\preceq$ for which the maximum cardinality of any $\preceq$-larger closed neighborhood of G is at most 9γ, namely $\preceq_{\mathsf{Leb}}$. The theorem follows from Theorem 8.4.8 and Corollary 8.4.7. □

Note that to apply the approximation algorithm, one does not need to know the ply of the given disk graph. The fact that the graph has a disk representation of ply γ only turns up in the analysis of the approximation factor.

8.4.2 A Constant Approximation Ratio

We can improve the approximation ratio further by using the shifting technique. We show that Minimum $\preceq_{\mathsf{Leb}}$-Dominating Set on n-vertex disk graphs of bounded ply, i.e. of ply $\gamma = \gamma(n) = \mathrm{o}(\log n)$, has an eptas. Because the $\preceq_{\mathsf{Leb}}$-factor is at most 6, this implies the existence of a $(6+\epsilon)$-approximation algorithm for Minimum Dominating Set on such disk graphs.

We use the shifting technique in the way outlined in Chapter 7. Assume that we are given a set of disks $\mathcal{D}$ such that the smallest disk has radius 1/2. We aim to find a small $\preceq_{\mathsf{Leb}}$-dominating set of $G = G[\mathcal{D}]$.

Partition the disks into levels. A disk of radius r has level j ($j \in \mathbb{Z}_{\geq 0}$) if $2^{j-1} \leq r < 2^j$. The level of the largest disk is denoted by l. The set $\mathcal{D}_{=j}$ is defined as the set of disks in $\mathcal{D}$ having level j. Similarly, we can define $\mathcal{D}_{\geq j}$ as the set of disks in $\mathcal{D}$ having level at least j, and so on.

For each level j, define a grid induced by horizontal lines $y = hk2^j$ and vertical lines $x = vk2^j$ ($h, v \in \mathbb{Z}$) for some odd integer $k \geq 7$, whose value we determine later. The grid formed in this way partitions the plane into squares of size $k2^j \times k2^j$, called *j-squares*. Furthermore, any j-square is contained in precisely one $(j+1)$-square and each $(j+1)$-square contains exactly four j-squares, denoted by $S_1, \ldots, S_4$. These four squares are *siblings* of each other. The set of disks intersecting a j-square S is denoted by $\mathcal{D}^S$, while the set of disks intersecting the boundary of S is denoted by $\mathcal{D}^{\mathrm{b}(S)}$. Similarly, $\mathcal{D}^{\mathrm{i}(S)} = \mathcal{D}^S - \mathcal{D}^{\mathrm{b}(S)}$ is the set of disks fully contained in the interior of S, $\mathcal{D}^{\mathrm{c}(S)}$ denotes the set of disks whose center is contained in S, and $\mathcal{D}^{+(S)} = \bigcup_{i=1}^4 \mathcal{D}^{\mathrm{b}(S_i)} - \mathcal{D}^{\mathrm{b}(S)}$ is the set of disks intersecting the boundary of at least one of the four children of S, but not the boundary of S itself. The meaning of combinations such as $\mathcal{D}_{\leq j}^{\mathrm{b}(S)}$ should be self-explaining. The level of a square S is denoted by $j(S)$.

Similarly, let $\mathcal{D}^{\mathrm{or}(S)}$ denote the set of disks having their center outside a j-square S and intersecting a band of width 2^j along the outer boundary of S. This band is called the *outer ring* of S. We also define several *inner rings*. Let $\mathcal{D}^{\mathrm{ir}_{j'}(S)} \subseteq \mathcal{D}^{\mathrm{c}(S)}$ denote the set of disks having their center inside S and intersecting a band of width $2^{j'}$ along the inner boundary of S. Observe that this implies that $\mathcal{D}^{\mathrm{ir}_{j(S)+\lceil \log k \rceil}(S)} = \mathcal{D}^{\mathrm{c}(S)}$. For convenience, we also define $\mathcal{D}^{\overline{\mathrm{ir}}(S)} = \bigcup_{j' \geq 0} \mathcal{D}_{\geq j'}^{\mathrm{ir}_{j'}(S)} = \bigcup_{j' \geq 0} \mathcal{D}_{=j'}^{\mathrm{ir}_{j'}(S)}$. Now define $\mathcal{D}^{\overline{+\mathrm{r}}(S)} = \bigcup_{i=1}^4 \mathcal{D}^{\overline{\mathrm{ir}}(S_i)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}$, extending the notion of $\mathcal{D}^{+(S)}$ we had before.

We now prove the following auxiliary theorem. Let $\mathcal{D}$ be a set of disks of ply γ and let OPT be a $\preceq_{\mathsf{Leb}}$-dominating set of $\mathcal{D}$ of minimum cardinality.

Theorem 8.4.10 *Let $\mathcal{D}$ be a set of disks of ply γ and $k \geq 7$ an odd positive integer. Then in $\mathrm{O}(k^2 n^2\, 2^{(80k-68)\gamma/\pi} 3^{(64k-60)\gamma/\pi})$ time, we can find a $\preceq_{\mathsf{Leb}}$-dominating set C of $\mathcal{D}$ such that $|C| \leq \sum_S \left(\left| OPT_{=j(S)}^{\mathrm{c}(S)} \right| + \left| OPT_{=j(S)}^{\mathrm{or}(S)} \right| \right)$, where the sum is over all squares S.*

We perform bottom-up dynamic programming on the j-squares. Observe that for each j-square S, disks in $\mathcal{D}_{\leq j}^{\mathrm{c}(S)}$ can be $\preceq_{\mathsf{Leb}}$-dominated by disks in $\mathcal{D}^{\mathrm{c}(S)}$ and $\mathcal{D}^{\mathrm{or}(S)}$. Following the approach developed in Chapter 7, we consider the status of disks in $\mathcal{D}_{>j}^{\mathrm{or}(S)}$. However, the outer ring of a j-square might partially overlap sibling j-squares, creating a problem when 'gluing' results together. Therefore we also consider the status of disks in the inner ring(s).

During the dynamic programming, we compute a $\preceq_{\mathsf{Leb}}$-dominating set of $\mathcal{D}^{c(S)} - \mathcal{D}_{>j}^{\overline{ir}(S)}$, given the status of disks in $\mathcal{D}_{>j}^{or(S)} \cup \mathcal{D}_{>j}^{\overline{ir}(S)}$ and using disks in $\mathcal{D}^{c(S)} - \mathcal{D}_{>j}^{\overline{ir}(S)}$ and $\mathcal{D}_{\leq j}^{or(S)}$. A disk in $\mathcal{D}_{>j}^{or(S)}$ is either in the dominating set, or it is not. A disk in $\mathcal{D}_{>j}^{\overline{ir}(S)}$ has three possible statuses: either it is in the dominating set, or it is $\preceq_{\mathsf{Leb}}$-dominated by a disk in $\mathcal{D}_{>j}^{or(S)} \cup \mathcal{D}_{>j}^{c(S)}$, or it is $\preceq_{\mathsf{Leb}}$-dominated by a yet undetermined disk. We define for each j-square S and any two disjoint sets $W_1 \subseteq \mathcal{D}_{>j}^{or(S)} \cup \mathcal{D}_{>j}^{\overline{ir}(S)}$ and $W_2 \subseteq \mathcal{D}_{>j}^{\overline{ir}(S)}$ the function $\mathsf{size}(S, W_1, W_2)$ as

$$\min\left\{ |T| \;\middle|\; T \subseteq \mathcal{D}_{=j}^{or(S)} \cup \left(\mathcal{D}^{c(S)} - \mathcal{D}_{>j}^{\overline{ir}(S)}\right); \right.$$
$$\left. W_1 \cup T \preceq_{\mathsf{Leb}}\text{-dominates } W_2 \cup \left(\mathcal{D}^{c(S)} - \mathcal{D}_{>j}^{\overline{ir}(S)}\right)\right\}$$

if $j = 0$ and

$$\min\left\{ |U| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \left(W_1 \cup U\right) \cap \left(\mathcal{D}_{>j-1}^{or(S_i)} \cup \mathcal{D}_{>j-1}^{\overline{ir}(S_i)}\right), X_i\right) \;\middle|\; \right.$$
$$U \subseteq \mathcal{D}_{>j-1}^{\overline{+r}(S)} \cup \mathcal{D}_{=j}^{or(S)} \cup \mathcal{D}_{=j}^{ir_j(S)}$$
$$\left. X_i = \left(\left(W_2 \cup \mathcal{D}_{>j-1}^{\overline{+r}(S)} \cup \mathcal{D}_{=j}^{ir_j(S)}\right) - N_{\succeq_{\mathsf{Leb}}}\left[W_1 \cup U\right]\right) \cap \mathcal{D}_{>j-1}^{\overline{ir}(S_i)} \right\}$$

if $j > 0$. Here the minimum over an empty set is ∞. Let $\mathsf{sol}(S, W_1, W_2)$ be the subset of $\mathcal{D}$ attaining $\mathsf{size}(S, W_1, W_2)$, or $\emptyset$ if $\mathsf{size}(S, W_1, W_2) = \infty$. The meaning of W_1 and W_2 is as follows. The disks in W_1 are dominators, whereas the disks in W_2 need to be $\preceq_{\mathsf{Leb}}$-dominated by disks in W_1 or $\mathcal{D}^{c(S)} - \mathcal{D}_{>j}^{\overline{ir}(S)}$.

Properties of the size- and sol-Functions

Functions size and sol are reasonably easy to compute, as we will show later. First, we prove that the size and sol functions attain the properties set forth in Theorem 8.4.10.

Lemma 8.4.11 $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset, \emptyset) \leq \sum_S \left(\left|\mathcal{C}_{=j(S)}^{c(S)}\right| + \left|\mathcal{C}_{=j(S)}^{or(S)}\right|\right)$, *where* $\mathcal{C}$ *is any* $\preceq_{\mathsf{Leb}}$*-dominating set.*

Proof: We apply induction on the level j and show that the following invariant holds for any j-square S:

$$\mathsf{size}\left(S, \left(\mathcal{C}_{>j}^{or(S)} \cup \mathcal{C}_{>j}^{\overline{ir}(S)}\right), \mathcal{D}_{>j}^{\overline{ir}(S)} - \mathcal{C}_{>j}^{\overline{ir}(S)} - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{c(S)} - \mathcal{C}^{or(S)}\right]\right)$$
$$\leq \left|\mathcal{C}_{>j}^{c(S)}\right| - \left|\mathcal{C}_{>j}^{\overline{ir}(S)}\right| + \sum_{S' \subseteq S} \left(\left|\mathcal{C}_{=j(S')}^{c(S')}\right| + \left|\mathcal{C}_{=j(S')}^{or(S')}\right|\right).$$

For $j = 0$, the invariant holds from the definition of size, as

$$\begin{aligned} |T| &\leq \left|\mathcal{C}^{\mathrm{or}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{c}(S)}\right| - \left|\mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j}\right| \\ &= \left|\mathcal{C}^{\mathrm{or}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{c}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{c}(S)}_{>j}\right| - \left|\mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j}\right|. \end{aligned}$$

So assume that $j > 0$ and that the invariant holds for all j-squares with $j' < j$. Then from the description of size and by applying induction,

$$\begin{aligned} &\mathsf{size}\Big(S, \left(\mathcal{C}^{\mathrm{or}(S)}_{>j} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j}\right), \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j} - \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j} - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{\mathrm{c}(S)} - \mathcal{C}^{\mathrm{or}(S)}\right]\Big) \\ \leq\;& \left|\mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{or}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{ir}_j(S)}_{=j}\right| \\ &+ \sum_{i=1}^{4} \mathsf{size}\Big(S_i, \left(\mathcal{C}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right), \\ &\qquad\qquad \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} - \mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{\mathrm{c}(S_i)} - \mathcal{C}^{\mathrm{or}(S_i)}\right]\Big) \\ \leq\;& \left|\mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{or}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{ir}_j(S)}_{=j}\right| \\ &+ \sum_{i=1}^{4}\left(\left|\mathcal{C}^{\mathrm{c}(S_i)}_{>j-1}\right| - \left|\mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right|\right) + \sum_{i=1}^{4}\sum_{S'_i \subseteq S_i}\left(\left|\mathcal{C}^{\mathrm{c}(S'_i)}_{=j(S'_i)}\right| + \left|\mathcal{C}^{\mathrm{or}(S'_i)}_{=j(S'_i)}\right|\right) \\ =\;& \left|\mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{or}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{ir}_j(S)}_{=j}\right| \\ &+ \left|\mathcal{C}^{\mathrm{c}(S)}_{>j-1}\right| - \left|\mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j-1}\right| - \left|\mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right| + \sum_{i=1}^{4}\sum_{S'_i \subseteq S_i}\left(\left|\mathcal{C}^{\mathrm{c}(S'_i)}_{=j(S'_i)}\right| + \left|\mathcal{C}^{\mathrm{or}(S'_i)}_{=j(S'_i)}\right|\right) \\ =\;& \left|\mathcal{C}^{\mathrm{c}(S)}_{>j}\right| - \left|\mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j}\right| + \left|\mathcal{C}^{\mathrm{c}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{or}(S)}_{=j}\right| \\ &+ \sum_{i=1}^{4}\sum_{S'_i \subseteq S_i}\left(\left|\mathcal{C}^{\mathrm{c}(S'_i)}_{=j(S'_i)}\right| + \left|\mathcal{C}^{\mathrm{or}(S'_i)}_{=j(S'_i)}\right|\right) \\ =\;& \left|\mathcal{C}^{\mathrm{c}(S)}_{>j}\right| - \left|\mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j}\right| + \sum_{S' \subseteq S}\left(\left|\mathcal{C}^{\mathrm{c}(S')}_{=j(S')}\right| + \left|\mathcal{C}^{\mathrm{or}(S')}_{=j(S')}\right|\right). \end{aligned}$$

The first inequality above is the crucial one. We give an explicit proof. Let $W_1 = \mathcal{C}^{\mathrm{or}(S)}_{>j} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j}$, $W_2 = \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j} - \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j} - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{\mathrm{c}(S)} - \mathcal{C}^{\mathrm{or}(S)}\right]$, and $U = \mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{C}^{\mathrm{or}(S)}_{=j} \cup \mathcal{C}^{\mathrm{ir}_j(S)}_{=j}$. We claim that the inequality holds for this choice of U.

First we show that $(W_1 \cup U) \cap \left(\mathcal{D}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right) = \mathcal{C}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}$ for any $i = 1, \ldots, 4$. Note that

$$W_1 \cup U \;=\; \mathcal{C}^{\mathrm{or}(S)}_{>j} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j} \cup \mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{C}^{\mathrm{or}(S)}_{=j} \cup \mathcal{C}^{\mathrm{ir}_j(S)}_{=j}$$

$$\begin{aligned} &= \mathcal{C}^{\mathrm{or}(S)}_{>j-1} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j-1} \cup \mathcal{C}^{\overline{+\mathrm{r}}(S)}_{>j-1} \\ &= \mathcal{C}^{\mathrm{or}(S)}_{>j-1} \cup \bigcup_{i=1}^{4} \mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}. \end{aligned}$$

But then $(W_1 \cup U) \cap \left(\mathcal{D}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right) = \mathcal{C}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}$ for any $i = 1, \ldots, 4$.

For the third parameter, we observe that for any $i = 1, \ldots, 4$

$$\begin{aligned} X_i &= \left(\left(W_2 \cup \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{ir}_j(S)}_{=j}\right) - N_{\succeq_{\mathsf{Leb}}}[W_1 \cup U]\right) \cap \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} \\ &= \Bigg(\left(\mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j} \cup \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{ir}_j(S)}_{=j}\right) - N_{\succeq_{\mathsf{Leb}}}\left[W_1 \cup U\right] \\ &\qquad - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{\mathrm{c}(S)} - \mathcal{C}^{\mathrm{or}(S)}\right]\Bigg) \cap \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} \\ &= \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} - N_{\succeq_{\mathsf{Leb}}}\left[W_1 \cup U\right] - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{\mathrm{c}(S)} - \mathcal{C}^{\mathrm{or}(S)}\right] \\ &\subseteq \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} - \mathcal{C}^{\overline{\mathrm{ir}}(S_i)}_{>j-1} - N_{\succeq_{\mathsf{Leb}}}\left[\mathcal{C} - \mathcal{C}^{\mathrm{c}(S_i)} - \mathcal{C}^{\mathrm{or}(S_i)}\right]. \end{aligned}$$

Because for any W and any $X_i \subseteq X'_i$, $\mathsf{size}(S_i, W, X_i) \le \mathsf{size}(S_i, W, X'_i)$, the first inequality is correct.

Since l is the level of the largest disk, for any j-square S with $j \ge l$, $\mathcal{C}^{\mathrm{or}(S)}_{>j} \cup \mathcal{C}^{\overline{\mathrm{ir}}(S)}_{>j} = \emptyset$, $\mathcal{D}^{\mathrm{c}(S)}_{>j} = \emptyset$, and $\mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j} = \emptyset$. Hence

$$\begin{aligned} \sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset, \emptyset) &\le \sum_{S;\, j(S)=l} \sum_{S' \subseteq S} \left(\left|\mathcal{C}^{\mathrm{c}(S')}_{=j(S')}\right| + \left|\mathcal{C}^{\mathrm{or}(S')}_{=j(S')}\right|\right) \\ &= \sum_{S} \left(\left|\mathcal{C}^{\mathrm{c}(S)}_{=j(S)}\right| + \left|\mathcal{C}^{\mathrm{or}(S)}_{=j(S)}\right|\right). \end{aligned}$$

This proves the lemma. □

It follows that if OPT is a minimum $\preceq_{\mathsf{Leb}}$-dominating set, then

$$\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset, \emptyset) \le \sum_{S} \left(\left|OPT^{\mathrm{c}(S)}_{=j(S)}\right| + \left|OPT^{\mathrm{or}(S)}_{=j(S)}\right|\right).$$

Lemma 8.4.12 $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S, \emptyset, \emptyset)$ *is a* $\preceq_{\mathsf{Leb}}$*-dominating set.*

Proof: For any j-square S and any two disjoint sets $W_1 \subseteq \mathcal{D}^{\mathrm{or}(S)}_{>j} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}$, $W_2 \subseteq \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}$, we claim that $W_1 \cup \mathsf{sol}(S, W_1, W_2)$ is a $\preceq_{\mathsf{Leb}}$-dominating set of $W_2 \cup (\mathcal{D}^{\mathrm{c}(S)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j})$ if $\mathsf{size}(S, W_1, W_2) \ne \infty$.

Apply induction on j. If $j = 0$, this follows trivially from the definition of size and sol. So assume that $j > 0$ and that the claim holds for all j'-squares with $j' < j$.

Suppose that $\mathsf{size}(S, W_1, W_2) \neq \infty$ for two disjoint sets $W_1 \subseteq \mathcal{D}^{\mathrm{or}(S)}_{>j} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}$, $W_2 \subseteq \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}$. Let $U \subseteq \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{or}(S)}_{=j} \cup \mathcal{D}^{\mathrm{ir}_j(S)}_{=j}$ attain the minimum in the definition of size for W_1 and W_2. Because $\mathsf{size}(S, W_1, W_2) \neq \infty$, it must be that $\mathsf{size}(S_i, W^i, X_i) \neq \infty$ for $i = 1, \ldots, 4$ as well, where

$$W^i = (W_1 \cup U) \cap \left(\mathcal{D}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right)$$

and

$$X_i = \left(\left(W_2 \cup \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{ir}_j(S)}_{=j}\right) - N_{\succeq_{\mathsf{Leb}}}\left[W_1 \cup U\right]\right) \cap \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}.$$

Then by induction, $W^i \cup \mathsf{sol}(S_i, W^i, X_i)$ is a $\preceq_{\mathsf{Leb}}$-dominating set of $X_i \cup \left(\mathcal{D}^{\mathrm{c}(S_i)} - \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right)$. Observe that

$$\begin{aligned}
&\bigcup_{i=1}^{4} W^i \cup \bigcup_{i=1}^{4} \mathsf{sol}(S_i, W^i, X_i) \\
&\quad = \bigcup_{i=1}^{4} \left((W_1 \cup U) \cap \left(\mathcal{D}^{\mathrm{or}(S_i)}_{>j-1} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right)\right) \cup \bigcup_{i=1}^{4} \mathsf{sol}(S_i, W^i, X_i) \\
&\quad \subseteq W_1 \cup U \cup \bigcup_{i=1}^{4} \mathsf{sol}(S_i, W^i, X_i) \\
&\quad = W_1 \cup \mathsf{sol}(S, W_1, W_2)
\end{aligned}$$

and

$$\bigcup_{i=1}^{4} \left(X_i \cup \left(\mathcal{D}^{\mathrm{c}(S_i)} - \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right)\right) = \bigcup_{i=1}^{4} X_i \cup \left(\mathcal{D}^{\mathrm{c}(S)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j-1} - \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right).$$

Since $W_1 \cup \mathsf{sol}(S, W_1, W_2)$ also $\preceq_{\mathsf{Leb}}$-dominates $N_{\succeq_{\mathsf{Leb}}}[W_1 \cup U]$, we can derive that $W_1 \cup \mathsf{sol}(S, W_1, W_2)$ $\preceq_{\mathsf{Leb}}$-dominates

$$\begin{aligned}
&\bigcup_{i=1}^{4} X_i \cup \left(\mathcal{D}^{\mathrm{c}(S)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j-1} - \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right) \cup N_{\succeq_{\mathsf{Leb}}}\left[W_1 \cup U\right] \\
&\quad \supseteq \bigcup_{i=1}^{4} \left(\left(W_2 \cup \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{ir}_j(S)}_{=j}\right) \cap \mathcal{D}^{\overline{\mathrm{ir}}(S_i)}_{>j-1}\right) \\
&\quad\quad \cup \left(\mathcal{D}^{\mathrm{c}(S)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j-1} - \mathcal{D}^{\overline{+\mathrm{r}}(S)}_{>j-1}\right) \\
&\quad = W_2 \cup \left(\mathcal{D}^{\mathrm{c}(S)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}\right).
\end{aligned}$$

From the previous lemma, we know that $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset, \emptyset) \neq \infty$. Hence $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S, \emptyset, \emptyset)$ is a $\preceq_{\mathsf{Leb}}$-dominating set of $\bigcup_{S;\, j(S)=l} \mathcal{D}^{\mathrm{c}(S)}$. Because each disk is in $\mathcal{D}^{\mathrm{c}(S)}$ for some l-square S, this is a $\preceq_{\mathsf{Leb}}$-dominating set of $\mathcal{D}$. □

Computing the size- and sol-Functions

We apply the methods outlined in Chapter 7. We show again that it is sufficient to size and sol for a limited number of j-squares.

The definition of nonempty and empty is slightly different than usual. We say that that a j-square S is *nonempty* if S or the outer ring of S is intersected by a level j disk and *empty* otherwise.

The definition of relevant remains the same, modulo the new definition of nonempty. A j-square S is *relevant* if one of its three siblings is nonempty or there is a nonempty square S' containing S, such that S' has level at most $j + \lceil \log k \rceil$ (so each nonempty j-square is relevant). Note that this definition induces $\mathrm{O}(k^2 n)$ relevant squares. A relevant square S is said to be a *relevant child* of another relevant square S' if $S \subset S'$ and there is no third relevant square S'', such that $S \subset S'' \subset S'$. Conversely, if S is a relevant child of S', S' is a *relevant parent* of S.

Lemma 8.4.13 *For each relevant* 0*-square* S*, all* size- *and* sol*-values for* S *can be computed in* $\mathrm{O}\big(nk\gamma\, 2^{(16k+32)\gamma/\pi} 3^{(40k-12)\gamma/\pi}\big)$ *time.*

Proof: We use area bounds to bound the cardinality of sets we are interested in. By Lemma 7.1.1, $\left|\mathcal{D}^{\mathrm{or}(S)}_{>j}\right| \le 16(k+2)\gamma/\pi$. To bound $\left|\mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}\right|$, note that $\left|\mathcal{D}^{\mathrm{ir}_{j+1}(S)}_{\ge j+1}\right| \le (20k-60)\gamma/\pi$ and that for any $j' > j$, $\left|\mathcal{D}^{\mathrm{ir}_{j'}(S)}_{\ge j'} - \mathcal{D}^{\mathrm{ir}_{j'-1}(S)}_{\ge j'-1}\right| \le (3 \cdot 2^{j-j'+3}k - 60)\gamma/\pi$. Hence

$$\left|\mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}\right| \le \left(\Big(20 + 3 \sum_{j'=j+2}^{\infty} 2^{j-j'+3}\Big)k - 60\right)\gamma/\pi \le (32k-60)\gamma/\pi.$$

Therefore we can enumerate all disjoint sets $W_1 \subseteq \mathcal{D}^{\mathrm{or}(S)}_{>j} \cup \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}$, $W_2 \subseteq \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}$ in $\mathrm{O}\big(2^{(16k+32)\gamma/\pi} 3^{(32k-60)\gamma/\pi}\big)$ time.

Using Lemma 7.1.2, the pathwidth of $\mathcal{D}^{\mathrm{or}(S)}_{=j} \cup \big(\mathcal{D}^{\mathrm{c}(S)} - \mathcal{D}^{\overline{\mathrm{ir}}(S)}_{>j}\big)$ can be bounded by $8(k+6)\gamma/\pi$. By adapting the algorithm of Corollary 5.3.9, we can find the set T required by the definition of size and sol in $\mathrm{O}(nk\gamma\, 3^{(8k+48)\gamma/\pi})$ time. The lemma follows. □

Now assume that the size- and sol-values of all relevant children of a j-square S are known.

Lemma 8.4.14 *For each relevant j-square S ($j > 0$) with relevant $(j-1)$-square children, in* $\mathrm{O}(2^{(80k-68)\gamma/\pi} 3^{(64k-60)\gamma/\pi})$ *time all* size- *and* sol*-values for* S *can be computed.*

Proof: Using similar ideas as in Lemma 8.4.13 and Lemma 7.1.1, we can show that $\left|\mathcal{D}^{\mathrm{or}(S)}_{\ge j}\right| \le (40k+60)\gamma/\pi$ and

$$\left|\mathcal{D}^{\overline{\mathrm{ir}}(S)}_{\ge j}\right| \le \left(\Big(40 + 3 \sum_{j'=j+1}^{\infty} 2^{j-j'+3}\Big)k - 60\right)\gamma/\pi \le (64k-60)\gamma/\pi.$$

Now bound $\left|\mathcal{D}_{\geq j}^{\overline{+\mathrm{r}}(S)}\right|$. Note that $\left|\bigcup_{i=1}^{4} \mathcal{D}_{\geq j}^{\mathrm{ir}_j(S_i)} - \mathcal{D}_{\geq j}^{\mathrm{ir}_j(S)}\right| \leq (32k - 128)\gamma/\pi$ and for any $j' > j$,

$$\begin{aligned}\left|\left(\bigcup_{i=1}^{4} \mathcal{D}_{\geq j'}^{\mathrm{ir}_{j'}(S_i)} - \mathcal{D}_{\geq j'}^{\mathrm{ir}_{j'}(S)}\right) - \left(\bigcup_{i=1}^{4} \mathcal{D}_{\geq j'-1}^{\mathrm{ir}_{j'-1}(S_i)} - \mathcal{D}_{\geq j'-1}^{\mathrm{ir}_{j'-1}(S)}\right)\right| \\ \leq \quad (2^{j-j'+3}k - 12) \cdot \gamma/\pi.\end{aligned}$$

Then

$$\begin{aligned}\left|\mathcal{D}_{\geq j}^{\overline{+\mathrm{r}}(S)}\right| &\leq \left(32k - 128 + \sum_{j'=j+1}^{\infty} 2^{j-j'+3}k\right) \cdot \gamma/\pi \\ &= (40k - 128)\gamma/\pi.\end{aligned}$$

The lemma now follows from the definition of size and sol. □

Lemma 8.4.15 *For each relevant j-square S ($j > 0$) with no relevant children of level $j-1$, all* size- *and* sol-*values for S can be computed in* $\mathrm{O}(n\, 2^{44\gamma/\pi} 3^{16\gamma/\pi})$ *time.*

Proof: Following the proof of Lemma 7.2.6, $\mathcal{D}_{\geq j}^{\overline{\mathrm{ir}}(S)} = \mathcal{D}_{\geq j+\lceil \log k \rceil}^{\overline{\mathrm{ir}}(S)}$. Then Lemma 7.1.1 shows that $\left|\mathcal{D}_{\geq j}^{\overline{\mathrm{ir}}(S)}\right| \leq \left|\mathcal{D}_{\geq j+\lceil \log k \rceil}^{\mathrm{c}(S)}\right| \leq 16\gamma/\pi$. Lemma 7.2.6 implies that $\left|\mathcal{D}_{\geq j}^{\overline{+\mathrm{r}}(S)}\right| = 0$ and $\mathcal{D}_{\geq j}^{\mathrm{or}(S)} = \mathcal{D}_{\geq j+\lceil \log k \rceil}^{\mathrm{or}(S)}$ and thus $\left|\mathcal{D}_{\geq j}^{\mathrm{or}(S)}\right| \leq 44\gamma/\pi$. Then from the proof of Lemma 7.2.6 and the definition of size and sol, we can compute all size- and sol-values in $\mathrm{O}(n\, 2^{44\gamma/\pi} 3^{16\gamma/\pi})$ time. □

Lemma 8.4.16 *The value of* $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset, \emptyset)$ *can be computed in time* $\mathrm{O}(k^2 n^2\, 2^{(80k-68)\gamma/\pi} 3^{(64k-60)\gamma/\pi})$.

Proof: Follows from Lemma 8.4.13, Lemma 8.4.14, Lemma 8.4.15, and the proof of Lemma 7.2.7. The number of relevant squares is $\mathrm{O}(k^2 n)$. □

Proof of Theorem 8.4.10: Follows directly from Lemmas 8.4.11, 8.4.12, and 8.4.16. □

The Approximation Algorithm

The shifting technique can now be applied as follows. For an integer a ($0 \leq a \leq k-1$), call a line of level j *active* if it has the form $y = (hk + a2^{l-j})2^j$ or $x = (vk + a2^{l-j})2^j$ ($h, v \in \mathbb{Z}$). The active lines partition the plane into j-squares as before, although shifted with respect to the shifting parameter a. However, we can still apply the algorithm of Theorem 8.4.10.

Let C_a denote the set returned by the algorithm for the j-squares induced by shifting parameter a ($0 \leq a \leq k-1$) and let $C_{\min}$ be a smallest such set.

Lemma 8.4.17 $|C_{\min}| \leq (1+24/k) \cdot |OPT|$, *where* OPT *is a minimum* $\preceq_{\mathsf{Leb}}$*-dominating set.*

Proof: Define $\mathcal{D}_a^{\mathrm{or}}$ as the set of disks intersecting the outer ring of a j-square S at their level, i.e. $\mathcal{D}_a^{\mathrm{or}} = \bigcup_S \mathcal{D}_{=j(S)}^{\mathrm{or}(S)}$. Clearly a disk of level j can be in $\mathcal{D}_a^{\mathrm{or}}$ for at most 8 values of a. Therefore $\sum_{a=0}^{k-1} |OPT \cap \mathcal{D}_a^{\mathrm{or}}| \leq 8 \cdot |OPT|$. Furthermore, for any fixed value of a, any level j disk can intersect the outer ring of at most 3 j-squares. It follows from Lemma 8.4.11 that

$$|C_a| \leq \sum_S \left(\left| \mathcal{C}_{=j(S)}^{\mathrm{c}(S)} \right| + \left| \mathcal{C}_{=j(S)}^{\mathrm{or}(S)} \right| \right) \leq |OPT| + 3|OPT \cap \mathcal{D}_a^{\mathrm{or}}|.$$

Then

$$k \cdot |C_{\min}| \leq \sum_{a=0}^{k-1} |C_a| \leq \sum_{a=0}^{k-1} (|OPT| + 3|OPT \cap \mathcal{D}_a^{\mathrm{or}}|) \leq (k+24) \cdot |OPT|.$$

Hence $|C_{\min}| \leq (1+24/k) \cdot |OPT|$. □

Combining Theorem 8.4.10 and Lemma 8.4.17, we obtain the following approximation scheme.

Theorem 8.4.18 *There is an eptas for Minimum* $\preceq_{\mathsf{Leb}}$*-Dominating Set on* n*-vertex disk graphs of bounded ply, i.e. of ply* $\gamma = \gamma(n) = \mathrm{o}(\log n)$.

Proof: Consider any $\epsilon > 0$. Choose k as the largest odd integer such that $(64k-60)\gamma/\pi \leq \log_3 n$. If $k < 7$, output $V(G)$. Otherwise, apply the algorithm of Lemma 8.4.10 and compute $C_{\min}$ in $\mathrm{O}(n^4 \log^3 n)$ time. Furthermore, if $\gamma = \gamma(n) = \mathrm{o}(\log n)$, there is a c_ϵ such that $k \geq 24/\epsilon$ and $k \geq 7$ for all $n \geq c_\epsilon$. Therefore, if $n \geq c_\epsilon$, it follows from Lemma 8.4.17 and the choice of k that $C_{\min}$ is a $(1+\epsilon)$-approximation to the optimum. Hence there is a fiptas$^\omega$ for Minimum $\preceq_{\mathsf{Leb}}$-Dominating Set on n-vertex disk graphs of bounded ply, i.e. of ply $\gamma = \gamma(n) = \mathrm{o}(\log n)$. The theorem follows from Theorem 2.2.4. □

Observe that the above theorem extends to intersection graphs of fat objects of any constant dimension and the weighted case. Because the $\preceq_{\mathsf{Leb}}$-factor is at most 6 for disk graphs, we also obtain the following result.

Theorem 8.4.19 *There is an algorithm that gives for any* $\epsilon > 0$ *a* $(6+\epsilon)$*-approximation for Minimum Dominating Set on disk graphs with* n *vertices and of bounded ply, i.e. of ply* $\gamma = \gamma(n) = \mathrm{o}(\log n)$*, in time* $f(1/\epsilon) \cdot n^{\mathrm{O}(1)}$ *for some computable function* f *of* $1/\epsilon$.

Similar constant-factor approximation algorithms exist for Minimum Dominating Set on intersection graphs of other fat objects of bounded ply, constant dimension, and constant $\preceq_{\mathsf{Leb}}$-factor. For example, a $(5+\epsilon)$-approximation algorithm on square graphs or a $(13+\epsilon)$-approximation algorithm on 3-dimensional ball graphs follow from Theorem 8.4.18.

8.4.3 A Better Constant

Although the above approach yields a constant-factor approximation algorithm for Minimum Dominating Set on disk graphs of bounded ply, we can also approximate it directly, i.e. without using $\preceq_{\mathsf{Leb}}$-dominating sets. This gives an easier algorithm with a better approximation ratio. To this end, we apply the shifting technique in a novel fashion.

Let $k \geq 9$ be an odd multiple of 3, let $\mathcal{D}$ be partitioned into levels and the plane into j-squares. We prove the following auxiliary theorem.

Theorem 8.4.20 *Let $\mathcal{D}$ be a set of disks of ply γ, $k \geq 9$ an odd multiple of 3, and OPT a minimum dominating set. Then we can obtain in time $\mathrm{O}(k^2 n^2\, 3^{32k\gamma/\pi} 2^{16k\gamma/\pi} 4^{16(k+1)\gamma/\pi})$ a set $C \subseteq \mathcal{D}$ dominating $\bigcup_S \mathcal{D}^{\mathrm{i}(S)}_{=j(S)}$ such that $|C| \leq \sum_S \left| OPT^S_{=j(S)} \right|$, where the union and sum are over all squares S.*

The set C is computed by performing bottom-up dynamic programming on the j-squares. For each j-square S, we consider each possible dominating set for $\mathcal{D}^{\mathrm{i}(S)}$, given the status of disks in $\mathcal{D}^{\mathrm{b}(S)}_{>j}$. A disk in $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ can have one of three statuses: either it is a dominator, or it is dominated by a disk in $\mathcal{D}^S$, or it is dominated by a yet undetermined disk. Now define for each j-square S and any two disjoint sets $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$ the function $\mathsf{size}(S, W_1, W_2)$ as

$$\begin{cases} \min\Big\{ \ |T| \ \Big| \ T \subseteq \mathcal{D}^{\mathrm{b}(S)}_{=j} \cup \mathcal{D}^{\mathrm{i}(S)}_{\geq j};\ W_1 \cup T \text{ dominates } \mathcal{D}^{\mathrm{i}(S)}_{\geq j} \cup W_2 \Big\} & \text{if } j = 0; \\ \min\Big\{ |U| + \displaystyle\sum_{i=1}^{4} \mathsf{size}\big(S_i, (W_1 \cup U)^{\mathrm{b}(S_i)}, X_i\big) \ \Big| & \\ \qquad U \subseteq \mathcal{D}^{+(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j},\ X_i \subseteq \mathcal{D}^{\mathrm{b}(S_i)}_{>j-1} & \\ \qquad \{X_1, \ldots, X_4\} \text{ decomposes } W_2 \cup (\mathcal{D}^{+(S)}_{>j-1} - U) \Big\} & \text{if } j > 0. \end{cases}$$

Here we define the minimum over an empty set to be ∞ and we say a family of pairwise disjoint sets $\{A_1, \ldots, A_m\}$ *decomposes* (or is a *decomposition* of) some set $\mathcal{A}$ if $A_i \subseteq \mathcal{A}$ for each i and $\bigcup_i A_i = \mathcal{A}$. Note that this definition explicitly allows empty sets.

Let $\mathsf{sol}(S, W_1, W_2)$ be the subfamily of $\mathcal{D}$ attaining $\mathsf{size}(S, W_1, W_2)$, or $\emptyset$ if $\mathsf{size}(S, W_1, W_2)$ is ∞. In the function parameters, the set W_1 is used for disks of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ that will be in the dominating set, while W_2 is used to denote the subset of $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ that should be dominated by a disk in D^S. Note that one actually only needs to consider sets $W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j} - N[W_1]$, but doing so would not improve the theoretical performance of the algorithm and might complicate its analysis.

Properties of the size- and sol-Functions

We start again by showing that size and sol are the functions that we need.

Lemma 8.4.21 $\sum_{S;\, j(S)=l} \mathsf{size}(S,\emptyset,\emptyset) \le \sum_S \left|\mathcal{C}^S_{=j(S)}\right|$, *where $\mathcal{C}$ is any dominating set.*

Proof: We prove using induction that the following inequality holds for all j-squares S:

$$\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}, N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j}\right) \le \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \sum_{S' \subseteq S} \left|\mathcal{C}^{S'}_{=j(S')}\right|.$$

Here $N^{\mathrm{ii}(S)}(X)$ is the set of disks $d \notin X$ such that d intersects some $d' \in X$ (i.e. $d \in N(X)$) and $d \cap d'$ intersects S.

The base case is trivial, since $\mathcal{C}^{\mathrm{b}(S)}_{>0} \cup \mathcal{C}^{\mathrm{i}(S)}_{>0} \cup \mathcal{C}^S_{=0} = \mathcal{C}^S$ clearly dominates $\mathcal{D}^{\mathrm{i}(S)}_{\ge j} \cup N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j}$. For the inductive step, we can show that

$$\begin{aligned}
&\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}, N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j}\right)\\
&\quad\le \left|\mathcal{C}^{+(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \mathcal{C}^{\mathrm{b}(S_i)}_{>j-1}, N^{\mathrm{ii}(S_i)}(\mathcal{C}^{S_i})^{\mathrm{b}(S_i)}_{>j-1}\right)\\
&\quad\le \left|\mathcal{C}^{+(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \left|\mathcal{C}^{\mathrm{i}(S_i)}_{>j-1}\right| + \sum_{i=1}^{4} \sum_{S'_i \subseteq S_i} \left|\mathcal{C}^{S'_i}_{=j(S'_i)}\right|\\
&\quad= \left|\mathcal{C}^{+(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \left|\mathcal{C}^{\mathrm{i}(S)}_{>j-1}\right| - \left|\mathcal{C}^{+(S)}_{>j-1}\right| + \sum_{i=1}^{4} \sum_{S'_i \subseteq S_i} \left|\mathcal{C}^{S'_i}_{=j(S'_i)}\right|\\
&\quad= \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \left|\mathcal{C}^S_{=j}\right| + \sum_{i=1}^{4} \sum_{S'_i \subseteq S_i} \left|\mathcal{C}^{S'_i}_{=j(S'_i)}\right|\\
&\quad= \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \sum_{S' \subseteq S} \left|\mathcal{C}^{S'}_{=j(S')}\right|.
\end{aligned}$$

The first inequality above is the crucial one and that it should hold is not obvious. We give an explicit proof.

Suppose that to obtain

$$\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}, N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j}\right)$$

using the definition of size, we consider $U = \mathcal{C}^{+(S)}_{>j-1} \cup \mathcal{C}^{\mathrm{b}(S)}_{=j}$. As for $i = 1, \ldots, 4$,

$$\begin{aligned}
(W_1 \cup U)^{\mathrm{b}(S_i)} &= \left(\mathcal{C}^{\mathrm{b}(S)}_{>j} \cup \mathcal{C}^{+(S)}_{>j-1} \cup \mathcal{C}^{\mathrm{b}(S)}_{=j}\right)^{\mathrm{b}(S_i)}\\
&= \left(\mathcal{C}^{\mathrm{b}(S)}_{>j-1} \cup \mathcal{C}^{+(S)}_{>j-1}\right)^{\mathrm{b}(S_i)}\\
&\overset{\text{by def.}}{=} \left(\bigcup_{m=1}^{4} \mathcal{C}^{\mathrm{b}(S_m)}_{>j-1}\right)^{\mathrm{b}(S_i)}\\
&= \mathcal{C}^{\mathrm{b}(S_i)}_{>j-1},
\end{aligned}$$

the second parameter of the inductive call is correct.

So what about the third parameter? We claim that

$$W_2 \cup (\mathcal{D}^{+(S)}_{>j-1} - U) \subseteq \bigcup_{i=1}^{4} N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S_i)}_{>j-1} \subseteq \bigcup_{i=1}^{4} N^{\mathrm{ii}(S_i)}(\mathcal{C}^{S_i})^{\mathrm{b}(S_i)}_{>j-1}.$$

Because $\mathcal{C}$ is a dominating set and every disk in $\mathcal{D}^{+(S)}_{>j-1}$ must be dominated by a disk in $\mathcal{D}^S$,

$$\begin{aligned} N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{+(S)}_{>j-1} &= N(\mathcal{C}^S)^{+(S)}_{>j-1} \\ &= \mathcal{D}^{+(S)}_{>j-1} - \mathcal{C}^{+(S)}_{>j-1} \\ &= \mathcal{D}^{+(S)}_{>j-1} - U. \end{aligned}$$

Then

$$\begin{aligned} W_2 \cup (\mathcal{D}^{+(S)}_{>j-1} - U) &= N^{ii(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j} \cup N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{+(S)}_{>j-1} \\ &\subseteq N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j-1} \cup N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{+(S)}_{>j-1} \\ &\overset{\text{by def.}}{=} \bigcup_{i=1}^{4} N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S_i)}_{>j-1}, \end{aligned}$$

thus proving the first part of the claim.

To prove the second part, consider any disk d in $N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S_i)}_{>j-1}$ for all $i \in I \subseteq \{1, \ldots, 4\}$. As $d \in N^{\mathrm{ii}(S)}(\mathcal{C}^S)$, there must be some $h \in \{1, \ldots, 4\}$ such that $d \in N^{\mathrm{ii}(S_h)}(\mathcal{C}^S)$, i.e. $d \in N^{\mathrm{ii}(S_h)}(\mathcal{C}^{S_h})$. Furthermore, it is clear that $h \in I$. But then $d \in N^{\mathrm{ii}(S_h)}(\mathcal{C}^{S_h})^{\mathrm{b}(S_h)}_{>j-1}$. This proves the claim.

Following the claim, there exists a decomposition $\{X_1, \ldots, X_4\}$ of $W_2 \cup (\mathcal{D}^{+(S)}_{>j-1} - U)$ such that $X_i \subseteq N^{\mathrm{ii}(S_i)}(\mathcal{C}^{S_i})^{\mathrm{b}(S_i)}_{>j-1}$ and hence $X_i \subseteq \mathcal{D}^{\mathrm{b}(S_i)}_{>j-1}$. Therefore for such X_i and the chosen set U,

$$\begin{aligned} &\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}, N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j}\right) \\ &\quad\leq |U| + \sum_{i=1}^{4} \mathsf{size}\Big(S_i, (W_1 \cup U)^{\mathrm{b}(S_i)}, X_i\Big) \\ &\quad= \left|\mathcal{C}^{+(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \mathcal{C}^{\mathrm{b}(S_i)}_{>j-1}, X_i\right) \\ &\quad\leq \left|\mathcal{C}^{+(S)}_{>j-1}\right| + \left|\mathcal{C}^{\mathrm{b}(S)}_{=j}\right| + \sum_{i=1}^{4} \mathsf{size}\left(S_i, \mathcal{C}^{\mathrm{b}(S_i)}_{>j-1}, N^{\mathrm{ii}(S_i)}(\mathcal{C}^{S_i})^{\mathrm{b}(S_i)}_{>j-1}\right), \end{aligned}$$

where the last inequality follows from $X_i \subseteq N^{\mathrm{ii}(S_i)}(\mathcal{C}^{S_i})^{\mathrm{b}(S_i)}_{>j-1}$. This proves the inequality of the previous page.

We now know that

$$\mathsf{size}\left(S, \mathcal{C}^{\mathrm{b}(S)}_{>j}, N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j}\right) \leq \left|\mathcal{C}^{\mathrm{i}(S)}_{>j}\right| + \sum_{S' \subseteq S} \left|\mathcal{C}^{S'}_{=j(S')}\right|.$$

Since l is the level of the largest disk, $\mathcal{C}^{\mathrm{i}(S)}_{>j} = \emptyset$ and $\mathcal{C}^{\mathrm{b}(S)}_{>j} = \emptyset$, and thus $N^{\mathrm{ii}(S)}(\mathcal{C}^S)^{\mathrm{b}(S)}_{>j} = \emptyset$ for all j-squares S with $j \geq l$. Hence

$$\sum_{S;\, j(S)=l} \mathsf{size}(S,\emptyset,\emptyset) \leq \sum_{S;\, j(S)=l} \sum_{S' \subseteq S} \left| \mathcal{C}^{S'}_{=j(S')} \right| = \sum_S \left| \mathcal{C}^S_{=j(S)} \right|.$$

This proves the lemma. $\square$

It follows that if OPT is a minimum dominating set, then

$$\sum_{S;\, j(S)=l} \mathsf{size}(S,\emptyset,\emptyset) \leq \sum_S \left| OPT^S_{=j(S)} \right|.$$

Lemma 8.4.22 $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S,\emptyset,\emptyset)$ *is a dominating set for* $\bigcup_S \mathcal{D}^{\mathrm{i}(S)}_{=j(S)}$.

Proof: For any j-square S and any two disjoint $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$, we claim that $W_1 \cup \mathsf{sol}(S, W_1, W_2)$ dominates $W_2 \cup \mathcal{D}^{\mathrm{i}(S)}_{>j} \cup \bigcup_{S' \subseteq S} \mathcal{D}^{\mathrm{i}(S')}_{=j(S')}$ if $\mathsf{size}(S, W_1, W_2) \neq \infty$. Apply induction on j. The case $j = 0$ is trivial, so assume that $j > 0$ and that the claim holds for all $j' < j$.

Suppose that $\mathsf{size}(S, W_1, W_2) \neq \infty$ for some j-square S and for disjoint $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$. Let $U, X_1, \ldots, X_4$ attain the minimum in the definition of size. Then $\mathsf{size}(S_i, (W_1 \cup U)^{\mathrm{b}(S_i)}, X_i) \neq \infty$ for $i = 1, \ldots, 4$. By induction, $(W_1 \cup U)^{\mathrm{b}(S_i)} \cup \mathsf{sol}(S_i, (W_1 \cup U)^{\mathrm{b}(S_i)}, X_i)$ dominates $X_i \cup \mathcal{D}^{\mathrm{i}(S_i)}_{>j-1} \cup \bigcup_{S'_i \subseteq S_i} \mathcal{D}^{\mathrm{i}(S'_i)}_{=j(S'_i)}$ for $i = 1, \ldots, 4$. Observe that

$$\begin{aligned}
&\bigcup_{i=1}^{4} \left((W_1 \cup U)^{\mathrm{b}(S_i)} \cup \mathsf{sol}\left(S_i, (W_1 \cup U)^{\mathrm{b}(S_i)}, X_i \right) \right) \\
&\qquad = W_1 \cup U \cup \bigcup_{i=1}^{4} \mathsf{sol}\left(S_i, (W_1 \cup U)^{\mathrm{b}(S_i)}, X_i \right) \\
&\qquad = W_1 \cup \mathsf{sol}(S, W_1, W_2)
\end{aligned}$$

and that

$$\begin{aligned}
&\bigcup_{i=1}^{4} \left(X_i \cup \mathcal{D}^{\mathrm{i}(S_i)}_{>j-1} \cup \bigcup_{S'_i \subseteq S_i} \mathcal{D}^{\mathrm{i}(S'_i)}_{=j(S'_i)} \right) \\
&\qquad = W_2 \cup \left(\mathcal{D}^{+(S)}_{>j-1} - U \right) \cup \bigcup_{i=1}^{4} \left(\mathcal{D}^{\mathrm{i}(S_i)}_{>j-1} \cup \bigcup_{S'_i \subseteq S_i} \mathcal{D}^{\mathrm{i}(S'_i)}_{=j(S'_i)} \right).
\end{aligned}$$

Because $W_1 \cup \mathsf{sol}(S, W_1, W_2)$ dominates U, $W_1 \cup \mathsf{sol}(S, W_1, W_2)$ dominates

$$W_2 \cup \mathcal{D}^{+(S)}_{>j-1} \cup \bigcup_{i=1}^{4} \left(\mathcal{D}^{\mathrm{i}(S_i)}_{>j-1} \cup \bigcup_{S'_i \subseteq S_i} \mathcal{D}^{\mathrm{i}(S'_i)}_{=j(S'_i)} \right) = W_2 \cup \mathcal{D}^{\mathrm{i}(S)}_{>j} \cup \bigcup_{S' \subseteq S} \mathcal{D}^{\mathrm{i}(S')}_{=j(S')}.$$

This proves the claim.

We know that $\sum_{S;\,j(S)=l} \mathsf{size}(S,\emptyset,\emptyset) \neq \infty$ from Lemma 8.4.21. Since each disk has level at most l, $\bigcup_{S;\,j(S)=l} \mathsf{sol}(S,\emptyset,\emptyset)$ dominates $\bigcup_S \mathcal{D}^{\mathrm{i}(S)}_{=j(S)}$. □

Computing the size- and sol-Functions

We use the same definitions of (non)empty and relevant (child) as in Chapter 7. That is, a j-square is *nonempty* if it is intersected by a level j disk and *empty* otherwise. A j-square S is *relevant* if one of its three siblings is nonempty or there is a nonempty square S' containing S, such that S' has level at most $j + \lceil \log k \rceil$ (so each nonempty j-square is relevant). Note that this definition induces $\mathrm{O}(k^2 n)$ relevant squares. A relevant square S is said to be a *relevant child* of another relevant square S' if $S \subset S'$ and there is no third relevant square S'', such that $S \subset S'' \subset S'$. Conversely, if S is a relevant child of S', then S' is a *relevant parent* of S.

Lemma 8.4.23 *For each relevant* 0*-square* S*, all* size- *and* sol-*values for* S *can be computed in* $\mathrm{O}\big(nk\gamma\, 3^{(24k+32)\gamma/\pi}\big)$ *time.*

Proof: We use the bounds of Lemma 7.1.1 and Lemma 7.1.2. Then $\left|\mathcal{D}^{\mathrm{b}(S)}_{>0}\right| \leq 16k\gamma/\pi$ and all disjoint $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>0}$ can be enumerated in $\mathrm{O}(3^{16k\gamma/\pi})$ time. Furthermore, as the pathwidth of $\mathcal{D}^{\mathrm{b}(S)}_{=j} \cup \mathcal{D}^{\mathrm{i}(S)}_{>j}$ can be bounded by $8(k+4)\gamma/\pi$, we can adapt the algorithm of Corollary 5.3.9 to find the appropriate minimum dominating set in $\mathrm{O}(nk\gamma\, 3^{(8k+32)\gamma/\pi})$ time. The lemma follows. □

Now assume that the size- and sol-values of all relevant children of a j-square S are known.

Lemma 8.4.24 *For each relevant* j*-square* S $(j > 0)$ *with relevant* $(j-1)$*-square children, in* $\mathrm{O}(3^{32k\gamma/\pi} 2^{16k\gamma/\pi} 4^{16(k+1)\gamma/\pi})$ *time all* size- *and* sol-*values for* S *can be computed.*

Proof: Using Lemma 7.1.1, we can show that $\left|\mathcal{D}^{\mathrm{b}(S)}_{\geq j}\right| \leq 32k\gamma/\pi$ and $\left|\mathcal{D}^{+(S)}_{\geq j}\right| \leq 16k\gamma/\pi$. Hence all disjoint $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$ and all $U \subseteq \mathcal{D}^{+(S)}_{>j-1} \cup \mathcal{D}^{\mathrm{b}(S)}_{=j}$ can be enumerated in $\mathrm{O}(3^{32k\gamma/\pi} 2^{16k\gamma/\pi})$ time. To enumerate all decompositions $\{X_1, \ldots, X_4\}$ of $W_2 \cup (\mathcal{D}^{+(S)}_{>j-1} - U)$ for fixed W_2 and U such that $X_i \subseteq \mathcal{D}^{\mathrm{b}(S_i)}_{>j-1}$, it suffices to consider decompositions of disks in an 'extended cross'. Following Lemma 7.1.1, the number of disks intersecting it is at most $16(k+1)\gamma/\pi$. The lemma follows. □

Lemma 8.4.25 *For each relevant* j*-square* S $(j > 0)$ *with no relevant children of level* $j-1$ *all* size- *and* sol-*values for* S *can be computed in* $\mathrm{O}(n\, 2^{64\gamma/\pi} 3^{32\gamma/\pi})$ *time.*

Proof: From the proof of Lemma 7.2.6, we know that $\mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}$, $\mathcal{D}^{+(S)}_{>j-1} = \emptyset$, and $\mathcal{D}^{\mathrm{b}(S)}_{=j} = \emptyset$. Then for any disjoint $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$, we can show that $\mathsf{size}(S, W_1, W_2)$ equals

$$\begin{cases} 0 & \text{if } S \text{ has no relevant children and} \\ & \quad W_1 \text{ dominates } W_2; \\ \infty & \text{if } S \text{ has no relevant children and} \\ & \quad W_1 \text{ doesn't dominate } W_2; \\ \min\limits_{\{X_{S''}\}} \sum\limits_{S''} \mathsf{size}\left(S'', W_1^{\mathrm{b}(S'')}, X_{S''}\right) & \text{otherwise.} \end{cases}$$

The sum is over all relevant children S'' of S and $\{X_{S''}\}$ decomposes W_2.

Since $\mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}$, $\left|\mathcal{D}^{\mathrm{b}(S)}_{>j}\right| \leq 32\gamma/\pi$. Then all disjoint $W_1, W_2 \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$ can be enumerated in $\mathrm{O}(3^{32\gamma/\pi})$ time.

For fixed W_1 and W_2, we can compute $\mathsf{size}(S, W_1, W_2)$ in $\mathrm{O}(n)$ time if S has no relevant children. Otherwise, number the relevant children of S arbitrarily, $S''_1, \ldots, S''_m$. Now compute $\mathsf{size}(S, W_1, W_2)$ using the following function s. For any $X \subseteq W_2$,

$$\begin{aligned} s^1(X) &= \mathsf{size}(S''_1, W_1^{\mathrm{b}(S'')}, X) \\ s^i(X) &= \min\nolimits_{X_{S''_i} \subseteq X} \{\mathsf{size}(S''_1, W_1^{\mathrm{b}(S'')}, X) + s^{i-1}(X - X_{S''_i})\} \end{aligned}$$

Then $\mathsf{size}(S, W_1, W_2) = s^m(W_2)$. One can thus compute $\mathsf{size}(S, W_1, W_2)$ in $\mathrm{O}(n\, 2^{64\gamma/\pi})$ time, as $m = \mathrm{O}(n)$. The lemma follows. □

Lemma 8.4.26 *The value of* $\sum_{S;\, j(S)=l} \mathsf{size}(S, \emptyset, \emptyset)$ *can be computed in time* $\mathrm{O}(k^2 n^2 \gamma\, 3^{32k\gamma/\pi} 2^{16k\gamma/\pi} 4^{16(k+1)\gamma/\pi})$.

Proof: Follows from Lemma 8.4.23, Lemma 8.4.24, Lemma 8.4.25, and the proof of Lemma 7.2.7. □

Proof of Theorem 8.4.20: Follows directly from Lemmas 8.4.21, 8.4.22, and 8.4.26. □

The Approximation Algorithm

The shifting technique can now be applied as follows. For an integer a ($0 \leq a \leq k-1$), call a line of level j *active* if it has the form $y = (hk + a2^{l-j})2^j$ or $x = (vk + a2^{l-j})2^j$ ($h, v \in \mathbb{Z}$). The active lines partition the plane into j-squares as before, although shifted with respect to the shifting parameter a. However, we can still apply the algorithm of Theorem 8.4.20.

Let C_a denote the set returned by the algorithm for the j-squares induced by shifting parameter a ($0 \leq a \leq k-1$). Instead of considering each set C_a individually, we join three such sets to ensure that we have a dominating set.

So let $C_i^3 = C_i \cup C_{i+k/3} \cup C_{i+2k/3}$ for each $i = 0, \ldots, k/3-1$. This is properly defined, as k is a multiple of 3. Denote the smallest such set by $C_{\min}^3$. We claim that $C_{\min}^3$ is a dominating set of cardinality at most $(3 + 36/k) \cdot |OPT|$, where OPT is a minimum dominating set.

To prove this claim, let $\mathcal{D}_a^b$ be the set of disks intersecting the boundary of a j-square S at their level, i.e. $\mathcal{D}_a^b = \bigcup_S \mathcal{D}_{=j(S)}^{\mathrm{b}(S)}$.

Lemma 8.4.27 *C_i^3 is a dominating set of $\mathcal{D}$ for any $i = 0, \ldots, k/3-1$.*

Proof: We claim that any disk is in at most two of the sets $\mathcal{D}_i^b, \mathcal{D}_{i+k/3}^b, \mathcal{D}_{i+2k/3}^b$. A level j disk is in $\mathcal{D}_a^b$ if and only if it intersects an active line of level j for a. We showed in Lemma 7.2.8 that any disk intersects an active horizontal line for at most two values of a and an active vertical line for at most two values of a. It is easy to see from the proof of this lemma that the intersections with an active horizontal line and similarly the intersections with an active vertical line must occur for consecutive values of a (modulo k). Since $k \geq 9$ is an odd multiple of 3, $k/3 > 1$, and thus $i, i + k/3, i + 2k/3$ are nonconsecutive integers (modulo k). It follows that any disk is in at most two of the sets $\mathcal{D}_i^b, \mathcal{D}_{i+k/3}^b, \mathcal{D}_{i+2k/3}^b$, as claimed.

Lemma 8.4.22 shows that C_a is a dominating set of $\bigcup_S \mathcal{D}_{=j(S)}^{\mathrm{i}(S)} = \mathcal{D} - \mathcal{D}_a^b$. Given the previous argument, $(\mathcal{D} - \mathcal{D}_i^b) \cup (\mathcal{D} - \mathcal{D}_{i+k/3}^b) \cup (\mathcal{D} - \mathcal{D}_{i+2k/3}^b) = \mathcal{D}$. Hence $C_i^3 = C_i \cup C_{i+k/3} \cup C_{i+2k/3}$ is a dominating set of $\mathcal{D}$. □

Lemma 8.4.28 *$|C_{\min}^3| \leq (3+36/k) \cdot |OPT|$, where OPT is a minimum dominating set.*

Proof: Following the proof of Lemma 7.2.8, a level j disk is in $\mathcal{D}_a^b$ for at most 4 different values of a. Therefore $\sum_{a=0}^{k-1} |OPT \cap \mathcal{D}_a^b| \leq 4 \cdot |OPT|$. Furthermore, for any fixed value of a, any level j disk can intersect at most 4 j-squares. It follows from Theorem 8.4.20 that

$$\begin{aligned} |C_a| &\leq \sum_S \left| OPT_{=j(S)}^S \right| \\ &\leq \sum_S \left(\left| OPT_{=j(S)}^{\mathrm{i}(S)} \right| + \left| OPT_{=j(S)}^{\mathrm{b}(S)} \right| \right) \\ &\leq |OPT| + 3|OPT \cap \mathcal{D}_a^b|. \end{aligned}$$

Then

$$\begin{aligned} \tfrac{1}{3}k \cdot |\mathcal{C}_{\min}^3| &\leq \sum_{i=0}^{k/3-1} |C_i^3| \\ &\leq \sum_{a=0}^{k-1} |C_a| \end{aligned}$$

$$\begin{aligned} &\leq \sum_{a=0}^{k-1} \left(|OPT| + 3|OPT \cap \mathcal{D}_a^b|\right) \\ &\leq (k+12) \cdot |OPT|. \end{aligned}$$

Hence $|C^3_{\min}| \leq (3 + 36/k) \cdot |OPT|$. □

Combining Theorem 8.4.20 and Lemma 8.4.28, we obtain the following approximation algorithm.

Theorem 8.4.29 *There is an algorithm that gives for any $\epsilon > 0$ a $(3+\epsilon)$-approximation for Minimum Dominating Set on disk graphs with n vertices and of bounded ply, i.e. of ply $\gamma = \gamma(n) = \mathrm{o}(\log n)$, in time $f(1/\epsilon) \cdot n^{\mathrm{O}(1)}$ for some computable function f of $1/\epsilon$.*

Proof: Consider any $\epsilon > 0$. Choose k as the largest odd multiple of 3 such that $32k\gamma/\pi \leq \log_3 n$. If $k < 9$, output $V(G)$. Otherwise, apply the algorithm of Theorem 8.4.20 and compute $C^3_{\min}$ in $\mathrm{O}(n^5 \log^3 n)$ time. Furthermore, if $\gamma = \gamma(n) = \mathrm{o}(\log n)$, there is a c_ϵ such that $k \geq 36/\epsilon$ and $k \geq 9$ for all $n \geq c_\epsilon$. Therefore, if $n \geq c_\epsilon$, it follows from Lemma 8.4.28 and the choice of k that $C^3_{\min}$ is a $(3+\epsilon)$-approximation to the optimum. The theorem now follows from the proof of Theorem 2.2.4. □

We can obtain analogous approximation algorithms on intersection graphs of fat objects of bounded ply and of any constant dimension.

8.5 Hardness of Approximation

We have seen that although Minimum Dominating Set is a challenging problem on intersection graphs of arbitrary-sized geometric objects, it still is approximable on a variety of classes of geometric intersection graphs. We show however that there are also classes of geometric intersection graphs for which no constant-factor approximation algorithm or approximation scheme can exist, under certain complexity assumptions.

We prove that Minimum Dominating Set on intersection graphs of convex polygons or of homothetic polygons is as hard as on general graphs. That is, it is not approximable within $(1-\epsilon)\ln n$ for any $\epsilon > 0$, unless $\mathrm{NP} \subset \mathrm{DTIME}(n^{\mathrm{O}(\log\log n)})$. This nicely complements Theorem 8.3.4, where we gave a constant-factor approximation algorithm on intersection graphs of homothetic convex polygons. Hence it seems that both convexity and homotheticity are essential properties of the objects when designing a constant-factor approximation algorithm.

We also gain further insight into Minimum Dominating Set on disk graphs. We show that on a collection of fat almost-disks, Minimum Dominating Set is as hard as on general graphs. Even if the fat objects have constant description complexity, the problem is still APX-hard.

Finally, we solve an open problem of Chlebík and Chlebíková [65] by proving that Minimum Dominating Set is APX-hard on rectangle intersection graphs. This result extends to intersection graphs of ellipses. We should note that all hardness results given here extend to Minimum Connected Dominating Set.

8.5.1 Intersection Graphs of Polygons

We consider the approximability of Minimum Dominating Set on intersection graphs of polygons. First we show that convexity of the objects is no guarantee for the existence of a constant-factor approximation algorithm. Instead of just looking at arbitrary convex polygons, we prove a stronger result.

Recall from Chapter 3 that a *polygon-circle graph* is the intersection graph of a set of polygons for which all corners lie on a fixed circle. Note that in a polygon-circle graph, all polygons are convex. This graph class is a generalization of *circle graphs*, which are intersection graphs of chords of a fixed circle. On circle graphs, we know that Minimum Dominating Set has a $(2+\epsilon)$-approximation algorithm [76], but no ptas unless P=NP [75]. The slight generalization to polygon-circle graphs however makes Minimum Dominating Set much more difficult.

Theorem 8.5.1 *Minimum Dominating Set on polygon-circle graphs is not approximable within* $(1-\epsilon)\ln n$ *for any* $\epsilon > 0$, *unless NP* $\subset$ *DTIME*($n^{\mathrm{O}(\log\log n)}$).

Proof: We give a gap-preserving reduction from Minimum Set Cover. Consider an instance $(\mathbb{U}, \mathcal{S})$ of Minimum Set Cover and assume that $\mathbb{U} = \bigcup \mathcal{S}$ and $\mathbb{U} = \{1, \ldots, n\}$. Fix a circle and $n+1$ points on this circle, numbered $p_1, \ldots, p_{n+1}$ in order of appearance on the circle. Construct a polygon P_j for each set $\mathcal{S}_j$ as the convex hull of the set $\{p_i \mid i \in \mathcal{S}_j\} \cup \{p_{n+1}\}$. Furthermore, place a tiny polygon around each point p_i such that these tiny polygons are pairwise disjoint.

Observe that any polygon dominated by a tiny polygon is also dominated by some polygon P_j. It is now easy to see that the optima of the minimum set cover instance and the constructed instance of Minimum Dominating Set on polygon-circle graphs are the same. As the construction can be computed in polynomial time, this gives a gap-preserving reduction. The theorem then follows from Feige's inapproximability result for Minimum Set Cover [108]. □

A direct consequence of Theorem 8.5.1 is an inapproximability result on intersection graphs of convex polygons.

Corollary 8.5.2 *Unless NP* $\subset$ *DTIME*($n^{\mathrm{O}(\log\log n)}$), *Minimum Dominating Set on intersection graphs of convex polygons cannot be approximated within* $(1-\epsilon)\ln n$ *for any* $\epsilon > 0$.

We give a similar result for intersection graphs of fat convex polygons later.

If the polygons are not convex, but translated copies of a fixed polygon, the approximability of Minimum Dominating Set does not change.

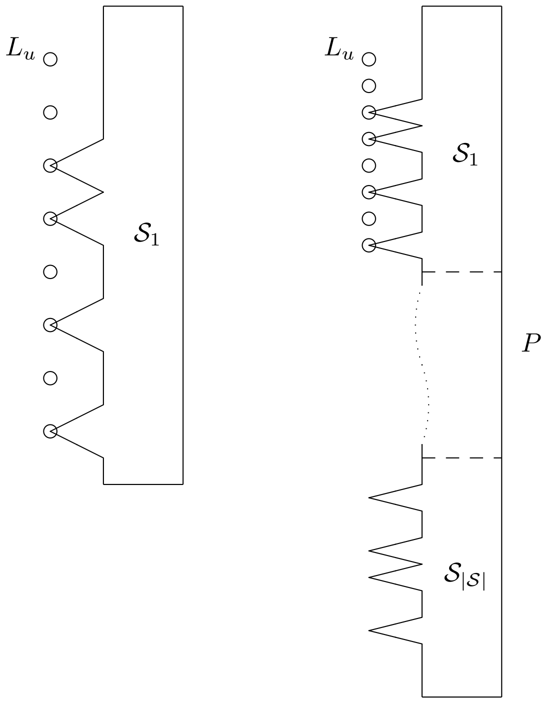

Figure 8.2: The left figure shows the rectangle constructed for $\mathcal{S}_1$. The right figure shows the combination of the rectangles for $\mathcal{S}_1, \ldots, \mathcal{S}_{|\mathcal{S}|}$. This is the base polygon P. The small circles represent the L_u, which are homothetic copies of P.

Theorem 8.5.3 *Minimum Dominating Set on intersection graphs of homothetic polygons is not approximable within* $(1-\epsilon)\ln n$ *for any* $\epsilon > 0$*, unless NP* $\subset$ *DTIME(*$n^{O(\log\log n)}$*).*

Proof: We give a similar reduction as in the proof of Theorem 8.5.1. Given an instance $(\mathbb{U}, \mathcal{S})$ of Minimum Set Cover, place for each $u \in \mathbb{U}$ a polygon L_u in the plane, such that these polygons are aligned in a column (see Figure 8.2). Now construct a rectangle next to this column. Deform the long side by placing small bulges on it such that the deformed rectangle intersects an L_u if and only if $u \in \mathcal{S}_1$ (see Figure 8.2). Do this for each set in $\mathcal{S}$ and stack these rectangles. This is the base polygon P. By taking a translated copy of P for each set in $\mathcal{S}$ and ensuring that the L_u are homothetic copies of P, we can build the same graph as in Theorem 8.5.1. The theorem follows. □

The hardness results of Corollary 8.5.2 and Theorem 8.5.3 complement Theorem 8.3.4, where we gave an $O(r^4)$-approximation algorithm on intersection graphs of homothetic convex polygons with r corners.

The approximability of Minimum Dominating Set on intersection graphs of convex polygons or of homothetic polygons with r corners has yet to be determined. The APX-hardness on circle graphs [75] implies APX-hardness on intersection graphs of convex polygons with two (or more) corners. Hence no ptas exists, unless P=NP [16]. Using the gadget of Theorem 8.5.1, we can give a slightly weaker result, but by an easier proof. We use that Minimum k-Set Cover is APX-hard for any $k \geq 3$ (by reduction from Minimum Vertex Cover on graphs of degree at most 3 [9]). *Minimum k-Set Cover* is the variation of Minimum Set Cover where each set has cardinality at most k.

Theorem 8.5.4 *Minimum Dominating Set on polygon-circle graphs of convex polygons with* r *corners is APX-hard for any* $r \geq 4$*. Hence it has no ptas, unless P=NP.*

Proof: We use the same gadget as in the proof of Theorem 8.5.1 and reduce from Minimum k-Set Cover, which is APX-hard for any $k \geq 3$. The gadget constructs polygons with at most $k+1$ corners. The theorem follows. □

Corollary 8.5.5 *Minimum Dominating Set on intersection graphs of convex polygons with* r *corners is APX-hard for any* $r \geq 4$*. Hence it has no ptas, unless P=NP.*

8.5.2 Intersection Graphs of Fat Objects

The approximation schemes for Maximum Independent Set and Minimum Vertex Cover on disk graphs (see Chapter 7) extend easily to intersection graphs of fat objects. It is unlikely that an approximation algorithm for Minimum Dominating Set extends this way, as on intersection graphs of fat objects that are almost-disks, Minimum Dominating Set becomes hard to approximate.

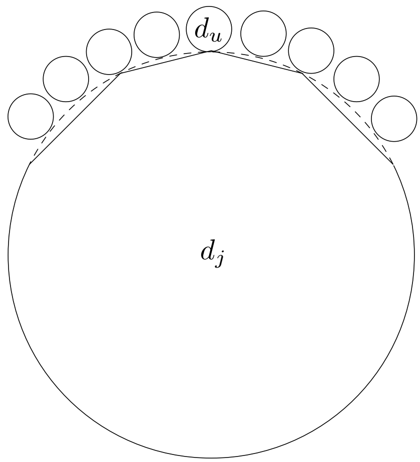

Figure 8.3: A cut-off disk d_j and the disks d_u for elements $u \in \mathbb{U}$ of Theorem 8.5.6.

Recall that a convex subset s of $\mathbb{R}^2$ is *α-fat* for some $\alpha \geq 1$ if the ratio between the radii of the smallest disk enclosing s and the largest disk inscribed in s is at most α [97].

Theorem 8.5.6 *For any $\alpha > 1$, Minimum Dominating Set on intersection graphs of α-fat objects is not approximable within $(1-\epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{O(\log\log n)}$).*

Proof: We reduce from Minimum Set Cover in a manner similar as in the proof of Theorem 8.5.1. For an instance $(\mathbb{U}, \mathcal{S})$ of Minimum Set Cover, construct an instance of Minimum Dominating Set as follows. Each $u \in \mathbb{U}$ corresponds to a 'small' disk d_u. Each $\mathcal{S}_j$ corresponds to a disk d_j with the top replaced by a polygonal structure such that d_j intersects d_u if and only if $u \in \mathcal{S}_j$ (see Figure 8.3). Packing the d_u close together makes the fatness of the construction arbitrarily close to 1. As any object dominated by a d_u is also dominated by a d_j for which $u \in S_j$, the optima of the two instances are equal. Moreover, the construction can be computed in polynomial time. The theorem then follows from Feige's result [108]. □

An object has *constant description complexity* if it is a semialgebraic set defined by a constant number of polynomial (in)equalities of constant maximum degree [97]. The objects in the proof of Theorem 8.5.6 that model the $\mathcal{S}_j$ are the intersection of a disk with a polygon and thus we can describe each such d_j by one quadratic inequality and $|\mathcal{S}_j| + 1$ linear inequalities. Hence the objects in the construction of Theorem 8.5.6 might not have constant description complexity. So for constant description complexity objects, better approximation ratios than $\ln n$ could be attained. However, we can still prove APX-hardness by reducing from Minimum k-Set Cover with the same gadget.

Theorem 8.5.7 *For any $\alpha > 1$, Minimum Dominating Set on intersection graphs of α-fat objects of constant description complexity is APX-hard. Hence it has no ptas, unless P=NP.*

These results say something about intersection graphs of fat objects in general, and of fat almost-disks in particular. But we can easily prove similar results for almost-squares, bounded aspect-ratio almost-rectangles, almost-triangles, etc. Basically, if we slightly relax the shape constraints for a given object, Minimum Dominating Set on the intersection graphs of such relaxed objects is hard to approximate. Moreover, the above results can be used to derive hardness of approximation results for Minimum Connected Dominating Set.

8.5.3 Intersection Graphs of Rectangles

Chlebík and Chlebíková [65] proved that Minimum Dominating Set is APX-hard on intersection graphs of three-dimensional axis-parallel boxes and asked whether this result can be extended to only two dimensions. We prove that this is indeed the case.

Theorem 8.5.8 *Minimum Dominating Set on rectangle intersection graphs is APX-hard. Hence it has no ptas, unless P=NP.*

Proof: We give an L-reduction [220] from Minimum Vertex Cover on graphs of degree three, which is known to be APX-hard [9]. Consider an arbitrary instance x of Minimum Vertex Cover on graphs of degree three. Let $G = (V, E)$ be the graph of x and denote the cardinality of the smallest vertex cover in G by k. Number the vertices of V arbitrarily $v_1, \ldots, v_n$, where $n = |V|$. Now construct for each vertex v_i a horizontal and a vertical rectangle R_i^h and R_i^v and connect them as shown in Figure 8.4. Call the big rectangle used in the connection of R_i^h and R_i^v the *big plate* P_i of i and the two small rectangles the *small plates* of i. This models the vertices. Next we model the edges. If $(v_i, v_j) \in E$ for certain $i < j$, then add a small rectangle $S_{i,j}$ in the intersection of rectangles R_i^v and R_j^h (see Figure 8.4). This gives the instance $f(x)$ of Minimum Dominating Set on rectangle intersection graphs. Observe that this is indeed a polynomial-time computable function (even if not only the graph, but also the rectangles are part of the output).

Let C be a vertex cover of G of cardinality k. Let $R^h[C] = \{R_i^h \mid v_i \in C\}$ be the set of horizontal rectangles induced by C and similarly let $R^v[C]$ be the set of vertical rectangles induced by C. Furthermore, let $\overline{P[C]} = \{P_i \mid v_i \notin C\}$ be the big plates for which the corresponding vertex is not in C.

We claim that $D = R^h[C] \cup R^v[C] \cup \overline{P[C]}$ is a dominating set of G'. Let r be an arbitrary rectangle. Suppose that r is an $S_{i,j}$ for a certain i, j. Since C is a vertex cover, $v_i \in C$ or $v_j \in C$. Assume w.l.o.g. that $v_i \in C$. Then

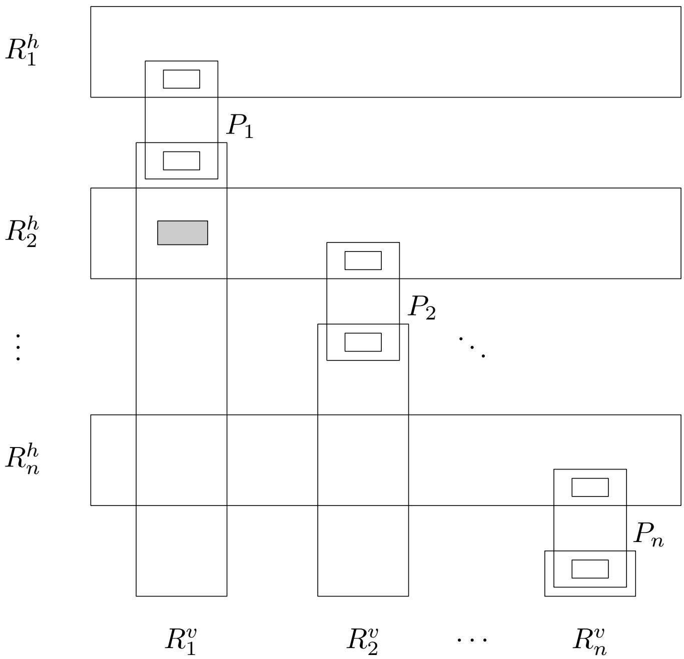

Figure 8.4: The intersection graph used in the proof of Theorem 8.5.8. If edge (v_1, v_2) is in E, then the shaded rectangle $S_{1,2}$ is in G'.

by construction, $R_i^h, R_i^v \in D$, and thus by construction of G', $S_{i,j}$ must be dominated. Suppose that r is a (big or small) plate of i. If $v_i \in C$, then $R_i^h, R_i^v \in D$, and thus the plate must be dominated. If $v_i \notin C$, then $P_i \in \overline{P[C]} \subseteq D$, and the plate is dominated. Using a similar argument, we can show that if r is R_i^h or R_i^v for certain i, it must be dominated. Hence D is a dominating set of G'.

Note that $|R^h[C]| = |R^v[C]| = |C| = k$. Furthermore, $|\overline{P[C]}| = n - k$. Since G has degree three, $k \geq n/4$. Hence

$$m^*(f(x)) \leq n - k + k + k \leq 4k + k = 5 \cdot m^*(x). \tag{8.1}$$

We now take a closer look at the cardinality of dominating sets of G'. Let D be an arbitrary dominating set of G'. Observe that the rectangles dominated by small plates and the $S_{i,j}$ are also dominated by the appropriate big plate or R_i^h respectively R_i^v. Hence we can replace these small plates and $S_{i,j}$'s and obtain a dominating set D' with $|D'| \leq |D|$, where all small plates and $S_{i,j}$ are dominated by big plates and rectangles of type R_i^h and R_i^v. Let $R^2[D'] = \{R_i^h, R_i^v \mid R_i^h, R_i^v \in D'\}$ be the rectangles for v_i for which both the horizontal and the vertical version occur in D', $R^1[D']$ the remaining rectangles of type R_i^h and R_i^v (i.e. rectangles for v_i for which only one version occurs in D'), and let $P[D']$ denote the big plates in D'. Furthermore, let $R[D'] = R^2[D'] \cup R^1[D']$. Note that $R^2[D'] \cap R^1[D'] = \emptyset$.

Consider $C = \{v_i \mid R_i^h \in D' \text{ or } R_i^v \in D'\}$. Since all $S_{i,j}$ are dominated by $R[D']$, C is a vertex cover. Observe that to dominate all plates of i, $P_i \in D'$, or $R_i^h, R_i^v \in D'$. This holds for all i. Thus $|P[D']| + |R^2[D']|/2 \geq n$. Also, as C is a vertex cover of G, $|R^1[D']| + |R^2[D']|/2 = |C| \geq k$.
Hence

$$\begin{aligned} |D'| &\geq |R^1[D']| + |R^2[D']| + |P[D']| \\ &\geq |R^1[D']| + |R^2[D']|/2 + n \\ &\geq k + n. \end{aligned} \tag{8.2}$$

Together with Equation 8.1, this implies that $m^*(f(x)) = n + k$.

Now suppose that $|D| = m^*(f(x)) + c$, for a certain $c \geq 0$. Then $|D'| \leq |D| = n + k + c$. Using Equation 8.2,

$$\begin{aligned} |R^1[D']| + |R^2[D']|/2 + n &\leq n + k + c \\ |R^1[D']| + |R^2[D']|/2 &\leq k + c \\ |C| &\leq m^*(x) + c \\ |C| - m^*(x) &\leq c. \end{aligned}$$

This gives an L-reduction from Minimum Vertex Cover on graphs of degree three to Minimum Dominating Set on rectangle intersection graphs with $\alpha = 5$ and $\beta = 1$. □

Note that this theorem holds even if the rectangles have to be axis-parallel or if no rectangle can be fully contained in another rectangle (by slightly changing the construction of Figure 8.4). Furthermore, the construction in the proof of Theorem 8.5.8 can be replicated using ellipses instead of rectangles. This gives the following theorem.

Theorem 8.5.9 *Minimum Dominating Set on ellipse intersection graphs is APX-hard. Hence it has no ptas, unless P = NP.*

The proof of Theorem 8.5.9 requires ellipses of relatively high eccentricity (of the order $\sqrt{1-n^{-2}}$) as 'R_i^h' and 'R_i^v'. Hence the proof does not immediately carry over to disk graphs.

The construction to prove the APX-hardness of Minimum Dominating Set in rectangle intersection graphs can be extended to prove the APX-hardness of Minimum Connected Dominating Set. In fact, the new construction generalizes the previous construction, as it can also be used to prove the APX-hardness of Minimum Dominating Set. Below, we give this generalized proof.

Theorem 8.5.10 *Minimum Connected Dominating Set on rectangle intersection graphs is APX-hard. Hence it has no ptas, unless P=NP.*

Proof: Consider again an arbitrary instance x of Minimum Vertex Cover on graphs of degree three. Let $G = (V = \{v_1, \ldots, v_n\}, E)$ be the graph of x and denote the cardinality of the smallest vertex cover of x by k. We keep the construction of Theorem 8.5.8 (see Figure 8.4) and extend it as follows (see Figure 8.5). For any big plate P_i, we add a horizontal plate P_i^h intersecting P_i and containing a single small rectangle, ensuring that P_i^h is in any connected dominating set. We also add three surrounding rectangles S_1, S_2, and S_3, each containing a single small rectangle, enforcing the presence of S_1, S_2, and S_3 in any connected dominating set. These rectangles are aligned such that S_1 intersects all horizontal rectangles R_i^h, S_2 intersects S_1 and all vertical rectangles R_i^v, and S_3 intersects S_2 and all horizontal plates P_i^h. The intersection graph G' of these rectangles is the function $f(x)$ for the L-reduction. It can be quickly verified that this is indeed a polynomial-time computable function (even if the rectangles are part of the output).

Let C be a vertex cover of G of cardinality k. Recall that $R^h[C] = \{R_i^h \mid v_i \in C\}$, $R^v[C] = \{R_i^v \mid v_i \in C\}$, and $\overline{P[C]} = \{P_i \mid v_i \notin C\}$. Let $P^h = \{P_i^h \mid i = 1, \ldots, n\}$. We claim that $D = R^h[C] \cup R^v[C] \cup \overline{P[C]} \cup P^h \cup \{S_1, S_2, S_3\}$ is a connected dominating set of G'. From the proof of Theorem 8.5.8 and the construction of G', it should be clear that D is a dominating set for G'.

To prove that D induces a connected subgraph of G', let $d, d' \in D$ be any two distinct rectangles in G'. We show there exists a path in $G'[D]$ between d and d'. If d or d' is in $R^h[C]$ ($R^v[C]$), it takes one step to reach S_1 (S_2). Similarly, if d or d' is in $\overline{P[C]}$, the appropriate horizontal plate can be used to reach S_3 in two steps. Thus from d or d', we can reach S_1, S_2, or S_3 in

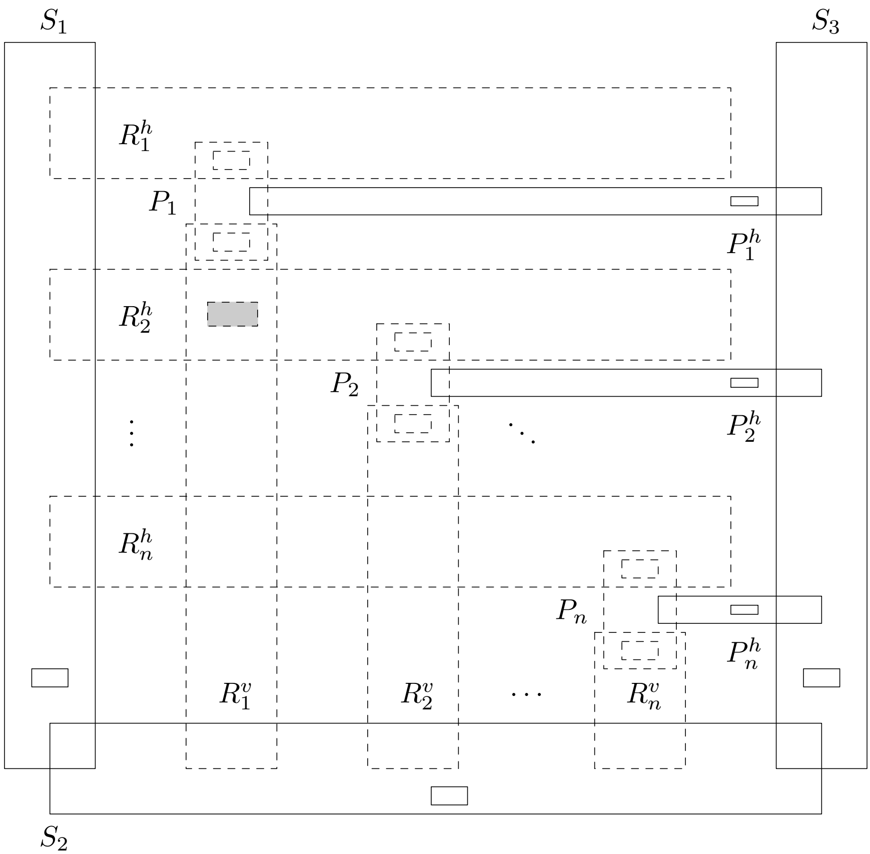

Figure 8.5: The intersection graph used in the APX-hardness proof of Minimum Connected Dominating Set. The rectangles that also appeared in Figure 8.4 have dashed boundaries.

the subgraph of G' induced by D in at most two steps. But since $\{S_1, S_2, S_3\}$ form a connected induced subgraph in G', this implies that D is a connected subgraph of G'. Hence D is a connected dominating set.

We now give an upper bound to $|D|$. Since the degree of G is three, $k \geq n/4$, and thus

$$\begin{aligned} m^*(f(x)) &\leq |D| \\ &= |R^h[C]| + |R^v[C]| + |\overline{P[C]}| + |P^h| + 3 \\ &= k + k + (n-k) + n + 3 \\ &= 2n + k + 3 \\ &\leq 9k + 3 \\ &\leq 12k \\ &= 12 \cdot m^*(x). \end{aligned} \tag{8.3}$$

Now let D be an arbitrary connected dominating set of G'. We may assume that D contains all P_i^h and S_1, S_2, S_3 (if not, the small rectangle contained in these rectangles is in D, which can be easily replaced by the bigger rectangle). Similarly, as already noted in the proof of Theorem 8.5.8, we may assume that all small plates and $S_{i,j}$ are dominated by big plates and rectangles R_i^h and R_i^v. Let $R^2[D'] = \{R_i^h, R_i^v \mid R_i^h, R_i^v \in D\}$ be the set of the rectangles for i for which both the horizontal and the vertical version occur in D, $R^1[D]$ the set of remaining rectangles of type R_i^h and R_i^v (i.e. rectangles for i for which only one version occurs in D), and let $P[D]$ denote the set of big plates in D. Furthermore, let $R[D] = R^2[D] \cup R^1[D]$. Note that $R^2[D] \cap R^1[D] = \emptyset$.

Consider $C = \{v_i \mid R_i^h \in D \text{ or } R_i^v \in D\}$. Since all $S_{i,j}$ are dominated by $R[D]$, C is a vertex cover. Observe that to dominate the small plates of i, $P_i \in D$, or both $R_i^h, R_i^v \in D$. This holds for all i. Therefore $|P[D]| + |R^2[D]|/2 \geq n$. Also, as C is a vertex cover for G, $|R^1[D]| + |R^2[D]|/2 = |C| \geq k$. Hence

$$\begin{aligned} |D| &\geq |R^1[D]| + |R^2[D]| + |P[D]| + |P^h| + 3 \\ &\geq |R^1[D]| + |R^2[D]|/2 + n + n + 3 \\ &\geq k + 2n + 3. \end{aligned} \tag{8.4}$$

Together with Equation 8.3, this implies that $m^*(f(x)) = 2n + k + 3$.

Now suppose that $|D| = m^*(f(x)) + c$, for a certain $c \geq 0$. Then $|D| = 2n + k + 3 + c$. Using Equation 8.4,

$$\begin{aligned} |R^1[D]| + |R^2[D]|/2 + 2n + 3 &\leq |D| = 2n + k + 3 + c \\ |R^1[D]| + |R^2[D]|/2 &\leq k + c \\ |C| &\leq m^*(x) + c \\ |C| - m^*(x) &\leq c. \end{aligned}$$

This gives an L-reduction with $\alpha = 12$ and $\beta = 1$. □

This reduction can also be extended to ellipse intersection graphs (where the ellipses have high eccentricity).

Part III

Approximating Geometric Coverage Problems

Overview

One of the most fundamental and best-known optimization problems is Minimum Set Cover. It has many applications, for instance in wireless network planning, as discussed in Section 1.2.2 and below. Given this particular application, it is natural to consider Minimum Set Cover in a geometric setting. Hence this part of the thesis is devoted to the approximability of the geometric version of Minimum Set Cover, as well as of several of its variants.

Problems

For sake of completeness, we start by defining Minimum Set Cover. Throughout this part, let $\mathbb{U}$ be a universe, $\mathcal{P} \subseteq \mathbb{U}$ a set of *elements*, and $\mathcal{S}$ a set of subsets of $\mathbb{U}$.

Definition III.1 *The* minimum set cover *problem is to find a smallest set $C \subseteq \mathcal{S}$ such that each element of $\mathcal{P}$ is contained in (covered by) a set in C.*

Minimum Set Cover where $\mathbb{U} = \mathbb{R}^d$ for some $d > 0$ will be called *Geometric Set Cover*. We will mostly discuss the case where $\mathbb{U} = \mathbb{R}^2$ and the sets in $\mathcal{S}$ are induced by simple geometric objects, such as disks or squares.

A first variation of Minimum Set Cover is its weighted case, where the sets in $\mathcal{S}$ are given a weight and we look for a cover of $\mathcal{S}$ of minimum total weight. We can extend further in this direction by assuming that each element u of $\mathcal{P}$ has a *profit* $p(u)$, each set $\mathcal{S}_i$ of $\mathcal{S}$ a *cost* $c(\mathcal{S}_i)$, and that we are given a budget B. This leads to the following problem.

Definition III.2 *The* budgeted maximum coverage *problem is find some set $C \subseteq \mathcal{S}$ of total cost at most B that maximizes the total profit of the elements covered by C.*

The geometric version is called *Geometric Budgeted Maximum Coverage*.

Two further variants of Minimum Set Cover that we will be particularly interested in are Unique Coverage and Minimum Membership Set Cover. In the former problem, one is given a collection of sets of elements from some universe and aims to select sets that maximize the number of elements contained in precisely one selected set. In the latter problem, the goal is to cover all elements of the universe while minimizing the maximum number of sets in which any element is contained. We discuss and formally define both problems below.

The unique coverage problem was proposed by Demaine et al. [83] and is mainly motivated by wireless network planning. Providers of wireless communication networks provide service for their customers. This can be achieved by placing a number of base stations that cover customer locations, which is a geometric set cover problem. However, if too many base stations cover a certain customer location, the resulting interference might cause this customer to receive no service at all. Ideally, each customer is serviced by exactly one base station and service is provided to as many customers as possible.

Definition III.3 *Given a set $C \subseteq \mathcal{S}$, an element $u \in \mathcal{P}$ is* uniquely covered *by C if there is precisely one $s \in C$ containing u. The* (maximum) unique coverage *problem is to find a set $C \subseteq \mathcal{S}$ that maximizes the number of uniquely covered elements of $\mathcal{P}$.*

Of course, it might be more profitable to provide service to certain customers. Furthermore, placing a base station is costly and providers generally have a limited budget. This gives rise to the budgeted unique coverage problem.

Definition III.4 *The* budgeted unique coverage *problem is find a set $C \subseteq \mathcal{S}$ of total cost at most B that maximizes the total profit of the elements uniquely covered by C.*

In practice, mobile devices can distinguish between signals from different base stations. However, this capability is limited and decreases with the number of base stations in range. Demaine et al. [83] model this by *satisfactions* $s_0 = 0$, $s_1 \geq s_2 \geq \cdots \geq 0$, where an element (customer) u yields *satisfaction-modulated profit* $s_i \cdot p(u)$ if it receives service from exactly i base stations.

Definition III.5 *The* budgeted low-coverage *problem is to find a set $C \subseteq \mathcal{S}$ of total cost at most B that maximizes the satisfaction-modulated profit of the elements covered by C.*

Another way to handle the limited capability of mobile devices to distinguish signals from different base stations, is to minimize the number of signals a device receives.

Definition III.6 *For a set $C \subseteq \mathcal{S}$, the* membership $\mathrm{mem}_C(u)$ *of an element $u \in \mathcal{P}$ is equal to the number of sets in C that contain u. The* maximum membership *is $\mathrm{mem}_C(\mathcal{P}) = \max_{u \in \mathcal{P}} \mathrm{mem}_C(u)$. Then the* minimum membership set cover *problem is to find a set $C \subseteq \mathcal{S}$ that covers all elements in $\mathcal{P}$ and minimizes $\mathrm{mem}_C(\mathcal{P})$.*

As mentioned before, we want to study these problems in their geometric version because of the close connection to wireless network planning. For instance, consider the case when the universe is the plane $\mathbb{R}^2$, $\mathcal{P}$ is a set of points corresponding to customer locations, and each $s \in \mathcal{S}$ is a geometric object modeling the broadcasting range of the corresponding base station. If all base stations are equivalent and we ignore obstacles to the signal, these geometric objects are unit disks.

One can make the problem more realistic by assuming that the base stations may have different broadcasting ranges and that they are hindered by obstacles, but the overlap of the broadcasting ranges of the potential base station locations is bounded. The latter assumption is reasonable, as in practice there are usually very few spots where a base station can or may be placed. We model this by a set of fat objects where any point in the plane is overlapped by a bounded number of objects, i.e. a set of fat objects of bounded ply.

Previous Work

Geometric Set Cover

Minimum Set Cover can be approximated within $1 + \ln|\mathcal{S}|$ by a greedy algorithm, even in the weighted case [156, 197, 66]. Hochbaum [149] gives a survey on these and other approximation algorithms for Minimum Set Cover. The greedy algorithm is also optimal. That is, Minimum Set Cover has no polynomial-time algorithm attaining an approximation ratio of $(1-\epsilon)\ln|\mathcal{S}|$ for any $\epsilon > 0$, unless NP $\subset$ DTIME$(n^{O(\log\log n)})$ [108].

Geometric Set Cover is NP-hard on unit squares and on unit disks [110, 157], even if the point set $\mathcal{P}$ corresponds to the centers of the squares or disks. By reducing from Minimum Dominating Set on unit disk/square graphs of bounded density, the NP-hardness continues to hold if the density is 1 [67].

The best approximation ratio for Geometric Set Cover is steadily being improved. Brönnimann and Goodrich [45] gave the first constant-factor approximation algorithm for Geometric Set Cover on unit disks, attaining an unspecified approximation ratio. Călinescu et al. [49] strengthened this result to a 108-approximation algorithm. Narayanappa and Vojtěchovský [218] built on the ideas of this algorithm to give a 72-approximation algorithm. The algorithm by Carmi, Katz, and Lev-Tov [52] has approximation ratio 38.

Recently a ptas for Geometric Set Cover on general disks was discovered. By a simple transformation, this problem is equivalent to the geometric version of *Minimum Hitting Set* (where a smallest subset of points hitting each object must be found) on three-dimensional half-spaces. Mustafa and Ray [217] give a ptas for this problem, as well as for *Geometric Hitting Set* on pseudo-disks.

The above algorithms are not known to be applicable to the weighted case. Ambühl et al. [13] give a 72-approximation algorithm for Weighted Geometric Set Cover on unit disks. A 2-approximation algorithm on unit squares is given by Mihalák [209].

Lev-Tov and Peleg [191] give a ptas for a special case of Geometric Set Cover on unit disks where the set $\mathcal{P}$ is a subset of the set of disk centers. Liao and Hu [193] consider the case where all points of $\mathcal{P}$ lie on the corners of a constant-size grid and give a ptas that extends to the weighted case.

Clarkson and Varadarajan [68] present a constant-factor approximation algorithm on pseudo-disks, attaining an unspecified approximation ratio.

Glaßer, Reitwießner, and Schmitz [125] consider a multi-objective version of Geometric Set Cover on unit disks, thus providing a trade-off between minimizing the number of selected disks and maximizing the number of covered points. They give a polynomial-time approximation scheme, meaning that for any (fixed) $\epsilon > 0$ an ϵ-optimal Pareto curve is output in polynomial time, i.e. for every dominating solution on the Pareto curve, the returned curve has an ϵ-approximate solution. It should be noted that Glaßer, Reitwießner, and Schmitz restrict to the case where the selected disks must have constant precision. Several inapproximability results are also given.

In three dimensions, Laue [188] showed that Geometric Set Cover has a constant-factor approximation algorithm on translated copies of a fixed polytope. The polytope need not be convex or fat.

Budgeted Maximum Coverage has a $(1 - \frac{1}{e})$-approximation algorithm in both the unit cost [264, 151, 149] and the general case [165]. Khuller, Moss, and Naor [165] proved that no polynomial-time algorithm can obtain an approximation ratio better than $(1 - \frac{1}{e})$, unless $\mathrm{NP} \subset \mathrm{DTIME}(n^{\mathrm{O}(\log \log n)})$.

As far as we know, the budgeted version of Minimum Set Cover has not been considered yet on unit squares or unit disks. Geometric Budgeted Maximum Coverage on unit disks or unit squares, even if the density is 1 and the point set $\mathcal{P}$ corresponds to the centers of the disks or squares, can easily be shown to be NP-hard by reduction from Geometric Set Cover.

Geometric Unique Coverage and Membership Set Cover

Demaine et al. [83] formulated the unique coverage problem and studied it in its general setting. They present a polynomial-time $\Omega(1/\log \rho) = \Omega(1/\log n)$ approximation algorithm using a greedy method, where n is the number of elements and ρ is one plus the ratio of the maximum number and the minimum number of sets in which an element is contained. They give hardness results to show that this algorithm is (near-)optimal. Demaine et al. give a gap-preserving reduction from a variant of Balanced Binary Independent Set, so any (in)approximability result for this problem holds for Unique Coverage as well. Hence, for any $\epsilon > 0$, it is hard to approximate Unique Coverage within ratio $\Omega(1/\log^{\sigma(\epsilon)} n)$, assuming that $\mathrm{NP} \not\subseteq \mathrm{BPTIME}(2^{n^{\epsilon}})$, where $\sigma(\epsilon)$ is some constant dependent on ϵ. Under the hypothesis that refuting random instances of 3SAT is hard on average, one can strengthen this to $\Omega(1/\log^{1/3-\epsilon} n)$ for any $\epsilon > 0$. By making a (plausible) assumption on the hardness of Balanced Binary Independent Set, a further strengthening to $\Omega(1/\log n)$ is possible.

The *unique hitting set* problem, where one tries to select elements to uniquely hit as many sets as possible, is equivalent to Unique Coverage. If sets have cardinality at most k, Guruswami and Trevisan [134] give an $\Omega(1/\log k)$-approximation algorithm. They also consider the more general 1-*in*-k-*SAT* problem and give a $1/e$-approximation algorithm on satisfiable instances.

Moser, Raman, and Sikdar [216] showed that if Unique Coverage is parameterized by the number k of elements to cover uniquely, it is in FPT. However, Budgeted Unique Coverage parameterized by k and the budget B is not in FPT, unless P=NP. If parameterized by B and the profits and costs are integer, the problem is not in FPT unless FPT=W[1].

As far as we know, the unique coverage problem as is and its extensions have not been studied in a geometric setting, although several related problems have. If one tries to maximize the number of points that are uniquely covered subject to the constraint that the selected objects are disjoint, this is a maximum-weight independent set problem, which has a ptas on fat objects of arbitrary ply (see Chapter 7). For a grid-based version of Unique Coverage on unit disks, where the disks are restricted to lie at grid points, a ptas has been announced by Lev-Tov and Peleg [190]. The eptas for fat objects of bounded ply given in Chapter 10 is significantly more general than this result. Recently, Glaßer, Reitwießner, and Schmitz [125] considered a multi-objective version of Geometric Unique Coverage on unit disks, where the selected disks must have constant precision. Their approximation scheme outputs an ϵ-optimal Pareto curve in polynomial time for any fixed $\epsilon > 0$.

Minimum Membership Set Cover has a $O(\ln n)$-approximation algorithm, but no polynomial-time $(1 - \epsilon) \ln n$-approximation algorithm for any $\epsilon > 0$, unless NP$\subset$ DTIME$(n^{O(\log \log n)})$ [186]. The minimum membership set cover problem has not been studied yet in its geometric setting.

Further Variations

There are many variations of Geometric Set Cover. For instance, suppose that the point set and the centers of the disks are given, but we are free to choose the radii of the disks. Lev-Tov and Peleg [191] give a ptas that minimizes the sum of the radii. Bilò et al. [34] provide a more general scheme that can also be applied if at most k radii may be nonzero. This problem, called *k-clustering*, is also considered by Alt et al. [11].

The related *k-center* clustering problem, where the maximum radius has to be minimized, is also well-studied and has a ptas in many cases. A ptas for the two-dimensional case was given by Agarwal and Procopiuc [3]. Bădoiu, Har-Peled, and Indyk [20] improved this ptas and generalized to arbitrary dimension. Agarwal and Procopiuc [3] give a nice overview of earlier results.

If the radius of the disks is fixed, but the positions of the centers of the disks may be chosen freely, we obtain the *geometric covering* problem. This problem is NP-hard on unit disks and on unit squares [110, 157] and a ptas is known both on unit disks [150] and on unit squares [128]. The variation where at most k disks may be selected also has a ptas [113].

When we are given both a set of points that should be covered by the objects and a set of points that should not be covered, we obtain a generalization of the geometric covering problem called the *class cover* problem. Cannon and Cowen [50] give a ptas on unit disks. Efrat et al. [96] give an approximation algorithm on ellipses, where the eccentricity of the ellipses may also be chosen freely. The approximation ratio is logarithmic in the optimum.

A closely related problem is *Geometric Piercing*, where given a set of objects, a minimum set of points should be found piercing (hitting) all objects. This problem has a ptas on fat objects [57].

Glaßer, Reith, and Vollmer [124] consider a maximum coverage problem where, for each (unit) disk and each point of $\mathcal{P}$ in the disk, we are given a 'signal strength'. Then a point of $\mathcal{P}$ is 'supplied' if it is covered by a disk whose signal strength for this point is higher than the sum of the signal strengths of other selected disks covering this point. If the given set of unit disks has constant precision, Glaßer, Reith, and Vollmer give a ptas for both the case where at most k disks can be selected and the number of supplied points must be maximized, and the case where at least l points must be supplied and the number of selected disks is minimized.

A problem similar to Geometric Unique Coverage is the problem to find a subset of a given set of objects that maximizes the total area that is uniquely covered. Chen et al. [63] provide a ptas for this problem on unit disks.

Chapter 9

Geometric Set Cover and Unit Squares

Geometric Set Cover can be approximated better than general Minimum Set Cover, but for many object types the approximability has not been settled yet. We give a ptas for Geometric Set Cover on unit squares, improving on the earlier 2-approximation algorithm [209]. This is the one of the first approximation schemes for Geometric Set Cover on two-dimensional objects (together with the recently appeared [217]) and the first that extends to the weighted case. The scheme in fact extends to the more general budgeted maximum coverage problem. Moreover, we prove that the scheme essentially has optimal running time (up to constants), unless the exponential time hypothesis is false.

Besides these positive algorithmic results, we also give several negative results. We show that on convex polygons, translated copies of a single polygon, rotated copies of a single convex polygon, and α-fat objects, Geometric Set Cover is as hard as Minimum Set Cover. These hardness results all carry over (mutatis mutandis) to the budgeted case. If the polygons have constant description complexity, Geometric Set Cover is still APX-hard on convex polygons. We also obtain APX-hardness results for Geometric Set Cover on axis-parallel rectangles and ellipses.

9.1 A ptas on Unit Squares

We consider Geometric Set Cover on unit squares and show that it has a ptas by applying the shifting technique.

So let $\mathcal{P}$ be a set of points and $\mathcal{S}$ a set of axis-aligned unit squares. For sake of notation, when referring to the (x, y)-coordinates of a square, we mean the coordinates of the bottom left corner of that square. For a square s, the x-coordinate of (the bottom left corner of) s is denoted by $x(s)$, while the y-coordinate is denoted by $y(s)$. By scaling and translating (as in Chapter 4), we can assume that no horizontal (vertical) boundary of a square is on the same line as the horizontal (vertical) boundary of another square. Furthermore, we can assume that all points are fully contained in the squares they are in, i.e. no point lies on the boundary of a square. Finally, we assume that none of the square or point coordinates are integers.

Consider the horizontal lines $y = h$ ($h \in \mathbb{Z}$). They partition the plane into horizontal slabs of height 1. Any point is contained in a slab and every square intersects precisely one line. Let $k \geq 1$ be some integer we determine later. For any k consecutive slabs, the points in these slabs must be covered by a subset of the squares intersecting the $k + 1$ horizontal lines defining those k slabs. Using the shifting technique, it suffices to prove that we can optimally solve Geometric Set Cover on unit squares if we restrict to k consecutive slabs and the $k + 1$ lines defining them.

Theorem 9.1.1 *For any instance of Geometric Set Cover on a set of unit squares $\mathcal{S}$ where all points of $\mathcal{P}$ are inside $k \geq 1$ consecutive height 1 horizontal slabs, one can find an optimal solution in $\mathrm{O}((3|\mathcal{S}|)^{4k+4}\,|\mathcal{P}|)$ time.*

The next few pages are devoted to proving this theorem.

The idea will be to apply a sweep-line algorithm. This requires that we somehow bound the number of squares of an optimal solution that intersect the sweep-line at a given sweep-line position. To this end, consider the subset of squares of an optimum solution intersecting a horizontal line $y = h$ for some $h \in \mathbb{Z}$. Any such square must appear on the lower or upper envelope of this subset, or all points it covers would be covered by other squares. This is because the union of a set of connected axis-aligned squares intersecting a common axis-parallel line cannot have any holes. Following this observation, for each position of the sweep-line and for each of the $k + 1$ integer horizontal lines, we should consider at most two squares intersecting the sweep-line: one that will appear on the upper envelope and one that will appear on the lower envelope of the final solution. Obviously, we could also have a single square both on the upper and lower envelope, or, in fact, have no square at all.

Although at first glance this approach seems feasible, there is a problem with the dynamic programming. A square might appear on the lower envelope for some position of the sweep-line and on the upper envelope for a later position. In fact, several other squares might appear on the lower or upper envelope before this square appears on the upper envelope. This makes it difficult to avoid counting certain squares twice. To circumvent this, we split the sweep-line into k parts, one part per slab. We move these parts at different speeds, but always in such a way that if a square appears both on the lower and the upper envelope, then the split sweep-line is positioned such that it intersects the square both at the point where the square appears on the lower and on the upper envelope.

Though intuitively it seems like this would work, the split sweep-line trick requires a rigorous proof. We do this by formalizing the sweep-line process.

Just as in any sweep-line algorithm, we maintain a data structure (the *front*) containing the squares that are 'active' at a given position of the sweep-lines. The difficulty in this sweep-line algorithm arises in maintaining the front and consequently in finding squares that can be validly inserted into the front. Therefore we start by defining the front that we use and the (four types of) insertions that we are allowed to perform.

Let $\mathcal{S}^l$ and $\mathcal{S}^r$ be two dummy sets of $k+1$ squares each, such that the squares in $\mathcal{S}^l$ ($\mathcal{S}^r$) are to the left (right) of all squares in $\mathcal{S}$ and each integer horizontal line intersects precisely one square of $\mathcal{S}^l$ and one square of $\mathcal{S}^r$. Let $\overline{\mathcal{S}} = \mathcal{S} \cup \mathcal{S}^l \cup \mathcal{S}^r$. Given some set $S \subseteq \overline{\mathcal{S}}$, let S_i denote the set of squares in S intersecting line i. Let $R_i \subseteq S_i$ be the set containing precisely:

- the rightmost square of S_i (denote it by s_i),
- those squares s that overlap part of the left boundary of s_i and whose right boundary is not fully covered by squares of S_i.

We now define a front. For a better understanding of the definition, imagine that the squares are being inserted in order of increasing x-coordinate and that we want to keep track of the upper and lower envelope of each line i.

Definition 9.1.2 *Let S be the union of $\mathcal{S}^l$ and some subset of $\overline{\mathcal{S}}$. Then a* front *$F = \{u_1, \ldots, u_{k+1}, l_1, \ldots, l_{k+1}, b_1, \ldots, b_{k+1}, x_1, \ldots, x_k\}$ for S has the following properties:*

- *$u_i, l_i \in R_i$ with $u_i = s_i$ or $l_i = s_i$,*
- *$y(s) \leq y(u_i)$ for any $s \in S_i$ to the right of u_i (i.e. with $x(s) > x(u_i)$),*
- *$y(s) \geq y(l_i)$ for any $s \in S_i$ to the right of l_i (i.e. with $x(s) > x(l_i)$),*
- *b_i is equal to:*
  - *the lowest square of S_i to the right of l_i if $x(u_i) > x(l_i)$,*
  - *the highest square of S_i to the right of u_i if $x(l_i) > x(u_i)$,*
  - *s_i if $x(u_i) = x(l_i)$ (i.e. if $u_i = l_i$),*
- *x_i is equal to the larger of the x-coordinate from which l_{i+1} starts appearing on the lower envelope of S_{i+1} and the x-coordinate from which u_i starts appearing on the upper envelope of S_i.*

An example is depicted in Figure 9.1.

Fronts are the representative of the current state of the sweep-line algorithm. The squares u_i and l_i track the current square on respectively the upper and the lower envelope of line i. The value of x_i is the x-coordinate of the part of the sweep-line between lines i and $i+1$. The square b_i is used in checking if a certain square may be inserted or not.

We can make two observations about a front. First, $y(u_i) \geq y(l_i)$ and, since $u_i, l_i \in R_i$, $|x(u_i) - x(l_i)| < 1$ for any $i = 1, \ldots, k+1$. Secondly, if $x(u_i) \geq x(l_i)$, then $y(u_i) \leq y(b_i) \leq y(l_i)$. If $x(u_i) \leq x(l_i)$, then $y(l_i) \leq y(b_i) \leq y(u_i)$.

For a given front, we distinguish four types of insertions that are possible. An upper-insertion for squares that will appear only on the upper envelope for some line, a lower-insertion for squares appearing only on the lower envelope,

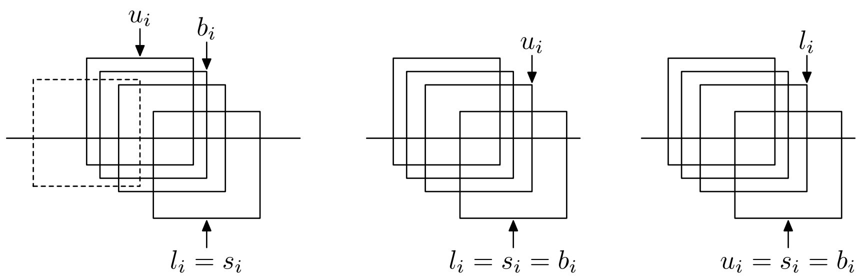

Figure 9.1: The left figure shows a set S_i. The dashed square is not in R_i and thus not in a front for S_i. The four solid squares are in R_i. From Definition 9.1.2, the labeling in the figure is correct. The middle figure shows the same set R_i, but with a different (and still correct) labeling. The labeling in the right figure however is incorrect.

and a middle-insertion or a skip-insertion for squares appearing on both envelopes. We define these four insertions, describe when they may be applied, and prove that any geometric set cover can be obtained using these insertions.

From now on, S will denote the union of $\mathcal{S}^l$ and some subset of $\overline{\mathcal{S}}$.

Definition 9.1.3 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k\}$. We say that s is* upper-insertable *into F if all of the following hold:*

1. *$y(s) > y(l_i)$ and if $x(l_i) > x(u_i)$, then $y(s) > y(b_i)$,*
2. *$x(s) \in (x(l_i), x(l_i) + 1]$ and $x(s) \in (x(u_i), x(u_i) + 1]$,*
3. *$x'_i > x_i$,*
4. *any point of $\mathcal{P}$ in $[x_i, x'_i] \times [i, i+1]$ is covered by u_i or l_{i+1},*

where x'_i is the x-coordinate from which s is on the upper envelope of $(S \cup \{s\})_i$.

Condition 1 ensures that s lies above l_i and all squares between u_i and l_i (represented by b_i), Condition 2 ensures that s appears on the upper envelope of $(S \cup \{s\})_i$, Condition 3 ensures that this appearance happens after u_i appears on the upper envelope, and Condition 4 ensures that we cover all points between two consecutive sweep-line positions. An example of upper-insertable squares and squares that are not upper-insertable is given in Figure 9.2.

Proposition 9.1.4 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k\}$. If Condition 2 of Definition 9.1.3 holds for F and s, then s appears on the upper envelope of $(S \cup \{s\})_i$ to the right of u_i.*

Proof: By Condition 2, $x(s) > \max\{x(u_i), x(l_i)\} = x(s_i)$, and thus s appears on the upper envelope of $(S \cup \{s\})_i$ to the right of u_i. $\square$

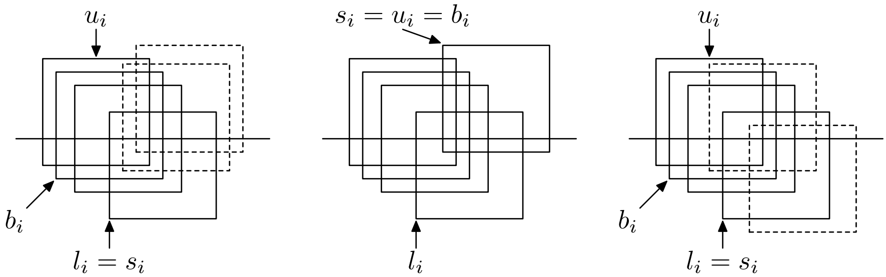

Figure 9.2: The left figure shows two (dashed) squares that are upper-insertable into the front of Figure 9.1. The middle figure shows the resulting front after upper-inserting the rightmost of these squares. The right figure shows two (dashed) squares that are not upper-insertable.

As a consequence of this proposition, the x-coordinate x'_i of Definition 9.1.3 does indeed exist (if Condition 2 holds).

Lemma 9.1.5 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k\}$ that is upper-insertable into F. Then between the appearance of u_i and the appearance of s on the upper envelope of $(S \cup \{s\})_i$ no other squares appear on the upper envelope of $(S \cup \{s\})_i$.*

Proof: If $u_i = s_i$, this follows from $x(s) > x(u_i) = x(s_i)$ and $x'_i > x_i$. So assume that $u_i \neq s_i$. Then $l_i = s_i$ and $x(l_i) > x(u_i)$. Recall the definition of a front and observe that b_i is the highest square of S_i to the right of u_i. As $x(l_i) - x(u_i) < 1$ and $y(b_i) < y(u_i)$, it suffices for s to lie above b_i (i.e. $y(s) > y(b_i)$) and for s to cover the x-range $[x(u_i)+1, x(l_i)+1]$ (i.e. $x(l_i) < x(s) < x(u_i)+1$). This holds from the definition of upper-insertable. □

Lemma 9.1.6 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k\}$ that is upper-insertable into F. Then $S \cup \{s\}$ has a front F' equal to F, except u_i is replaced by s, x_i is set to x'_i, where x'_i is equal to $x(s)$ if $y(s) > y(u_i)$ and to $x(u_i) + 1$ otherwise, and if $x(u_i) \leq x(l_i)$ or $y(s) \leq y(b_i)$, b_i is set to s.*

Proof: Since $x(s) > \max\{x(u_i), x(l_i)\} = x(s_i)$ by Condition 2 of Definition 9.1.3, we can replace u_i by s. Note that l_i can remain the same by Condition 1 and 2. By Lemma 9.1.5, x'_i is indeed the x-coordinate from which s appears on the upper envelope of $(S \cup \{s\})_i$. From Condition 3, x_i should be set to x'_i. If $x(u_i) \leq x(l_i)$, then as $x(s) > x(l_i)$, b_i should be set to s. If $x(u_i) > x(l_i)$, then b_i should only be changed if s lies below b_i, i.e. if $y(s) \leq y(b_i)$. Then the front F' is indeed a front for $S \cup \{s\}$. □

Constructing the front F' from F as prescribed in the lemma statement is called the *upper-insertion* of s into F.

Definition 9.1.7 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{2, \ldots, k+1\}$. We say that s is* lower-insertable *into F if all of the following hold:*

1. $y(s) < y(u_i)$ *and if* $x(u_i) > x(l_i)$*, then* $y(s) < y(b_i)$,
2. $x(s) \in (x(l_i), x(l_i) + 1]$ *and* $x(s) \in (x(u_i), x(u_i) + 1]$,
3. $x'_{i-1} > x_{i-1}$,
4. *any point of* $\mathcal{P}$ *in* $[x_{i-1}, x'_{i-1}] \times [i-1, i]$ *is covered by* u_{i-1} *or* l_i.

Here x'_{i-1} is the x-coordinate by which s is on the lower envelope of $(S \cup \{s\})_i$.

Lemma 9.1.8 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{2, \ldots, k+1\}$ that is lower-insertable into F. Then $S \cup \{s\}$ has a front F' equal to F, except l_i is replaced by s, x_{i-1} is set to x'_{i-1}, where x'_{i-1} is equal to $x(s)$ if $y(s) < y(l_i)$ and to $x(l_i) + 1$ otherwise, and if $x(l_i) \leq x(u_i)$ or $y(s) \geq y(b_i)$, b_i is set to s. Furthermore, between the appearance of l_i and the appearance of s on the lower envelope of $(S \cup \{s\})_i$ no other squares appear on the lower envelope of $(S \cup \{s\})_i$.*

Constructing the front F' from F as prescribed in the lemma statement is called the *lower-insertion* of s into F.

We define middle-insertable, which combines upper- and lower-insertable, except that we drop the constraints that $y(s) > y(l_i)$ and $y(s) < y(u_i)$.

Definition 9.1.9 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k+1\}$. We say that s is* middle-insertable *into F if all of the following hold:*

1. *if* $x(l_i) > x(u_i)$*, then* $y(s) > y(b_i)$*, and if* $x(u_i) > x(l_i)$*, then* $y(s) < y(b_i)$,
2. $x(s) \in (x(l_i), x(l_i) + 1]$ *and* $x(s) \in (x(u_i), x(u_i) + 1]$,
3. $x'_i > x_i$ *(if* $i \neq k+1$*) and* $x'_{i-1} > x_{i-1}$ *(if* $i \neq 1$*)*,
4. *any point of* $\mathcal{P}$ *in* $[x_i, x'_i] \times [i, i+1]$ *(if* $i \neq k+1$*) or* $[x_{i-1}, x'_{i-1}] \times [i-1, i]$ *(if* $i \neq 1$*) is covered by* u_{i-1}*,* l_i*,* u_i*, or* l_{i+1}*,*

where x'_i (x'_{i-1}) is the x-coordinate from which s appears on the upper (lower) envelope of $(S \cup \{s\})_i$.

Lemma 9.1.10 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k+1\}$ that is middle-insertable into F. Then $S \cup \{s\}$ has a front F' equal to F, except u_i, l_i, and b_i are replaced by s, x_i is set to x'_i (if $i \neq k+1$), where x'_i is equal to $x(s)$ if $y(s) > y(u_i)$ and to $x(u_i) + 1$ otherwise, and x_{i-1} is set to x'_{i-1} (if $i \neq 1$), where x'_{i-1} is equal to $x(s)$ if $y(s) < y(l_i)$ and to $x(l_i) + 1$ otherwise. Furthermore, between the appearance of u_i (l_i) and the appearance of s on the upper (lower) envelope of $(S \cup \{s\})_i$ no other squares appear on the upper (lower) envelope of $(S \cup \{s\})_i$.*

Constructing the front F' from F as prescribed in the lemma statement is called the *middle-insertion* of s into F.

Definition 9.1.11 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k+1\}$. We say that s is* skip-insertable *into F if all of the following hold:*

1. *$u_i = l_i$,*

2. *$x(s) > 1 + \max\{x(u_i), x(l_i)\}$,*

3. *(if $i \neq k+1$) $x(s) > x_i$ and (if $i \neq 1$) $x(s) > x_{i-1}$*

4. *any point of $\mathcal{P}$ in $[x_i, x(s)] \times [i, i+1]$ (if $i \neq k+1$) or in $[x_{i-1}, x(s)] \times [i-1, i]$ (if $i \neq 1$) is covered by u_{i-1}, l_i, u_i, or l_{i+1}.*

Lemma 9.1.12 *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k+1\}$ that is skip-insertable into F. Then $S \cup \{s\}$ has a front F' equal to F, except u_i, l_i, and b_i are replaced by s, x_i is set to $x(s)$ (if $i \neq k+1$), and x_{i-1} is set to $x(s)$ (if $i \neq 1$). Furthermore, between the appearance of u_i (l_i) and the appearance of s on the upper (lower) envelope of $(S \cup \{s\})_i$ no other squares appear on the upper (lower) envelope of $(S \cup \{s\})_i$.*

Constructing the front F' from F as prescribed in the lemma statement is called the *skip-insertion* of s into F.

In general, we call an upper-/lower-/middle-/skip-insertion an *insertion* and we say s is *insertable* if it is upper-/lower-/middle-/skip-insertable. A *valid insertion* is the upper- (respectively lower-/middle-/skip-) insertion of a square that is upper- (respectively lower-/middle-/skip-) insertable.

We now prove that any set cover can be obtained using a sequence of valid insertions. Denote by F^l and F^r the fronts for $\mathcal{S}^l$ and $\overline{\mathcal{S}}$.

Lemma 9.1.13 *Assume $\mathcal{P} = \emptyset$. Let S be some set such that $S = \mathcal{S}^l \cup S_i \cup \mathcal{S}^r$ for some $i \in \{1, \ldots, k+1\}$ and any square in S_i appears on the lower or the upper envelope of S_i. Then there is a sequence of $|S_i| + k - 1$ valid insertions starting from F^l, leading to fronts $F^l = F_0, F_1, \ldots, F_{|S_i|+k-1} = F^r$ such that for any square $s \in S_i$, there is a front F_j containing s.*

Proof: We assume that if $i = 1$, then no squares of S_i appear only on the lower envelope of S_i. Similarly, if $i = k+1$, assume that no squares of S_i appear only on the upper envelope of S_i. Order the squares in $S_i \backslash \mathcal{S}^l$ by increasing x-coordinate, i.e. $s_1, \ldots, s_{|S_i|-1}$. Note that the squares appearing on the upper envelope form an increasing subsequence of S_i. Similarly, the squares appearing on the lower envelope form an increasing subsequence. We claim that one can obtain the requested sequence of valid insertions by inserting s_j into F_{j-1} for all $j = 1, \ldots, |S_i| - 1$ as follows: if s_j appears

- only on the upper envelope of S_i, then s_j is upper-insertable and will be upper-inserted;
- only on the lower envelope of S_i, then s_j is lower-insertable and will be lower-inserted;
- on the upper and lower envelope of S_i and a square of S_i covers part of its left boundary, then s_j is middle-insertable and will be middle-inserted;
- on the upper and lower envelope of S_i and no square of S_i covers part of its left boundary, then s_j is skip-insertable and will be skip-inserted.

We prove this by induction on the number j of inserted squares.

Suppose that $j = 0$. Since s_1 is the leftmost square of $S_i \backslash \mathcal{S}^l$, it appears on both envelopes of S_i and no square of S_i covers part of its left boundary. By the definition of $F^l = F_0$, s_1 is skip-insertable into F_0 and can be skip-inserted.

Assume that $j > 0$ and consider the current front F_j. If s_{j+1} appears on both envelopes of S_i and no square of S_i covers part of its left boundary, then $x(s_{j+1}) > 1 + x(s_{j'})$ for any $j' < j + 1$. As squares are inserted in order of increasing x-coordinate, in F_j, $x(s_{j+1}) > 1 + \max\{x(u_i), x(l_i)\}$. Since $S = \mathcal{S}^l \cup S_i \cup \mathcal{S}^r$, this implies that (if $i \neq k + 1$) $x(s_{j+1}) > x_i$ and (if $i \neq 1$) $x(s_{j+1}) > x_{i-1}$. Finally, observe that s_j must appear on both envelopes of S_i and thus must have been middle- or skip-inserted. But then $u_i = l_i = s_j$. Hence s_{j+1} is skip-insertable into F_j.

If s_{j+1} appears only on the upper envelope of S_i, then there must be squares appearing on the lower envelope of S_i covering the bottom left corner of s_{j+1}. By induction, the rightmost such square must be l_i. Hence $y(s_{j+1}) > y(l_i)$ and $x(s_{j+1}) \in (x(l_i), 1 + x(l_i)]$. But then there are squares covering (part of) the left boundary of s_{j+1} appearing on the upper envelope of S_i. By induction, the right-most such square must be u_i and thus $x(s_{j+1}) \in (x(u_i), 1 + x(u_i)]$. As squares are inserted in order of increasing x-coordinate, s_{j+1} appears on the upper envelope of S_i after u_i. Then $x'_i > x_i$. Finally, suppose that $x(l_i) > x(u_i)$. By induction, any square $s \in S_i$ with $x(u_i) < x(s) \leq x(l_i)$, and in particular b_i, does not appear on the upper envelope of S_i. Therefore $y(s_{j+1}) > y(b_i)$. Hence s_{j+1} is upper-insertable into F_j.

The cases when s_{j+1} appears only on the lower envelope of S_i or when s_{j+1} appears on both envelopes of S_i and (part of) its left boundary is covered are similar. Finally, apply skip-insertions to insert the squares of $\mathcal{S}^r$. □

Lemma 9.1.14 *Assume $\mathcal{P} = \emptyset$. Let S be some subset of $\overline{\mathcal{S}}$ containing $\mathcal{S}^l \cup \mathcal{S}^r$, such that for the set S_i of squares in S intersecting line i for $i \in \{1, \ldots, k+1\}$, any square in S_i appears on the upper or lower envelope of S_i. Then there is a sequence of $|S| - k - 1$ valid insertions starting from $F_0 = F^l$, leading to $F_1, \ldots, F_{|S|-k-1} = F^r$ such that for any square $s \in S$, there is a front F_j containing s.*

Proof: Following the proof of the previous lemma, we can insert the squares intersecting each horizontal line in order of increasing x-coordinate. However, we should interleave the sequences of the different lines. For any $i = 1, \ldots, k$, consider the squares appearing on the upper envelope of S_i and the lower envelope of S_{i+1}. Order these squares according to the x-coordinate from which they appear on the upper envelope of S_i or on the lower envelope of S_{i+1} respectively. Combining these two orders, we can extend this to an order by which to insert the squares of S. We claim that the j-th square s_j according to this order is insertable into F_{j-1} and that after inserting s_j, all squares $s_{j'}$ with $j' > j$ are still insertable.

We prove this by induction on the number j of inserted squares. Trivially, all squares are insertable into F_0. Now consider any $j \geq 0$. By induction, s_{j+1} is insertable into F_j. Let F_{j+1} be the front arising from the insertion of s_{j+1} into F_j. Suppose that $s_{j+1} \in S_i$ for some i. As squares of S_i are inserted in order of increasing x-coordinate, it follows from Lemma 9.1.13 that all $s_{j'} \in S_i$ with $j' > j+1$ are still insertable into F_{j+1}.

To see that remaining squares in $S_{i'}$ for $i' \neq i$ are still insertable, it suffices to see that from the perspective of such a square s', only a change to $x_{i'}$ (if $i' \neq k+1$) or $x_{i'-1}$ (if $i' \neq 1$) can affect its insertability. We may thus assume that s' appears on the lower envelope of S_{i+1} (if $i \neq k+1$) or the upper envelope of S_{i-1} (if $i \neq 1$). Assume without loss of generality that $i \neq k+1$, s' appears on the lower envelope of S_{i+1}, and s_{j+1} appears on the upper envelope of S_i. Keeping in mind the definition of x_i (Definition 9.1.2), as all noninserted squares on the lower envelope of S_{i+1} start appearing on this envelope at a larger x-coordinate than the x-coordinate from which s_{j+1} appears on the upper envelope of S_i, the condition that $x'_i > x_i$ will still hold for s' in F_{j+1}.

The other case, when $i \neq 1$, s' appears on the lower envelope of S_{i-1}, and s_{j+1} appears on the lower envelope of S_i, is similar. We thus found a sequence of valid insertions as in the lemma statement. □

Lemma 9.1.15 *Let S be any smallest subset of $\overline{\mathcal{S}}$ containing $\mathcal{S}^l \cup \mathcal{S}^r$ and covering all points in $\mathcal{P}$. Then there is a sequence of $|S|-k-1$ valid insertions starting from F^l, leading to $F_1, \ldots, F_{|S|-k-1} = F^r$ such that for any square $s \in S$, there is a front F_j containing s.*

Proof: This follows from the preceding lemma. Note that in the definitions of insertable, the coverage constraints are satisfied by S. □

The converse of this lemma is also true.

Lemma 9.1.16 *Let $l \geq 0$. Then any sequence of $l+k+1$ valid insertions starting from F^l and resulting in F^r corresponds to a set $S \subseteq \mathcal{S}$ of cardinality l covering all points in $\mathcal{P}$.*

Proof: Take S to be the set of inserted squares, except those in $\mathcal{S}^r$. Because we only performed valid insertions, the set S covers all points in $\mathcal{P}$. □

Proof of Theorem 9.1.1: Construct a directed graph G with $V(G)$ equal to the set of all fronts and a directed edge from front F to F' if F' can be obtained from F by a single valid insertion. From the definition of a front, $|V(G)| = \mathrm{O}(|\overline{\mathcal{S}}|^{4k+3})$. As each front allows for at most $4|\overline{\mathcal{S}}|$ valid insertions, $|E(G)| = \mathrm{O}(|\overline{\mathcal{S}}|^{4k+4})$. Because the validity of an insertion can be checked in $\mathrm{O}(|\mathcal{P}|)$ time, G can be constructed in $\mathrm{O}(|\overline{\mathcal{S}}|^{4k+4}|\mathcal{P}|)$ time.

From Lemma 9.1.15 and 9.1.16, a shortest path in G from F^l to F^r corresponds to a minimum subset of $\mathcal{S}$ covering all points in $\mathcal{P}$. Using breadth-first search, a shortest path can be found in $\mathrm{O}(|E(G)|) = \mathrm{O}(|\overline{\mathcal{S}}|^{4k+4})$ time. Observe that $|\overline{\mathcal{S}}| = |\mathcal{S}| + |\mathcal{S}^l| + |\mathcal{S}^r| \leq 3|\mathcal{S}|$, because if no square intersects a certain line, we may ignore this line. Then the running time of the algorithm is $\mathrm{O}((3|\mathcal{S}|)^{4k+4}|\mathcal{P}|)$. □

Combining Theorem 9.1.1 with the shifting technique, we obtain a ptas for Geometric Set Cover on unit squares. For each integer $0 \leq a \leq k-1$, let L_a denote the set of squares intersecting a line $y \equiv a \pmod{k}$. Moreover, for each $b \in \mathbb{Z}$, let $\mathcal{P}_a^b$ denote the set of points between lines $y = bk + a$ and $y = (b+1)k + a$. Apply the algorithm of Theorem 9.1.1 to each such set $\mathcal{P}_a^b$ and denote the returned set of squares by C_a^b. Then let $C_a = \bigcup_{b \in \mathbb{Z}} C_a^b$ and let $C_{\min}$ denote a smallest such set. Trivially, each C_a is a geometric set cover for $\mathcal{P}$, as $\bigcup_{b \in \mathbb{Z}} \mathcal{P}_a^b = \mathcal{P}$ for any value of a.

Lemma 9.1.17 *$|C_{\min}| \leq (1 + 1/k) \cdot |OPT|$, where OPT is a minimum geometric set cover.*

Proof: Let $\mathcal{S}_a^b$ denote the set of all squares in $\mathcal{S}$ covering at least one point in $\mathcal{P}_a^b$ for some a, b. We may assume that $C_a^b \subseteq \mathcal{S}_a^b$. Now observe that $OPT \cap \mathcal{S}_a^b$ is a cover for $\mathcal{P}_a^b$. Hence $|C_a^b| \leq |OPT \cap \mathcal{S}_a^b|$ and

$$|C_a| \leq \sum_{b \in \mathbb{Z}} |OPT \cap \mathcal{S}_a^b| \leq |OPT| + |OPT \cap L_a|$$

for any $0 \leq a \leq k-1$. A square is in L_a for precisely one value of a. Then

$$k \cdot |C_{\min}| \leq \sum_{a=0}^{k-1} |C_a| \leq \sum_{a=0}^{k-1} \Big(|OPT| + |OPT \cap L_a|\Big) = (k+1) \cdot |OPT|.$$

Therefore $|C_{\min}| \leq (1 + 1/k) \cdot |OPT|$. □

Theorem 9.1.18 *There is a ptas for Geometric Set Cover on unit squares.*

Proof: Consider some $\epsilon > 0$ and let $k = \max\{1, \lceil 1/\epsilon \rceil\}$. Following Theorem 9.1.1, we can compute $C_{\min}$ in $\mathrm{O}(k|\mathcal{P}| \cdot (3|\mathcal{S}|)^{4k+4}|\mathcal{P}|)$ time. From the choice of k and Lemma 9.1.17, this is a $(1 + \epsilon)$-approximation. The theorem follows immediately. □

9.1.1 Geometric Budgeted Maximum Coverage

The above ptas easily extends to the weighted case of Geometric Set Cover, by weighting the graph constructed in the proof of Theorem 9.1.1. We can however extend to the more general budgeted case as well.

Let $\mathcal{S}$ be a set of unit squares, $\mathcal{P}$ a set of points, c a cost function over $\mathcal{S}$, p a nonnegative profit function over $\mathcal{P}$, and B a budget. Let $p_{\max}$ denote the maximum profit of any single point. We define the function $cov(s)$ as the set of points in $\mathcal{P}$ covered by a square $s \in \mathcal{S}$. This notation extends to $cov(S)$ for a set $S \subseteq \mathcal{S}$. Abusing notation, we will use $p(S)$ to denote $p(cov(S))$.

Let $k \geq 2$ be an integer we determine later. Use slabs as before.

Theorem 9.1.19 *For any instance of Geometric Budgeted Maximum Coverage on a set of unit squares $\mathcal{S}$ where all points are inside $k-1$ consecutive height 1 horizontal slabs and all profits are positive integers, one can find for all $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$ a cheapest set of profit at least r (if one exists) in $O((3|\mathcal{S}|)^{4k} (|\mathcal{P}| \cdot p_{\max})^2)$ time.*

Proof: We modify the algorithm described above. Assume the cost of squares in $\mathcal{S}^l \cup \mathcal{S}^r$ to be zero. Remove the coverage constraints from the four definitions of insertable. Then, as in the proof of Theorem 9.1.1, we construct a directed graph G with $V(G)$ equal to the set of all fronts and an edge from F to F' if F' can be obtained from F by a single valid insertion.

Alter this graph G as follows. For any edge in $E(G)$ from some front F to a front F', we replace the edge by a path. The number of edges of the path is equal to the total profit of the points covered by the insertion. For example, for an upper insertion of a square s intersecting line i, this is the total profit of the points covered by u_i or l_{i+1} in $[x_i, x'_i] \times [i, i+1]$. The cost of inserting s is modeled by assigning a weight of $c(s)$ to the first edge of the path and assigning weight 0 to all other edges.

Now the number of edges on a F^l–F^r path minus $k+1$ is equal to the profit of the solution corresponding to this path. Its cost is equal to the weight of the path. Hence we aim to find for any $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$ a lightest path of length at least r. A straightforward dynamic programming algorithm for this problem takes $O(|E(G)| \cdot |\mathcal{P}| \cdot p_{\max}) = O((3|\mathcal{S}|)^{4k}(|\mathcal{P}| \cdot p_{\max})^2)$ time. □

By slightly changing the dynamic programming algorithm of Theorem 9.1.19, we can also deal with points of zero profit.

We now apply the shifting technique and scaling to obtain a ptas. Start by assuming integer profits. For each integer $0 \leq a \leq k-1$, let N_a denote the set of points between lines $y = bk + a$ and $y = bk + a + 1$ for any $b \in \mathbb{Z}$. Moreover, for any $b \in \mathbb{Z}$, let $\mathcal{P}_a^b$ be the set of points between lines $y = bk + a + 1$ and $y = (b+1)k + a$.

For any $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$, let $C_a^b(r)$ denote the set returned by the algorithm of Theorem 9.1.19, applied on $\mathcal{S}$ and $\mathcal{P}_a^b$, attaining profit at least r. We assume that $c(C_a^b(r)) = \infty$ if $C_a^b(r)$ profit at least r cannot be attained.

Let the nonempty sets $\mathcal{P}_a^b$ are numbered arbitrarily $\mathcal{P}_a^0, \ldots, \mathcal{P}_a^{l_a}$, and let $C_a^0, \ldots, C_a^{l_a}$ be the corresponding solutions. Define

$$\begin{aligned} \mathsf{s}_a(0,r) &= c(C_a^0(r)) \\ \mathsf{s}_a(b,r) &= \min_{0 \le r' \le r} \{c(C_a^b(r')) + s_a(b-1, r-r')\} \end{aligned}$$

for $1 \le b \le l_a$ and $0 \le r \le |\mathcal{P}| \cdot p_{\max}$. Observe that computing s_a takes $\mathrm{O}(|\mathcal{P}| \cdot (|\mathcal{P}| \cdot p_{\max})^2)$ time.

Let C_a denote a set attaining $\max_{0 \le r \le |\mathcal{P}| \cdot p_{\max}} \{r \mid \mathsf{s}_a(l_a, r) \le B\}$ and let $C_{\max}$ denote a most profitable such set. By definition, $c(C_{\max}) \le B$.

Lemma 9.1.20 *$p(C_{\max}) \ge (1 - 1/k) \cdot p(OPT)$, where OPT is an optimal solution.*

Proof: Let $\mathcal{S}_a^b$ denote the set of squares in $\mathcal{S}$ covering at least one point in $\mathcal{P}_a^b$. Then it can be easily seen that

$$c\left(C_a^b\left(p\left(cov\left(OPT \cap \mathcal{S}_a^b\right) \cap \mathcal{P}_a^b\right)\right)\right) \le c(OPT \cap \mathcal{S}_a^b)$$

for any $0 \le a \le k-1$ and $0 \le b \le l_a$. Because for fixed a the sets $\mathcal{S}_a^b$ are pairwise disjoint, $\sum_{b=0}^{l_a} c(OPT \cap \mathcal{S}_a^b) \le B$. Then it follows from the definition of s and by induction that

$$p(C_a) \ge \sum_{b=0}^{l_a} p(cov(OPT \cap \mathcal{S}_a^b) \cap \mathcal{P}_a^b).$$

Since

$$\sum_{b=0}^{l_a} p(cov(OPT \cap \mathcal{S}_a^b) \cap \mathcal{P}_a^b) = p(OPT) - p(cov(OPT) \cap N_a)$$

and any point is in N_a for precisely one value of a,

$$\begin{aligned} k \cdot p(C_{\max}) &\ge \sum_{a=0}^{k-1} p(C_a) \\ &\ge \sum_{a=0}^{k-1} \Big(p(OPT) - p(cov(OPT) \cap N_a)\Big) \\ &= (k-1) \cdot p(OPT). \end{aligned}$$

Hence $p(C_{\max}) \ge (1 - 1/k) \cdot p(OPT)$. □

Theorem 9.1.21 *There is a ptas for Geometric Budgeted Maximum Coverage on unit squares.*

Proof: Consider some $\epsilon > 0$ and let $k = \max\{2, \lceil 1/\epsilon \rceil\}$. To deal with non-integer profits and to achieve polynomial running time, we first scale the profits. Define the integer profit function p' by $p'(u) = \left\lfloor \frac{|\mathcal{P}| \cdot p(u)}{\epsilon \cdot p_{\max}} \right\rfloor$ for any $u \in \mathcal{P}$. Now apply the above algorithm with p' and compute $C_{\max}$. Following Lemma 9.1.20, $p'(C_{\max}) \geq (1 - 1/k) \cdot p'(OPT)$, where OPT is an optimal solution with profit function p. Hence, as $p(OPT) \geq p_{\max}$,

$$\begin{aligned} p(C_{\max}) &\geq \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \cdot p'(C_{\max}) \\ &\geq (1 - 1/k) \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \cdot p'(OPT) \\ &\geq (1 - 1/k) \cdot \left(p(OPT) - |\mathcal{P}| \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \right) \\ &\geq (1 - 1/k) \cdot (1 - \epsilon) \cdot p(OPT) \\ &\geq (1 - \epsilon)^2 \cdot p(OPT). \end{aligned}$$

The running time is

$$\mathrm{O}(k|\mathcal{P}| \cdot (3|\mathcal{S}|)^{4k} \, |\mathcal{P}|^4/\epsilon^2 + k|\mathcal{P}| \cdot |\mathcal{P}|^4/\epsilon^2) = \mathrm{O}(k|\mathcal{P}| \cdot (3|\mathcal{S}|)^{4k} \, |\mathcal{P}|^4/\epsilon^2),$$

because $p'_{\max} \leq |\mathcal{P}|/\epsilon$. This gives the ptas. □

9.1.2 Optimality and Relation to Domination

Geometric Set Cover and the geometric version of Minimum Dominating Set are closely related. We exploit this relation here to give a ptas for Minimum-Weight Dominating Set on unit square graphs and to show that the algorithms for (the budgeted version of) Geometric Set Cover developed above are essentially optimal.

Observe that two squares of side length 1 centered on points p and p' intersect if and only if p is contained in the square of side length 2 centered on p' and p' is contained in the square of side length 2 centered on p. Hence given a collection of unit squares $\mathcal{S}$, the minimum dominating set problem on $G[\mathcal{S}]$ is equivalent to the Geometric Set Cover problem on $\mathcal{S}'$ and the centers of $\mathcal{S}'$ as point set, where $\mathcal{S}'$ is obtained from $\mathcal{S}$ using the preceding observation (see also Mihalák [209]). Then the following result is immediate from Theorem 9.1.21.

Theorem 9.1.22 *There is a ptas for Minimum-Weight Dominating Set on unit square graphs.*

Recall from the discussion of Section 6.3.5 that the techniques developed earlier were not sufficient to give a ptas for the weighted case of Minimum Dominating Set. Theorem 9.1.22 therefore is the first ptas for Minimum-Weight Dominating Set on intersection graphs of two-dimensional objects.

Another consequence of the above reduction from Minimum Dominating Set on unit square graphs to Geometric Set Cover on unit squares is the following. Recall from Section 6.4 that the exponential time hypothesis states that n-variable 3SAT cannot be decided in $2^{o(n)}$ time.

Theorem 9.1.23 *If there exist constants $\delta \geq 1$, $0 < \beta < 1$ such that Geometric Set Cover or Geometric Budgeted Maximum Coverage on unit squares of density d have 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.*

This is immediate from Theorem 6.4.4. Note that the algorithms of Theorem 9.1.18 and 9.1.21 are optimal for dense instances, but might still be slightly improved for nondense instances. It seems though that the analysis of the above algorithms does not improve when assuming bounded density. We believe however that some small changes to the algorithm are sufficient to make it optimal in this sense.

Similarly, we can show from Theorem 6.4.9 that Geometric Set Cover and Geometric Budgeted Maximum Coverage on unit squares have no eptas.

Theorem 9.1.24 *Geometric Set Cover and Geometric Budgeted Maximum Coverage on n unit squares of density $d = d(n) = \Omega(n^{\alpha})$ for some constant $0 < \alpha \leq 1$ cannot have an eptas, unless FPT=W[1].*

This again gives an indication that there is little chance to improve on the ptas of Theorem 9.1.22.

Finally, we note that Theorem 9.1.23 and Theorem 9.1.24 also hold (mutatis mutandis) on unit disks.

9.2 Hardness of Approximation

Not much is clear yet about the approximability of Geometric Set Cover. The approximation scheme of the previous section all but settled its approximability on unit squares. This gives hope for the existence of a ptas on unit disks. For more general objects however, we know almost nothing. In this section, we give several hardness results, showing that Geometric Set Cover is as hard as Minimum Set Cover in some cases and APX-hard in others.

Theorem 9.2.1 *Geometric Set Cover is not approximable within $(1-\epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{O(\log\log n)}$), on the following objects:*

- *convex polygons,*
- *translated copies of a single polygon,*
- *rotated copies of a single convex polygon,*
- *α-fat objects for any $\alpha > 1$.*

Geometric Set Cover is APX-hard on the following objects:

- *convex polygons with r corners, where $r \geq 4$,*
- *α-fat objects of constant description complexity for any $\alpha > 1$,*
- *rectangles,*
- *ellipses.*

Proof: The reductions are essentially the same as those in Section 8.5. The gadgets proposed there need only be slightly modified. In short, each object that we used to model an element $u \in \mathbb{U}$ for the universe $\mathbb{U}$ of the minimum set cover instance that we are reducing from, will be a point instead. For rectangles and ellipses, we use the gadget of the proof of Theorem 8.5.8, but any small plate or $S_{i,j}$ will be a point instead of a rectangle.

The $\ln n$-hardness on rotated copies of a single polygon follows by applying the same ideas as in Theorem 8.5.3, but on slices of a single large disk. Rotations of the disks then allow for the same kind of construction. □

The $\ln n$-hardness on translated copies of a single polygon seems somewhat at odds with Laue's [188] constant-factor approximation algorithm on translated copies of a fixed three-dimensional polytope. However, the complexity of the polygon used in the hardness result depends on the number of polygons (i.e. on the number of sets in the minimum set cover instance). Hence the polygon may not been assumed to be fixed. When looking closely at Laue's result, one can observe that the approximation ratio of his algorithm actually depends (linearly) on the complexity of the polytope if it is nonconvex.

Using an idea by Khuller, Moss, and Naor [165] to reduce from Minimum Set Cover to Budgeted Maximum Coverage, we can also prove hardness results for Geometric Budgeted Maximum Coverage.

Theorem 9.2.2 *Geometric Budgeted Maximum Coverage is not approximable within ratio better than $(1 - 1/e)$, unless NP $\subset$ DTIME($n^{O(\log \log n)}$), on the following objects:*

- *convex polygons,*
- *translated copies of a single polygon,*
- *rotated copies of a single convex polygon,*
- *α-fat objects for any $\alpha > 1$.*

Note that the APX-hardness results of Theorem 9.2.1 cannot be transferred to Geometric Budgeted Geometric Coverage using this trick. The underlying reductions are not from Minimum Set Cover, but from Minimum k-Set Cover and Minimum Vertex Cover, for which the idea of Khuller, Moss, and Naor does not appear to work.

Chapter 10

Geometric Unique and Membership Coverage Problems

We present the first study on the approximability of geometric versions of the unique coverage problem and the minimum membership set cover problem. We prove that Unique Coverage (and thus Budgeted Low-Coverage as well) remains NP-hard on unit disks and give constant-factor approximation algorithms for both problems on unit disks. The results extend to unit squares. We then show that Budgeted Low-Coverage has an eptas on fat objects of bounded ply, but prove that without the bounded ply assumption, the problem is as hard to approximate as in its general setting.

We then consider the geometric version of Minimum Membership Set Cover. We prove that approximating the problem within ratio less than 2 is NP-hard on unit disks and unit squares, and give a 5-approximation algorithm on unit squares that runs in polynomial time if the optimal objective value is bounded by a constant.

10.1 Unique Coverage

We consider the complexity and approximability of Unique Coverage on unit disks and on unit squares. We first introduce some notation. Suppose that we are given a set $\mathcal{S}$ of geometric objects in the plane and a set of points $\mathcal{P} \subseteq \mathbb{R}^2$. Then for any $X \subseteq \mathcal{S}$, $cov(X)$ respectively $uc(X)$ denotes the set of points in $\mathcal{P}$ covered respectively covered uniquely by X. If S is a square in the plane, $uc_S(X)$ is defined as the set of points of $\mathcal{P}$ that lie inside S and are uniquely covered by X.

Theorem 10.1.1 *Unique Coverage on unit disks and unit squares is NP-hard.*

Proof: We reduce from Independent Set on planar graphs of maximum degree 3, which is known to be NP-hard [114]. The starting point of the construction is similar to one used to prove NP-hardness of Independent Set on unit disk graphs [67]. For a planar graph G of degree 3, create a rectilinear embedding of G. This is an embedding of G onto a unit grid, such that each vertex is mapped to a unique corner of the grid and each edge is mapped to a path in the grid, where all such paths are disjoint, except possibly at their ends.

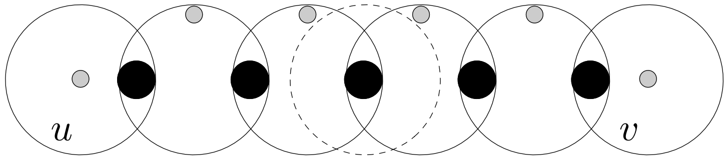

Figure 10.1: The edge gadget for some $(u, v) \in E(G)$, shown together with the vertex disks for u and v. If the path in the rectilinear embedding corresponding to (u, v) consists of more than one straight line segment, we can easily adapt the gadget.

Valiant [255, Theorem 2] has shown that such an embedding exists having area $O(|V(G)|^2)$. Now replace each vertex by a disk with a point in its center and each edge by a gadget, shown in Figure 10.1.

The set of (solid) disks connecting u and v for any edge $(u, v) \in E(G)$ must have even cardinality and is denoted by $\mathcal{D}^{(u,v)}$. A single (dashed) disk contains the middle black blob. The small gray points correspond to a single point. All disks of the construction, except the middle disk of each edge gadget (drawn dashed in Figure 10.1), contain a unique gray point. Each big black blob corresponds to a collection of $t = |V(G)| + \sum_{(u,v) \in E(G)} |\mathcal{D}^{(u,v)}|$ points.

We claim that in the constructed instance of Unique Coverage, a set of unit disks can uniquely cover at least

$$k' := k + \sum_{(u,v) \in E(G)} \left(\tfrac{1}{2} |\mathcal{D}^{(u,v)}| + t \cdot \left(|\mathcal{D}^{(u,v)}| + 1 \right) \right)$$

points if and only if G has an independent set of cardinality at least k. For the 'only if' part, note that any set of disks uniquely covering at least k' points must uniquely cover all black blobs. This can only be done using exactly $\frac{1}{2}|\mathcal{D}^{(u,v)}|$ disks of each edge gadget (which uniquely cover $\frac{1}{2} \sum_{(u,v) \in E(G)} |\mathcal{D}^{(u,v)}|$ gray points in total) and at least k vertex disks. By the construction of the edge gadgets, these vertex disks form an independent set, which indeed has cardinality at least k. We can now easily verify the 'if' part of the claim.

A similar argument can be given to demonstrate NP-hardness on unit squares. Moreover, the ply of the above construction is 3 and thus the NP-hardness extends to this case. □

10.1.1 Approximation Algorithm on Unit Disks

Let $\mathcal{D}$ be a set of equally sized disks and $\mathcal{P}$ a set of points, both in $\mathbb{R}^2$. By scaling, we may assume that all disks in $\mathcal{D}$ have radius 1/2. We aim to find a set $C \subseteq \mathcal{D}$ maximizing the number of uniquely covered points of $\mathcal{P}$. We apply the shifting technique (see Chapter 6) in a novel way.

Lemma 10.1.2 *Suppose for points in a square of size $\delta \times \delta$ ($0 < \delta \leq 1$), the unique coverage problem on unit disks has a polynomial time $1/c$-approximation algorithm. Then there is a polynomial time $(1/c) \cdot \delta^2/(1+\delta)^2$-approximation algorithm for the general unique coverage problem on unit disks.*

Proof: Let OPT be an optimal solution to the unique coverage problem on unit disks for some set of disks $\mathcal{D}$ and set of points $\mathcal{P}$. Pick two numbers a, b uniformly at random from $[0, 1+\delta)$. Consider the set of squares $\mathcal{S} = \{[a + h + h\delta, a + h + (h+1)\delta) \times [b + v + v\delta, b + v + (v+1)\delta) \mid h, v \in \mathbb{Z}\}$. Each square has size $\delta \times \delta$. As the squares in $\mathcal{S}$ have pairwise distance greater than 1, no unit disk can cover a point in more than one square of this set. Hence we may consider these squares to be 'independent'.

The probability that a point of $\mathcal{P}$ is in a square of $\mathcal{S}$ is $\delta^2/(1+\delta)^2$. Hence

$$\mathbf{E}\left[\textstyle\sum_{S \in \mathcal{S}} |uc_S(OPT)|\right] = |uc(OPT)| \cdot \delta^2/(1+\delta)^2.$$

By assumption, we can find a $1/c$-approximation of the unique coverage problem on unit disks for each of the squares in $\mathcal{S}$ in polynomial time. Let $\mathsf{size}(S)$ denote the number of points uniquely covered by the solution of the algorithm for a particular square $S \in \mathcal{S}$. Then $\mathsf{size}(S) \geq 1/c \cdot |uc_S(OPT)|$. As we can assume that the solutions produced by the $1/c$-approximation algorithm contain only disks intersecting the square it was invoked on, it follows that

$$\begin{aligned} \mathbf{E}\left[\textstyle\sum_{S \in \mathcal{S}} \mathsf{size}(S)\right] &\geq (1/c) \cdot \mathbf{E}\left[\textstyle\sum_{S \in \mathcal{S}} |uc_S(OPT)|\right] \\ &= (1/c) \cdot |uc(OPT)| \cdot \delta^2/(1+\delta)^2. \end{aligned}$$

This approach can be derandomized. Choices of a, b for which the same set of points is in the squares of $\mathcal{S}$ give an approximation of the same quality. Hence it suffices to look at the $\mathrm{O}(|\mathcal{P}|^2)$ values of a, b for which a square boundary hits a point of $\mathcal{P}$ and thus we can consider all relevant values in polynomial time. The solution with the highest $\sum_{S \in \mathcal{S}} \mathsf{size}(S)$ is a $(1/c) \cdot \delta^2/(1+\delta)^2$-approximation of the optimum. □

The proof of the next theorem uses some ideas from Ambühl et al. [13].

Theorem 10.1.3 *Unique Coverage on unit disks has a $1/18$-approximation algorithm running in $\mathrm{O}(|\mathcal{P}|^3 |\mathcal{D}|^8)$ time.*

Proof: We prove that there is a polynomial-time $1/2$-approximation algorithm for Unique Coverage on unit disks for size $1/2 \times 1/2$ squares. Together with Lemma 10.1.2, this proves the theorem.

Consider a size $1/2 \times 1/2$ square S containing a set of points $\mathcal{P}_S$ and intersected by a set of disks $\mathcal{D}_S$ of radius $1/2$. We may assume that no disk covers all points of $\mathcal{P}_S$, for such a disk would constitute an optimal solution. Construct a mapping of disks in $\mathcal{D}_S$ to one of the four boundaries of S. The mapping assigns a disk to the boundary of S it overlaps most, breaking ties arbitrarily. If we can solve the unique coverage problem on unit disks optimally for both pairs of opposing boundaries of S, the best solution gives a 1/2-approximation of the optimum for Unique Coverage on unit disks for S.

Let OPT_S be an optimal solution to the unique coverage problem on $\mathcal{D}_S$ and $\mathcal{P}_S$. Consider two opposing boundaries b, b' of S (say b is the top boundary, b' the bottom boundary) and the sets of disks $\mathcal{D}^b, \mathcal{D}^{b'}$ assigned to them. Observe that for any disk $d \in \mathcal{D}^b$ the projection of $d \cap S$ onto b is equal to $d \cap b$, or d would overlap another boundary of S more and the mapping would assign it to this boundary. Hence any disk in $\mathcal{D}^b \cap OPT_S$ meets the lower envelope of $\mathcal{D}^b \cap OPT_S$ (for a set of disks, we say that a disk d from the set *meets* or *is on* the lower envelope of the set if part of that envelope is formed by the boundary of d), or any point of $\mathcal{P}_S$ it covers would already be covered by a disk on the lower envelope. Furthermore, for any disk d_2 on the lower envelope of $\mathcal{D}^b \cap OPT_S$, all points of $\mathcal{P}_S$ uniquely covered by d_2 are in $d_2 - d_1 - d_3$, where d_1 (d_3) is the disk lying directly to the left (right) of d_2 on the lower envelope of $\mathcal{D}^b \cap OPT_S$. The same properties hold for disks in $\mathcal{D}^{b'}$ and the upper envelope of $\mathcal{D}^{b'} \cap OPT_S$.

We now design a sweep-line algorithm with a vertical sweep-line, moving from left to right and stopping at each point of $\mathcal{P}_S$. Index the points of $\mathcal{P}_S$ in order of nondecreasing x-coordinate, with u_1 being the leftmost and u_m the rightmost point. By the above arguments, it suffices to consider three disks of both boundaries at each position of the sweep-line. Therefore when the sweep-line is at u_i, we say that a triple $\mathsf{d} = (d_1, d_2, d_3)$ of disks in $\mathcal{D}^b \cup \{b\}$ is *proper* if

- if $d_i = d_j$ for some $1 \leq i < j \leq 3$, then $d_i = d_j = b$;
- d_1, d_2, and d_3 appear in this order on the lower envelope of $d_1 \cup d_2 \cup d_3$;
- the intersection point of the sweep-line and d_2 lies on this envelope.

We use the 'disk' b to model the boundary, i.e. the situation when no disk intersects the sweep-line. If d is proper, then a proper triple d' is a *predecessor* of d if $\mathsf{d} = \mathsf{d}'$ or $\mathsf{d} = (d'_2, d'_3, d'_4)$ for some $d'_4 \in \mathcal{D}^b \cup \{b\}$. These notions can be defined analogously for triples $\mathsf{e} = (e_1, e_2, e_3)$ with respect to b'.

Now define for any point u_i and any pair of proper triples d and e a function h such that

$$h_i(\mathsf{d}, \mathsf{e}) = \begin{cases} \iota_i(\mathsf{d}, \mathsf{e}) & \text{if } i = 1; \\ \iota_i(\mathsf{d}, \mathsf{e}) + \max h_{i-1}(\mathsf{d}', \mathsf{e}') & \text{if } i > 1, \end{cases}$$

where $\iota_i(\mathsf{d}, \mathsf{e})$ is 1 if u_i is in $d_2 - d_1 - d_3$ or in $e_2 - e_1 - e_3$ but not in both, and 0 otherwise. The maximum is over all proper triples d' and e' that are predecessors of d and e respectively.

The maximum value of h_m over all proper triples for u_m is the optimal solution for this pair of opposing boundaries of S. Because h can be computed in $O(|\mathcal{P}|\,|\mathcal{D}|^8)$ time, the theorem follows from Lemma 10.1.2. □

10.1.2 Budgets and Satisfactions

The above algorithm extends easily to the case where points have associated profits and we aim to maximize the total profit of uniquely covered points. However for Budgeted Low-Coverage on unit disks, we need to change the approach. In the following, let $p_{\max}$ and $c_{\max}$ denote the maximum profit and the maximum cost respectively. We may assume that the budget is at least the maximum cost and that an optimal solution attains profit at least $s_1 \cdot p_{\max}$.

Theorem 10.1.4 *If $s_z > 0$ and $s_{z+1} = 0$ for some fixed z and the profits and satisfactions are integers, then Budgeted Low-Coverage on unit disks has a $1/18$-approximation algorithm running in $\mathrm{O}(|\mathcal{P}|^3 \cdot (s_1 \cdot p_{\max})^2 \cdot |\mathcal{D}|^{4z+4})$ time.*

Proof: Recall the proof of Lemma 10.1.2. Let $\mathcal{S}$ be the set of squares under consideration. We use similar ideas as in Theorem 10.1.3 to compute for both pairs of opposing boundaries of each square $S \in \mathcal{S}$ and for each $r = 0, \dots, |\mathcal{P}| \cdot s_1 \cdot p_{\max}$ a set of disks such that the total satisfaction-modulated profit of covered points is at least r and the total cost is minimized.

Instead of triples, we consider tuples of $2z + 1$ disks. Similar to Theorem 10.1.3, we can define the notions of proper and predecessor tuples. We then define for any $r = 0, \dots, |\mathcal{P}| \cdot s_1 \cdot p_{\max}$, any point u_i, and any pair of proper tuples $\mathsf{d} = (d_1, \dots, d_{2z+1})$, $\mathsf{e} = (e_1, \dots, e_{2z+1})$ a function h such that $h_i(\mathsf{d}, \mathsf{e}, r)$ equals

$$\begin{cases} \min\Big\{\vartheta(d_{z+1}, d'_{z+1}) \cdot c(d_{z+1}) \\ \qquad +\vartheta(e_{z+1}, e'_{z+1}) \cdot c(e_{z+1}) \\ \qquad +h_{i-1}\Big(\mathsf{d}', \mathsf{e}', r - s_{\gamma_i(\mathsf{d},\mathsf{e})} \cdot p(u_i)\Big)\Big\} & \text{if } i > 1; \\ \sum_{i=1}^{z+1} (c(d_i) + c(e_i)) & \text{if } i = 1 \text{ and } s_{\gamma_i(\mathsf{d},\mathsf{e})} \cdot p(u_i) \geq r; \\ \infty & \text{otherwise.} \end{cases}$$

The minimum is over all proper triples d' and e' that are predecessors of d and e respectively. The value of $\gamma_i(\mathsf{d}, \mathsf{e})$ is equal to the number of disks in $\{d_1, \dots, d_{2z+1}, e_1, \dots, e_{2z+1}\}$ containing u_i. The indicator function $\vartheta(\cdot, \cdot)$ is 1 if its parameters are distinct and 0 otherwise.

We are then interested in h_m for all $r = 0, \dots, |\mathcal{P}| \cdot s_1 \cdot p_{\max}$. These values can clearly be computed in $\mathrm{O}(|\mathcal{P}|^2 \cdot s_1 \cdot p_{\max} \cdot |\mathcal{D}|^{4z+4})$ time.

For fixed S and r, let $\mathsf{cost}(S, r)$ denote the minimum cost over both pairs of opposing boundaries. Then $\max\{\sum_{S \in \mathcal{S}} r_S \mid \sum_{S \in \mathcal{S}} \mathsf{cost}(S, r_S) \leq B\}$ gives a $1/2$-approximation of $\sum_{S \in \mathcal{S}} sp_S(OPT)$, where $sp_S(OPT)$ is the satisfaction-modulated profit accrued by OPT in square S. This maximum is just an instance of Multiple-Choice Knapsack [60, 120], which we have to solve in a way that avoids having B in the running time. We first compute for each $r = 0, \dots, |\mathcal{P}| \cdot s_1 \cdot p_{\max}$ a function

$$g(r) = \min\left\{\textstyle\sum_{S \in \mathcal{S}} \mathsf{cost}(S, r_S) \mid \sum_{S \in \mathcal{S}} r_S = r\right\}$$

and then choose the largest r such that $g(r) \leq B$. Using dynamic programming, this can be done in $\mathrm{O}(|\mathcal{P}|^3 \cdot (s_1 \cdot p_{\max})^2)$ time, as the number of nonempty squares of $\mathcal{S}$ is at most $|\mathcal{P}|$. Applying the shifting technique as in Lemma 10.1.2 yields a 1/18-approximation in $\mathrm{O}(|\mathcal{P}|^3 \cdot (s_1 \cdot p_{\max})^2 \cdot |\mathcal{D}|^{4z+4})$ time. □

Theorem 10.1.5 *If $s_z > 0$ and $s_{z+1} = 0$ for some fixed z, then for any $\epsilon > 0$, there is a $\mathrm{O}(|\mathcal{P}|^7 \cdot |\mathcal{D}|^{4z+4}/\epsilon^4)$ time $(1-\epsilon)/18$-approximation algorithm for Budgeted Low-Coverage on unit disks.*

Proof: We extend the algorithm of Theorem 10.1.4 to deal with noninteger profits and satisfactions and to achieve polynomial running-time by applying scaling. Scale the profits by $\frac{|\mathcal{P}|}{\epsilon \cdot p_{\max}}$, i.e. define $p'(u) = \left\lfloor \frac{|\mathcal{P}| \cdot p(u)}{\epsilon \cdot p_{\max}} \right\rfloor$, and scale the satisfactions by $\frac{|\mathcal{P}|}{\epsilon \cdot s_1}$, i.e. define $s'_i = \left\lfloor \frac{|\mathcal{P}| \cdot s_i}{\epsilon \cdot s_1} \right\rfloor$. We first give an auxiliary inequality. For any $u \in \mathcal{P}$ and any i,

$$\begin{aligned}
& s_i \cdot p(u) - \left(\frac{\epsilon \cdot s_1}{|\mathcal{P}|} \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \right) \cdot s'_i \cdot p'(u) \\
&= s_i \cdot p(u) - \frac{\epsilon \cdot s_1}{|\mathcal{P}|} \cdot \left\lfloor \frac{|\mathcal{P}| \cdot s_i}{\epsilon \cdot s_1} \right\rfloor \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \cdot \left\lfloor \frac{|\mathcal{P}| \cdot p(u)}{\epsilon \cdot p_{\max}} \right\rfloor \\
&\leq s_i \cdot p(u) - \frac{\epsilon \cdot s_1}{|\mathcal{P}|} \cdot \left(\frac{|\mathcal{P}| \cdot s_i}{\epsilon \cdot s_1} - 1 \right) \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \cdot \left(\frac{|\mathcal{P}| \cdot p(u)}{\epsilon \cdot p_{\max}} - 1 \right) \\
&= s_i \cdot p(u) - \left(s_i - \frac{\epsilon \cdot s_1}{|\mathcal{P}|} \right) \cdot \left(p(u) - \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \right) \\
&\leq 2 \cdot \frac{\epsilon \cdot s_1 \cdot p_{\max}}{|\mathcal{P}|}.
\end{aligned}$$

Now for a set of disks $\mathcal{D}'$, let $sp(\mathcal{D}')$ denote the satisfaction-modulated profit achieved by $\mathcal{D}'$ under the original profits and satisfaction and let $s'p'(\mathcal{D}')$ be the satisfaction-modulated profit achieved by $\mathcal{D}'$ under the scaled profits and scaled satisfactions.

Apply the algorithm of Theorem 10.1.4 to obtain a set of disks C that is a 1/18-approximation of the optimum under the scaled profits and satisfactions. Let OPT denote a set of disks giving an optimum satisfaction-modulated profit under the budget and the original profits and satisfactions. Then

$$\begin{aligned}
sp(C) &\geq \frac{\epsilon \cdot s_1}{|\mathcal{P}|} \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \cdot s'p'(C) \\
&\geq \frac{\epsilon \cdot s_1}{|\mathcal{P}|} \cdot \frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|} \cdot \tfrac{1}{18} \cdot s'p'(OPT) \\
&\geq \tfrac{1}{18} \cdot \left(sp(OPT) - 2 \cdot \frac{\epsilon \cdot s_1 \cdot p_{\max}}{|\mathcal{P}|} \cdot |\mathcal{P}| \right) \\
&\geq \tfrac{1}{18} \cdot \left(sp(OPT) - 2 \cdot \epsilon \cdot sp(OPT) \right).
\end{aligned}$$

Note that $s'_1, p'_{\max} \leq |\mathcal{P}|/\epsilon$. The theorem follows from Theorem 10.1.4. □

If the satisfactions are different for each point, but all still nonincreasing, a similar algorithm may be used. It remains an interesting open problem if Budgeted Low-Coverage on unit disks is approximable for arbitrary z and/or satisfactions that are not nonincreasing.

10.1.3 Approximation Algorithm on Unit Squares

As with the geometric set cover problem, it seems that Unique Coverage is easier to approximate on unit squares than on unit disks. On unit disks, we had to consider instances on size $1/2 \times 1/2$ squares to be able to restrict the attention to the upper and lower envelopes. On unit squares however, this is no longer necessary. We use some ideas from Mihalák [209].

Theorem 10.1.6 *If* $s_z > 0$ *and* $s_{z+1} = 0$ *for some fixed* z*, then for any* $\epsilon > 0$*, there is a* $\mathrm{O}(|\mathcal{P}|^7 \cdot |\mathcal{D}|^{4z+4}/\epsilon^4)$*-time* $(1-\epsilon)/2$*-approximation algorithm for Budgeted Low-Coverage on unit squares.*

Proof: First assume that the profits and satisfactions are integer. Partition the plane into horizontal slabs of height 1. This induces a partition of the points of $\mathcal{P}$ as well. We claim that for any slab S and for each $r = 0, \ldots, |\mathcal{P}| \cdot s_1 \cdot p_{\max}$, we can find a set of squares such that the total satisfaction-modulated profit of covered points in S is at least r and the total cost is minimized.

To see this, consider a slab S and the set of points $\mathcal{P}_S$ contained in it. Any square covering a point of $\mathcal{P}_S$ must intersect either the top or the bottom boundary of S. Let $\mathcal{D}^t$ denote the set of squares intersecting the top boundary of S and $\mathcal{D}^b$ the set of squares intersecting the bottom boundary. Observe that in any optimal solution OPT, any square in $OPT \cap \mathcal{D}^t$ must intersect the lower envelope of $OPT \cap D^t$. Moreover, all points yielding a profit that are covered by a square $d_{z+1} \in OPT \cap \mathcal{D}^t$ are in $\bigcup_{i \in \{2,\ldots,2z\}} d_i - d_1 - d_{2z+1}$, where $d_1, \ldots, d_z$ and $d_{z+2}, \ldots, d_{2z+1}$ are the squares respectively to the left and to the right of d_{z+1} on the lower envelope of $OPT \cap \mathcal{D}^t$. Similar observations can be made about $C \cap \mathcal{D}^b$ and its upper envelope. Then we can just apply the algorithm of Theorem 10.1.3 and Theorem 10.1.4. This takes $\mathrm{O}(|\mathcal{P}|^2 \cdot s_1 \cdot p_{\max} \cdot |\mathcal{D}|^{4z+4})$ time.

Now let $\mathcal{S}^1$ denote the set of slabs whose bottom boundary is $y = i$ for some even integer i and let $\mathcal{S}^2$ denote the set of remaining slabs. Use the Multiple-Choice Knapsack algorithm of Theorem 10.1.4 to compute for both $\mathcal{S}^1$ and $\mathcal{S}^2$ a set of squares such that the total satisfaction-modulated profit of covered points is maximal and the total cost is at most B. The most profitable of these two solutions then gives a 1/2-approximation algorithm. This takes $\mathrm{O}(|\mathcal{P}|^3 \cdot (s_1 \cdot p_{\max})^2)$ time, as the number of nonempty slabs is at most $|\mathcal{P}|$.

Finally, we apply scaling as in Theorem 10.1.5 to deal with noninteger profits and satisfactions and to obtain a polynomial running time. □

10.2 Unique Coverage on Disks of Bounded Ply

If the disks have arbitrary size, but have bounded ply, we can improve on the algorithm of the previous section and give an eptas for Unique Coverage on disks. Recall that the *ply* of a set of disks is the maximum over all points in the plane of the number of disks strictly containing this point. To obtain the eptas, we apply the shifting technique, which we used before in the case of disks of bounded ply (see Chapter 7). However, the complexity of the budgeted unique coverage problem on disks forces major changes in this approach. In particular, it is nontrivial to enforce the global budget constraint; we handle this by creating dynamic programming tables that are additionally indexed by profit values and contain the cheapest cost for achieving a certain profit in a given square. The best choice of profit values for disjoint squares can then be addressed as a multiple-choice knapsack problem. Furthermore, relating the profit of the algorithm's solution to that of a modified optimal solution (Lemma 10.2.5) is significantly more difficult than for the problems studied in Chapter 7. In the proofs below we focus on those aspects that are different from the approach in Chapter 7.

The setup of the algorithm is as follows. Let $\mathcal{D}$ be a set of disks and $\mathcal{P}$ a set of points. By scaling, we may assume that all disks have radius at least 1/2. Partition the disks into levels, where a disk with radius r has level $j \in \mathbb{Z}_{\geq 0}$ if $2^{j-1} \leq r < 2^j$. Then we can define $\mathcal{D}_{=j}$ as the set of disks having level exactly j, $\mathcal{D}_{\geq j}$ as the set of disks having level at least j, etc. We use l to denote the level of the largest disk.

For each level j, partition the plane using a grid induced by horizontal lines $y = hk2^j$ and vertical lines $x = vk2^j$ $(h, v \in \mathbb{Z})$, where $k \geq 5$ is an odd integer determined later. The squares of this partition for level j are called *j-squares*. Any j-square is contained in precisely one $j+1$-square, while each $j+1$-square contains exactly four j-squares, denoted $S_1, \ldots, S_4$ and called *siblings*. For a j-square S, let $\mathcal{D}^S$ denote the set of disks intersecting S and let $\mathcal{D}^{\mathrm{b}(S)}$ denote the set of disks intersecting the boundary of S. As a shorthand, let $\mathcal{D}^{\mathrm{i}(S)} = \mathcal{D}^S - \mathcal{D}^{\mathrm{b}(S)}$ be the set of disks fully contained inside S and $\mathcal{D}^{+(S)} = \bigcup_{i=1}^4 \mathcal{D}^{\mathrm{b}(S_i)} - \mathcal{D}^{\mathrm{b}(S)}$ be the set of disks intersecting the boundary of some S_i, but not the boundary of S. Combinations such as $\mathcal{D}^{\mathrm{b}(S)}_{>j}$ are self-explanatory. Let $\mathcal{P}_S$ denote the set of points contained in a j-square S and let $j(S)$ denote the level of a square S.

We prove the following auxiliary result.

Theorem 10.2.1 *Let $\mathcal{D}$ be a set of disks of ply γ, $k \geq 5$ an odd integer, and OPT a subset of $\mathcal{D}$ such that $p(uc(OPT))$ is maximum under $c(OPT) \leq B$. Then in time $\mathrm{O}((k^2|\mathcal{D}| + |\mathcal{P}|) \cdot k\,|\mathcal{D}| \cdot 2^{32k\gamma/\pi}(|\mathcal{P}| \cdot p_{\max})^2)$, we can find a set $\mathcal{D}' \subseteq \mathcal{D}$ such that $c(\mathcal{D}') \leq B$ and $p(uc(\mathcal{D}')) \geq p(uc(\bigcup_S OPT^{\mathrm{i}(S)}_{=j(S)}))$.*

To prove this theorem, we apply dynamic programming on the squares. Define for each j-square S, each set $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$, and each $r \in \{0, \ldots, |\mathcal{P}| \cdot p_{\max}\}$ the function $\mathsf{cost}(S, W, r)$ as

$$\mathsf{cost}(S,W,r) = \begin{cases} \min\left\{ c(T) \,\middle|\, p(uc_S(T \cup W)) \geq r;\ T \subseteq \mathcal{D}^{\mathrm{i}(S)}_{\geq j} \right\} & \text{if } j = 0; \\ \min\Big\{ c(U) + \sum_{i=1}^{4} \mathsf{cost}\big(S_i, (U \cup W)^{\mathrm{b}(S_i)}, r_i\big) \,\Big| \\ \qquad\qquad \sum_{i=1}^{4} r_i = r;\ U \subseteq \mathcal{D}^{+(S)}_{>j-1} \Big\} & \text{if } j > 0. \end{cases}$$

Here the minimum over an empty set is ∞. Let $\mathsf{sol}(S, W, r)$ denote the subset of $\mathcal{D}$ attaining $\mathsf{cost}(S, W, r)$ if $\mathsf{cost}(S, W, r) \neq \infty$, or $\emptyset$ otherwise. Note that we would actually only need to define $\mathsf{cost}(S, W, r)$ and $\mathsf{sol}(S, W, r)$ for subsets W of $\mathcal{D}^{\mathrm{b}(S)}_{>j} \cap \bigcup_{S' \supset S} \mathcal{D}^{\mathrm{i}(S')}_{=j(S')}$. This will turn up in the analysis of the approximation ratio of the algorithm, but it is not very important in the analysis of the (worst-case) running time.

10.2.1 Properties of the cost- and sol-Functions

We start by proving the function cost is indeed close to the optimum. Define

$$up(S) = \bigcup_{S' \supset S} OPT^{\mathrm{i}(S')}_{=j(S')} \qquad \text{and} \qquad down(S) = \bigcup_{S' \subseteq S} OPT^{\mathrm{i}(S')}_{=j(S')}.$$

We give some properties of *up* and *down* that will be auxiliary to later lemmas.

Proposition 10.2.2 *For any j-square S and any child square S_i of S,*

$$\left(OPT^{+(S)}_{>j-1} \cup (up(S))^{\mathrm{b}(S)} \right)^{\mathrm{b}(S_i)} = (up(S_i))^{\mathrm{b}(S_i)}.$$

Proof: Observe that

$$\begin{aligned} &\left(OPT^{+(S)}_{>j-1} \cup (up(S))^{\mathrm{b}(S)} \right)^{\mathrm{b}(S_i)} \\ &\qquad = \left(OPT^{+(S)}_{>j-1} \cup \left(\bigcup_{S' \supset S} OPT^{\mathrm{i}(S')}_{=j(S')} \right)^{\mathrm{b}(S)} \right)^{\mathrm{b}(S_i)} \\ &\qquad = \left(OPT^{+(S)}_{=j} \cup \bigcup_{S' \supset S} OPT^{\mathrm{i}(S')}_{=j(S')} \right)^{\mathrm{b}(S_i)} \\ &\qquad = \left(\bigcup_{S'_i \supset S_i} OPT^{\mathrm{i}(S'_i)}_{=j(S'_i)} \right)^{\mathrm{b}(S_i)} \\ &\qquad = (up(S_i))^{\mathrm{b}(S_i)}. \end{aligned}$$

The proposition follows. □

Corollary 10.2.3 *It holds that* $\bigcup_{i=1}^{4} (up(S_i))^{\mathrm{b}(S_i)} = (up(S))^{\mathrm{b}(S)} \cup OPT^{+(S)}_{>j-1}$ *for any j-square S.*

Proposition 10.2.4 *For any j-square S,*

$$OPT^{+(S)}_{>j-1} \cup \bigcup_{i=1}^{4} \left(OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i) \right) = OPT^{\mathrm{i}(S)}_{>j} \cup down(S).$$

Proof: Observe that

$$\begin{aligned} & OPT^{+(S)}_{>j-1} \cup \bigcup_{i=1}^{4} \left(OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i) \right) \\ = & \left(OPT^{+(S)}_{>j} \cup \bigcup_{i=1}^{4} OPT^{\mathrm{i}(S_i)}_{>j} \right) \cup \bigcup_{i=1}^{4} down(S_i) \\ & \cup \left(OPT^{+(S)}_{=j} \cup \bigcup_{i=1}^{4} OPT^{\mathrm{i}(S_i)}_{=j} \right) \\ = & OPT^{\mathrm{i}(S)}_{>j} \cup \bigcup_{i=1}^{4} down(S_i) \cup OPT^{\mathrm{i}(S)}_{=j} \\ = & OPT^{\mathrm{i}(S)}_{>j} \cup down(S). \end{aligned}$$

The proposition follows. □

We are now ready to prove that cost is close to optimal.

Lemma 10.2.5 *It holds that*

$$\max \left\{ \sum_{S;\, j(S)=l} r_S \;\middle|\; \sum_{S;\, j(S)=l} \mathsf{cost}(S, \emptyset, r_S) \leq B \right\} \geq p \left(uc \left(\bigcup_S OPT^{\mathrm{i}(S)}_{=j(S)} \right) \right).$$

Proof: To prove the lemma, we claim that for any j-square S

$$\begin{aligned} & \mathsf{cost}\left(S, (up(S))^{\mathrm{b}(S)}, p\left(uc_S \left((up(S))^{\mathrm{b}(S)} \cup OPT^{\mathrm{i}(S)}_{>j} \cup down(S) \right) \right) \right) \\ & \qquad \leq \quad c\left(OPT^{\mathrm{i}(S)}_{>j} \right) + \sum_{S' \subseteq S} c\left(OPT^{\mathrm{i}(S')}_{=j(S')} \right). \end{aligned}$$

The intuition behind this formula is that if we consider the set W the optimum uses and the profit r attained by the optimum for S, the cost attained by $\mathsf{cost}(S, W, r)$ is at most the cost needed by the optimum. We prove it inductively on j. For $j = 0$, it is easily verified. Consider some $j > 0$ and assume inductively that the above statement holds for any $j' < j$. Observe that

$$\bigcup_{i=1}^{4} uc_{S_i} \left((up(S_i))^{\mathrm{b}(S_i)} \right) = uc_S \left(\bigcup_{i=1}^{4} (up(S_i))^{\mathrm{b}(S_i)} \right).$$

As $\mathcal{P}_{S_1}, \ldots, \mathcal{P}_{S_4}$ are pairwise disjoint sets, it follows from Corollary 10.2.3 and Proposition 10.2.4 that

$$\begin{aligned} & \sum_{i=1}^{4} p\left(uc_{S_i} \left((up(S_i))^{\mathrm{b}(S_i)} \cup OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i) \right) \right) \\ = \; & p\left(uc_S \left(\bigcup_{i=1}^{4} \left((up(S_i))^{\mathrm{b}(S_i)} \cup OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i) \right) \right) \right) \\ = \; & p\left(uc_S \left((up(S))^{\mathrm{b}(S)} \cup OPT^{+(S)}_{>j-1} \cup \bigcup_{i=1}^{4} \left(OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i) \right) \right) \right) \\ = \; & p\left(uc_S \left((up(S))^{\mathrm{b}(S)} \cup OPT^{\mathrm{i}(S)}_{>j} \cup down(S) \right) \right). \end{aligned}$$

Then by induction, the definition of cost, and Proposition 10.2.2,

$$\begin{aligned}
&\mathsf{cost}\left(S, (up(S))^{\mathrm{b}(S)}, p\left(uc_S\left((up(S))^{\mathrm{b}(S)} \cup OPT^{\mathrm{i}(S)}_{>j} \cup down(S)\right)\right)\right) \\
&\quad\leq\ c\Big(OPT^{+(S)}_{>j-1}\Big) \\
&\qquad + \sum_{i=1}^{4} \mathsf{cost}\Bigg(S_i, \left(OPT^{+(S)}_{>j-1} \cup (up(S))^{\mathrm{b}(S)}\right)^{\mathrm{b}(S_i)}, \\
&\qquad\qquad p\Big(uc_{S_i}\Big((up(S_i))^{\mathrm{b}(S_i)} \cup OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i)\Big)\Big)\Bigg) \\
&\quad=\ c\Big(OPT^{+(S)}_{>j-1}\Big) \\
&\qquad + \sum_{i=1}^{4} \mathsf{cost}\Bigg(S_i, (up(S_i))^{\mathrm{b}(S_i)}, \\
&\qquad\qquad p\Big(uc_{S_i}\Big((up(S_i))^{\mathrm{b}(S_i)} \cup OPT^{\mathrm{i}(S_i)}_{>j-1} \cup down(S_i)\Big)\Big)\Bigg) \\
&\quad\leq\ c\left(OPT^{+(S)}_{>j-1}\right) + \sum_{i=1}^{4}\left(c\left(OPT^{\mathrm{i}(S_i)}_{>j-1}\right) + \textstyle\sum_{S_i' \subseteq S_i} c\left(OPT^{\mathrm{i}(S_i')}_{=j(S_i')}\right)\right) \\
&\quad=\ c\left(OPT^{\mathrm{i}(S)}_{>j}\right) + \sum_{S' \subseteq S} c\left(OPT^{\mathrm{i}(S')}_{=j(S')}\right),
\end{aligned}$$

proving the claim. Because l is the level of the largest disk, $up(S) = \emptyset$ and $OPT^{\mathrm{i}(S)}_{>l} = \emptyset$ for any l-square S. Hence

$$\begin{aligned}
\sum_{S;\, j(S)=l} \mathsf{cost}\left(S, \emptyset, p\left(uc_S\left(down(S)\right)\right)\right) &\leq \sum_{S;\, j(S)=l}\ \sum_{S' \subseteq S} c\left(OPT^{\mathrm{i}(S')}_{=j(S')}\right) \\
&\leq c(OPT) \\
&\leq B
\end{aligned}$$

and thus

$$\begin{aligned}
&\max\left\{ \sum_{S;\, j(S)=l} r_S \;\middle|\; \sum_{S;\, j(S)=l} \mathsf{cost}(S, \emptyset, r_S) \leq B \right\} \\
&\qquad\qquad \geq \sum_{S;\, j(S)=l} p\left(uc_S\left(down(S)\right)\right) \\
&\qquad\qquad = p\left(uc\left(\textstyle\bigcup_S OPT^{\mathrm{i}(S)}_{=j(S)}\right)\right).
\end{aligned}$$

This proves the lemma. □

It follows immediately that for any $\{r_S\}_{S;\, j(S)=l}$ attaining the maximum of Lemma 10.2.5, $\bigcup_{S;\, j(S)=l} \mathsf{sol}(S, \emptyset, r_S)$ is a set of disks of cost at most B for which the total profit of the points uniquely covered by this set is at least $p\left(uc\left(\bigcup_S OPT^{\mathrm{i}(S)}_{=j(S)}\right)\right)$, as requested.

10.2.2 Computing the cost- and sol-Functions

We say that a j-square S is *nonempty* if it contains a level j disk and *empty* otherwise. A j-square S is said to be *relevant* if one of its three siblings is nonempty, there is a nonempty $S' \supseteq S$ of level at most $j + \lceil \log k \rceil$, or $j = 0$ and S contains at least one point of $\mathcal{P}$. This implies that any nonempty square is relevant and that there are at most $\mathrm{O}(k^2\,|\mathcal{D}| + |\mathcal{P}|)$ relevant squares. Note that the definition of relevant used here is slightly different from the definition used in Chapter 7.

A relevant square S is a *relevant child* of another relevant square S' if $S \subset S'$ and there is no relevant square S'' with $S \subset S'' \subset S'$. If S is a relevant child of S', S' is a *relevant parent* of S. We show that to compute cost it is sufficient to consider only relevant squares.

Lemma 10.2.6 *For each relevant level* 0 *square* S*, all* cost- *and* sol*-values can be computed in time* $\mathrm{O}(k|\mathcal{D}^S|\,|\mathcal{P}|^2\,p_{\max}\,2^{32k\gamma/\pi})$.

Proof: Since $|\mathcal{D}^{\mathrm{b}(S)}_{>0}| \leq 16k\gamma/\pi$ by Lemma 7.1.1, enumerating all $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>0}$ takes $\mathrm{O}(2^{16k\gamma/\pi})$ time.

We show how to compute $\min\{c(T) \mid p(uc_S(T \cup W)) \geq r;\ T \subseteq \mathcal{D}^{\mathrm{i}(S)}_{\geq 0}\}$ for a particular set $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>0}$ and each $r = 0, \ldots, |\mathcal{P}| \cdot p_{\max}$. Partition S into k vertical slabs of width exactly one and assign a point of $\mathcal{P}$ to a slab if the point is contained in the slab (w.l.o.g. no point lies on a slab boundary). Observe that any disk of $\mathcal{D}^{\mathrm{i}(S)}_{\geq 0}$ covering a point in a certain slab must intersect the left or the right boundary of this slab. Now order the slabs from left to right. Define for each slab $i = 1, \ldots, k$, each subset X of the set of disks intersecting the right boundary of slab i, and each $r = 0, \ldots, |\mathcal{P}| \cdot p_{\max}$ a function $a(i, X, r)$ as

$$\begin{aligned}\min\bigl\{\, & c(X) + a(i-1, Y, r - p(uc_i(X \cup Y \cup W))) \;\bigm|\; \\ & Y \subseteq \text{disks intersecting left boundary of } i, r \geq p(uc_i(X \cup Y \cup W))\bigr\}.\end{aligned}$$

Here $uc_i(X \cup Y \cup W)$ is the set of points in slab i uniquely covered by $X \cup Y \cup W$. We set $a(0, \emptyset, 0) = 0$ and $a(0, \emptyset, r) = \infty$ for each $r \neq 0$ (as we only consider disks in $\mathcal{D}^{\mathrm{i}(S)}$, no disk intersects the left boundary of the first slab). Observe that the minima we are looking for are in $a(k, \emptyset, r)$. Furthermore, we can compute them in $\mathrm{O}(k|\mathcal{D}^S|\,|\mathcal{P}|^2\,p_{\max}\,2^{16k\gamma/\pi})$ time, as computing $p(uc_i(X \cup Y \cup W))$ takes $\mathrm{O}((|X| + |Y| + |W|) \cdot |\mathcal{P}|)$ time and a vertical line in S intersects at most $8k\gamma/\pi$ disks of $\mathcal{D}^{\mathrm{i}(S)}$ (see Lemma 7.1.1). □

To compute cost and sol for j-squares with $j > 0$, we require the following auxiliary proposition. The problem described in the proposition statement is an instance of Multiple-Choice Knapsack and may be solved using a similar method as the one in Theorem 10.1.4.

Proposition 10.2.7 *For any j-square S, given $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$ and $U \subseteq \mathcal{D}^{+(S)}_{>j-1}$, we can compute*

$$\min\left\{\sum_{i=1}^{4}\mathsf{cost}(S_i, (U \cup W)^{\mathrm{b}(S_i)}, r_i) \mid \sum_{i=1}^{4} r_i = r\right\}$$

for all $r = 0, \ldots, |\mathcal{P}| \cdot p_{\max}$ in $\mathrm{O}((|\mathcal{P}| \cdot p_{\max})^2)$ time.

Proof: We compute for all $r' = 0, .., |P| \cdot p_{\max}$ the values of

$$g(1, r') = \mathsf{cost}(S_1, (U \cup W)^{\mathrm{b}(S_1)}, r')$$

and, for all $i = 2, 3, 4$, the values of

$$g(i, r') = \min\{g(i-1, r'-r'') + \mathsf{cost}(S_i, (U \cup W)^{\mathrm{b}(S_i)}, r'') \mid r'' = 0, \ldots, r'\}.$$

This takes time $\mathrm{O}((|P| \cdot p_{\max})^2)$. The values we need are the values $g(4, r)$ for $r = 0, \ldots, |\mathcal{P}| \cdot p_{\max}$. $\square$

Lemma 10.2.8 *For each relevant j-square S ($j > 0$) which has relevant $(j-1)$-square children, all* cost*- and* sol*-values for S can be computed in time $\mathrm{O}(2^{32k\gamma/\pi}(|\mathcal{P}| \cdot p_{\max})^2)$.*

Proof: This follows from $|\mathcal{D}^{\mathrm{b}(S)}_{>j}|, |\mathcal{D}^{+(S)}_{>j-1}| \leq 16k\gamma/\pi$ (see Lemma 7.1.1) and Proposition 10.2.7. $\square$

Lemma 10.2.9 *For each relevant j-square S ($j > 0$) with no relevant level $j-1$ children, in $\mathrm{O}((|\mathcal{D}|+|\mathcal{P}|) \cdot 2^{32\gamma/\pi}(|\mathcal{P}| \cdot p_{\max})^2)$ time all* cost*- and* sol*-values can be computed.*

Proof: It is easy to show from Lemma 7.2.6 that

$$\mathsf{cost}(S, W, r) =$$

$$\min\left\{\sum_{\text{relevant child } S' \text{ of } S} \mathsf{cost}(S', W^{b(S')}, r_{S'}) \;\middle|\; \sum_{\text{relevant child } S' \text{ of } S} r_{S'} = r\right\}$$

for any $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$ and any $r = 0, \ldots, |\mathcal{P}| \cdot p_{\max}$. Furthermore, $\mathcal{D}^{\mathrm{b}(S)}_{>j} = \mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}$ and thus $\left|\mathcal{D}^{\mathrm{b}(S)}_{>j}\right| = \left|\mathcal{D}^{\mathrm{b}(S)}_{\geq j+\lceil \log k \rceil}\right| \leq 32\gamma/\pi$. For a fixed $W \subseteq \mathcal{D}^{\mathrm{b}(S)}_{>j}$, we follow the approach of Proposition 10.2.7 and extend it to deal with all relevant children of S. The increase in time complexity is linear in the number of relevant children of S. As a relevant square has $\mathrm{O}(|\mathcal{D}|+|\mathcal{P}|)$ relevant children, the time bound follows. $\square$

Proof of Theorem 10.2.1: Combining Lemmas 10.2.6, 10.2.8, and 10.2.9, we can compute the cost- and sol-values for all relevant squares in $\mathrm{O}((k^2|\mathcal{D}|+|\mathcal{P}|) \cdot k\,|\mathcal{D}| \cdot 2^{32k\gamma/\pi}(|\mathcal{P}| \cdot p_{\max})^2)$ time. By using an extension to Proposition 10.2.7 as in Lemma 10.2.9, $\max\left\{\sum_{S;\, j(S)=l} r_S \;\middle|\; \sum_{S;\, j(S)=l} \mathsf{cost}(S, \emptyset, r_S) \leq B\right\}$ can be computed in $\mathrm{O}((|\mathcal{D}| + |\mathcal{P}|) \cdot (|\mathcal{P}| \cdot p_{\max})^2)$ time. The theorem then follows from Lemma 10.2.5. $\square$

10.2.3 The Approximation Algorithm

Given the above algorithm, the shifting technique is applied as follows. Given an integer a ($0 \le a \le k-1$), a line of level j is *active* if it is of the form $y = (hk + a2^{l-j})2^j$ or $x = (vk + a2^{l-j})2^j$ for $h, k \in \mathbb{Z}$. Active lines partition the plane into j-squares as before, shifted with respect to a. Hence we can still apply Theorem 10.2.1. Let C_a denote the set of disks output by the algorithm for the j-squares induced by a and let $C_{\max}$ be the set among such sets with the maximum profit of uniquely covered points of $\mathcal{P}$.

Lemma 10.2.10 *$p(uc(C_{\max})) \ge (1 - 4/k) \cdot p(uc(OPT))$, where OPT is a solution for which $p(uc(OPT))$ is maximum under $c(OPT) \le B$.*

Proof: Let $\mathcal{D}_a^b = \bigcup_S \mathcal{D}_{=j(S)}^{\mathrm{b}(S)}$ be the set of disks intersecting the boundary of a j-square S induced by a at their level. Observe that for a fixed value of a the points uniquely covered by OPT by disks in $OPT \cap \mathcal{D}_a^b$ are precisely the points uniquely covered by $OPT \cap \mathcal{D}_a^b$ not covered by disks in $OPT - \mathcal{D}_a^b$. Then it follows from Theorem 10.2.1 that

$$\begin{aligned} p(uc(C_a)) &\ge p\left(uc\left(\textstyle\bigcup_S OPT_{=j(S)}^{\mathrm{i}(S)}\right)\right) \\ &= p\left(uc\left(OPT - (OPT \cap \mathcal{D}_a^b)\right)\right) \\ &\ge p\left(uc(OPT)\right) - p\left((uc(OPT \cap \mathcal{D}_a^b) - cov(OPT - \mathcal{D}_a^b)\right). \end{aligned}$$

Following the proof of Lemma 7.2.8, a disk is in $\mathcal{D}_a^b$ for at most 4 values of a. Hence

$$\textstyle\sum_{a=0}^{k-1} p\left(uc(OPT \cap \mathcal{D}_a^b) - cov(OPT - \mathcal{D}_a^b)\right) \le 4 \cdot p(uc(OPT)).$$

Then

$$\begin{aligned} &k \cdot p(uc(C_{\max})) \\ &\quad\ge \textstyle\sum_{a=0}^{k-1} p(uc(C_a)) \\ &\quad\ge \textstyle\sum_{a=0}^{k-1} \left(p(uc(OPT)) - p\left(uc(OPT \cap \mathcal{D}_a^b) - cov(OPT - \mathcal{D}_a^b)\right)\right) \\ &\quad\ge k \cdot p(uc(OPT)) - 4 \cdot p(uc(OPT)). \end{aligned}$$

Therefore $p(uc(C_{\max})) \ge (1 - 4/k) \cdot p(uc(OPT))$. □

In the following, we denote by x the given instance of Budgeted Unique Coverage and by $|x|$ the length of some natural encoding of x. We can clearly assume that $|x| \ge |\mathcal{D}| + |\mathcal{P}|$, where $\mathcal{D}$ and $\mathcal{P}$ denote the given sets of disks and points, respectively.

Theorem 10.2.11 *There is an eptas for Budgeted Unique Coverage on disks of bounded ply, i.e. of ply $\gamma = \gamma(|x|) = \mathrm{o}(\log |x|)$.*

Proof: Consider any $\epsilon > 0$. Choose k as the largest odd integer such that $32k\gamma/\pi \leq \log|x|$. If $k < 5$, output $\emptyset$ (or any other arbitrary solution of cost at most B). Otherwise, scale the profits by $\frac{\epsilon \cdot p_{\max}}{|\mathcal{P}|}$, similar to Theorem 9.1.21. Following the analysis of this theorem and applying Theorem 10.2.1 and Lemma 10.2.10, we obtain a $(1-4/k)\cdot(1-\epsilon)$-approximation of the optimum in

$$O(k^2(k^2\,|\mathcal{D}| + |\mathcal{P}|)\cdot|\mathcal{D}|\cdot 2^{32k\gamma/\pi}(|\mathcal{P}|^2/\epsilon)^2)$$

time. By the choice of k, this is bounded by

$$O(\log^2|x|\cdot(|\mathcal{D}|\log^2|x| + |\mathcal{P}|)\cdot|\mathcal{D}|\cdot|x|\cdot(|\mathcal{P}|^2/\epsilon)^2).$$

Hence in time polynomial in the size of the input and $1/\epsilon$, a feasible solution is computed. Furthermore, there is a c_ϵ such that $k \geq 4/\epsilon$ and $k \geq 5$ for all $|x| \geq c_\epsilon$. Therefore if $|x| \geq c_\epsilon$, we obtain a $(1-\epsilon)^2$-approximation of the optimum. The theorem then follows from Theorem 2.2.4. □

This theorem can be easily extended to Budgeted Low-Coverage on fat objects of bounded ply. We simply adjust the algorithm of Lemma 10.2.6 to modulate the profits by the satisfactions and apply scaling as in Theorem 10.1.5.

Theorem 10.2.12 *There is an eptas for Budgeted Low-Coverage on fat objects of bounded ply, i.e. of ply $\gamma = \gamma(|x|) = o(\log|x|)$.*

If the objects are arbitrary unit disks, but the density of the set of points (i.e. the maximum number of points of $\mathcal{P}$ in any 1×1 box) is bounded, then we can reduce the set of unit disks to a set of unit disks of bounded ply.

Lemma 10.2.13 *For any instance of Budgeted Low-Coverage on a set $\mathcal{D}$ of unit disks and a set $\mathcal{P}$ of points for which the density is bounded by some constant $d > 0$, there is an equivalent instance on a set $\mathcal{D}'$ of unit disks of ply at most $324d^2$ and the same set $\mathcal{P}$ of points.*

Proof: Remove all disks from $\mathcal{D}$ that do not cover any point of $\mathcal{P}$. Then iteratively remove disks d from $\mathcal{D}$ for which there is a disk $d' \in \mathcal{D}$ such that $d \cap \mathcal{P} = d' \cap \mathcal{P}$ and $c(d) \geq c(d')$. Let $\mathcal{D}'$ be the resulting set of disks. Clearly, the instance of Budgeted Low-Coverage on $\mathcal{D}'$ and $\mathcal{P}$ is equivalent to the instance on the original sets $\mathcal{D}$ and $\mathcal{P}$. Moreover, we claim that the ply of this set $\mathcal{D}'$ is at most $324d^2$.

To see this, we apply an argument reminiscent of one by Hochbaum and Maass [150]. Let $\delta > 0$ denote the smallest distance over all disks $d \in \mathcal{D}'$ of d and a point of $\mathcal{P} - d$. For any disk $d \in \mathcal{D}'$, let $\widehat{d}$ denote the disk of radius $1/2 + \min\{\delta, 1/4\}$ centered at the same point as d. We may assume that for any disk $d \in \mathcal{D}'$ there are two points $u, u' \in \mathcal{P}$ such that u and u' lie on the boundary of d or $\widehat{d}$. Otherwise we can move d to such a position while keeping $d \cap \mathcal{P}$ the same. By the construction of $\mathcal{D}'$, this implies that any pair of points $u, u' \in \mathcal{P}$ can 'induce' at most 8 disks of $\mathcal{D}'$.

Now consider some point $p \in \mathbb{R}^2$. By the preceding arguments, any disk of $\mathcal{D}'$ overlapping p is induced by points that are within distance 3/2 of p. Because the density is d, there are at most $9d$ such points. Including disks that cover a single point of $\mathcal{P}$, the number of disks of $\mathcal{D}'$ overlapping p is at most $8\binom{9d}{2} + 9d \leq 324d^2$. $\square$

The lemma clearly extends to unit-size fat objects.

Now use Lemma 10.2.13 in conjunction with Theorem 10.2.12 to improve on the results of Section 10.1 if the point set has bounded density.

Theorem 10.2.14 *There is an eptas for Budgeted Low-Coverage on unit fat objects and a set of points of bounded density, i.e. of density* $d = d(|x|) = \mathrm{o}(\sqrt{\log |x|})$.

Finally, observe that the algorithms of Theorem 10.2.12 and Theorem 10.2.14 extend to Geometric Budgeted Maximum Coverage, which was already discussed in Chapter 9. To approximate this problem, adjust the algorithm of Lemma 10.2.6 to include the profit of all covered points, not only of uniquely covered points.

Theorem 10.2.15 *There exists an eptas for Geometric Budgeted Maximum Coverage on fat objects of bounded ply, i.e. of ply* $\gamma = \gamma(|x|) = \mathrm{o}(\log |x|)$, *and on unit fat objects and a point set of bounded density, i.e. of density* $d = d(|x|) = \mathrm{o}(\sqrt{\log |x|})$. *Here* $|x|$ *is the size of the input.*

10.3 Geometric Membership Set Cover

We study the geometric version of Minimum Membership Set Cover. We give an approximation algorithm for Geometric Membership Set Cover on unit squares (squares with side length 1), achieving a constant approximation ratio. Its running time is polynomial only if the optimum objective value is bounded by a constant.

Consider an instance of Minimum Membership Set Cover on a set $\mathcal{S}$ of unit squares and a set $\mathcal{P}$ of points. We partition the plane into horizontal slabs of height 1 and compute a separate solution for each such slab. The solution for a slab must cover all points in the slab while ensuring that the maximum membership of points inside and outside the slab is bounded. In the end, the output is the union of the solutions for the different slabs.

To solve the problem for a slab M, we observe that the unit squares of any minimal solution consist of squares intersecting the top boundary of M, all of which are on the lower envelope of their union, and of squares intersecting the bottom boundary of M, all of which are on the upper envelope of their union. This enables a sweep-line approach in which we maintain $2\ell+1$ squares around the current position on each of the two envelopes, where ℓ is a given bound on the maximum membership. The details are presented in the following lemma.

Lemma 10.3.1 *Let $\mathcal{P}$ be a set of points and $\mathcal{S}$ be a set of unit squares. Let M denote the slab contained between the lines $y = 0$ (inclusive) and $y = 1$ (exclusive). Let $\mathcal{P}_M = \mathcal{P} \cap M$ and $\overline{\mathcal{P}_M} = \mathcal{P} - \mathcal{P}_M$. For any constant ℓ, there is a polynomial-time algorithm that either asserts that there is no $C \subseteq \mathcal{S}$ that covers $\mathcal{P}$ with $\mathrm{mem}_C(\mathcal{P}) \leq \ell$, or computes a set $C \subseteq \mathcal{S}$ that covers all points in $\mathcal{P}_M$ and satisfies $\mathrm{mem}_C(\mathcal{P}_M) \leq \ell$ (points inside M are covered at most ℓ times) and $\mathrm{mem}_C(\overline{\mathcal{P}_M}) \leq 2\ell$ (points outside M are covered at most 2ℓ times).*

Proof: Let $\mathcal{S}_t$ ($\mathcal{S}_b$) be the set of squares in $\mathcal{S}$ that intersect the top (bottom) boundary of M. We may assume that no square intersects both boundaries. Then all squares that intersect M are in either $\mathcal{S}_t$ or $\mathcal{S}_b$. We use a vertical sweep-line that moves from left to right and stops at all points in $\mathcal{P}_M$. Let $\mathcal{P}_M = \{u_1, \ldots, u_k\}$, with points indexed in order of nondecreasing x-coordinate. For a given position of the sweep-line, a $(2\ell + 1)$-tuple $\mathsf{s}_t = (s_1, \ldots, s_{2\ell+1})$ of distinct squares from $\mathcal{S}_t$ is a *proper tuple* if all squares are on the lower envelope of their union (in the order $s_1, s_2, \ldots, s_{2\ell+1}$) and the intersection of the sweep-line and this envelope is with $s_{\ell+1}$. We allow any prefix and/or suffix of a proper tuple to consist of dummy objects that do not contain any points in order to represent the case that fewer than ℓ squares appear before or after the current position on the lower envelope. We ignore this technicality below. Proper tuples s_b from $\mathcal{S}_b$ are defined analogously.

If the sweep-line is at the point u_i, then a pair $(\mathsf{s}_t, \mathsf{s}_b)$ of proper tuples is called *admissible* if

- the point u_i is covered by the union of the squares in s_t and s_b, and
- the maximum membership of the union of the squares in s_t and s_b (taking into account all points of $\mathcal{P}$, also the ones outside M) is at most ℓ.

For each point $u_i \in \mathcal{P}_M$, we consider the set A_i of all admissible pairs $(\mathsf{s}_t, \mathsf{s}_b)$. We say that $(\mathsf{s}_t, \mathsf{s}_b) \in A_i$ and $(\mathsf{s}'_t, \mathsf{s}'_b) \in A_{i+1}$ are *compatible* if

- $\mathsf{s}_t = \mathsf{s}'_t$ or $\mathsf{s}'_t = (s_2, \ldots, s_{2\ell+1}, s_{2\ell+2})$ for some new square $s_{2\ell+2}$, and
- the analogous condition holds for s_b and s'_b.

Then we check if a sequence of compatible tuples exists from some $(\mathsf{s}_t, \mathsf{s}_b) \in A_1$ to some $(\mathsf{s}'_t, \mathsf{s}'_b) \in A_k$. If there is such a sequence π, the union of the squares from all tuples of the sequence is the solution C. Otherwise, the algorithm outputs that there is no solution with maximum membership at most ℓ.

We see that the algorithm is correct as follows. If the algorithm outputs a solution C, it is clear that C covers $\mathcal{P}_M$. To bound the maximum membership, note that the solution consists of a set C_t of squares from $\mathcal{S}_t$ that all meet the lower envelope of their union, and a set C_b of squares from $\mathcal{S}_b$ that all meet the upper envelope of their union. Observe that $C = C_t \cup C_b$. We imagine the squares in C_t to be ordered from left to right as they appear on their lower envelope, and similarly for C_b. Terms such as 'before,' 'after,' and 'between' refer to this order. We now argue separately about the maximum membership for points in $\mathcal{P}_M$ and in $\overline{\mathcal{P}_M}$ and show that the membership for points in $\mathcal{P}_M$ is at most ℓ and that for points in $\overline{\mathcal{P}_M}$ it is at most 2ℓ.

For a point u_i in $\mathcal{P}_M$, note that the squares in C_t that contain u_i are consecutive on the lower envelope of C_t. If there were two squares containing u_i and another square in between that does not contain u_i, then the latter square would not be on the lower envelope. Furthermore, if C_t contains any square covering u_i, then also the square that is on the envelope at the x-coordinate of u_i must contain u_i. Hence, if g squares from C_t contain u_i, then the proper tuple s_t of the admissible pair $(\mathsf{s}_t, \mathsf{s}_b)$ that was chosen for the sweep-line at u_i must contain at least $\min\{g, \ell+1\}$ squares containing u_i. The analogous statement holds for C_b and s_b. Therefore, if more than ℓ squares from $C_t \cup C_b$ were to contain u_i, then at least $\ell + 1$ of the squares in the pair $(\mathsf{s}_t, \mathsf{s}_b)$ would contain u_i, and thus the pair would not be admissible. Hence points in $\mathcal{P}_M$ are covered at most ℓ times, i.e. $\text{mem}_C(\mathcal{P}_M) \leq \ell$.

Now consider a point u in $\overline{\mathcal{P}_M}$ that lies below M (the reasoning for points above M is analogous). The point u cannot be contained in any square in C_t. Consider the set C_u of squares in C_b that contain u. We claim that C_u consists of at most two subsets of consecutive (i.e. consecutive on the upper envelope of C_b) squares in C_b. This can be seen as follows. Consider the first (i.e. leftmost) square x in C_b that contains u. Let y be the first square after x in C_b that does not contain u. If y is entirely to the right of u, then no further square in C_b can contain u, and thus C_u is one consecutive subset of C_b (starting with x and ending with the square just before y).

So assume that there is a square after y in C_b that contains u. Then y must be entirely above u. Let z be the first square after y that contains u. Clearly, z must be to the lower right of y (i.e. z can be obtained from y by shifting y right and down). All further squares (after z) whose x-range contains the x-coordinate of u must be to the lower right of z and hence contain u. For if one of them, say w, was to the upper right of z, then z could not be on the upper envelope of C_b, as z would be below the upper envelope of $\{y, w\}$. By repeating this argument for squares to the right of z, we can show that C_u consists of at most two consecutive subsets of C_b.

By construction, the number of consecutive squares from C_b containing u is bounded by ℓ. Otherwise, the sequence π would include a pair of tuples having $\ell+1$ consecutive squares containing u, but such a pair would not be admissible. As there are at most two consecutive subsets of C_b containing u, we have that u is contained in at most 2ℓ squares from C_b. Hence $\text{mem}_C(\overline{\mathcal{P}_M}) \leq 2\ell$.

We have shown that if the algorithm outputs a solution C, then C covers all points in $\mathcal{P}_M$ and satisfies $\text{mem}_C(\mathcal{P}_M) \leq \ell$ and $\text{mem}_C(\overline{\mathcal{P}_M}) \leq 2\ell$. On the other hand, if there is a solution that covers $\mathcal{P}_M$ and has maximum membership at most ℓ, then the $(2\ell + 1)$-tuples of consecutive squares on the two envelopes allow to construct a valid candidate sequence for the algorithm, and thus the algorithm will indeed output a solution. This implies that if the algorithm does not output a solution, then there is no solution that covers $\mathcal{P}_M$ and has membership at most ℓ (and thus also no solution that covers $\mathcal{P}$ and has membership at most ℓ).

For the running time, note that each A_i contains $O(|\mathcal{S}|^{4\ell+2})$ pairs of tuples. Moreover, for each admissible pair in A_i, there are $O(|\mathcal{S}|^2)$ compatible pairs in A_{i+1}. It follows that we can check for the existence of a sequence π of compatible pairs in $O(|\mathcal{P}| \cdot |\mathcal{S}|^{4\ell+4})$ time. □

Theorem 10.3.2 *There is a polynomial-time* 5*-approximation algorithm for instances of Geometric Membership Set Cover on unit squares if the optimal objective value is bounded by an arbitrary constant L.*

Proof: For a given constant ℓ, the following procedure either computes a solution with maximum membership at most 5ℓ or asserts that no solution with maximum membership at most ℓ exists. Partition the plane into horizontal slabs of unit height. For each slab M that contains at least one point from $\mathcal{P}$, run the algorithm of Lemma 10.3.1 to compute a cover $C_M \subseteq \mathcal{S}$ for the points inside M with maximum membership at most ℓ for points in M and at most 2ℓ for points outside M. If for one of the slabs the algorithm of Lemma 10.3.1 outputs that there is no cover with maximum membership at most ℓ, return that the whole instance has no solution with objective value at most ℓ. Otherwise, return the union of the solutions C_M computed for all slabs M. Note that the squares in the solution computed for a slab M can only cover points in M and in the slabs directly above and below M. A point in M is covered at most ℓ times by squares in the solution computed for M, at most 2ℓ times by squares in the solution computed for the slab directly above M, and at most 2ℓ times by squares in the solution computed for the slab directly below M. This shows that every point in $\mathcal{P}$ is covered at most 5ℓ times.

Now run the above procedure for $\ell = 1, 2, \ldots, L$. The first time the procedure returns a cover, we output that cover and terminate. If the procedure does not return a cover for any of the calls, we output that the instance does not have a solution with maximum membership at most L. □

The approximation algorithm does not seem to extend to unit disks directly. One problem is that a point outside a slab could be contained in several unit disks that are not consecutive on the envelope of the selected unit disks. Hence even if the maximum membership of consecutive unit disks on an envelope is bounded by ℓ, the maximum membership of the overall solution could be large.

10.4 Hardness of Approximation

We give several hardness results for both Geometric Unique Coverage and Geometric Membership Set Cover.

We first consider Geometric Unique Coverage. Recall that the approximability of Unique Coverage has not been fully settled yet [83]. We can show however that in some cases the approximability of Geometric Unique Coverage is equal to approximability of general Unique Coverage.

Theorem 10.4.1 *There is a gap-preserving reduction from Unique Coverage to Geometric Unique Coverage on the following objects:*

- *convex polygons,*
- *translated copies of a single polygon,*
- *rotated copies of a single convex polygon,*
- *α-fat objects for any $\alpha > 1$.*

Proof: We use similar reductions as in Section 8.5 in the way outlined in the proof of Theorem 9.2.1. □

This theorem implies for instance that for any $\epsilon > 0$, it is hard to approximate Geometric Unique Coverage on the above object types within ratio $\Omega(1/\log^{\sigma(\epsilon)} n)$, assuming that NP $\not\subseteq$ BPTIME(2^{n^ϵ}), where $\sigma(\epsilon)$ is some constant dependent on ϵ.

For Minimum Membership Set Cover, we can give more results. Recall that Minimum Membership Set Cover is not approximable within $(1-\epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{O(\log\log n)}$) [186]. In a more restricted setting, the problem remains APX-hard. Let *Minimum Membership k-Set Cover* be the variant of Minimum Membership Set Cover where each set has cardinality at most k.

Theorem 10.4.2 *Minimum Membership k-Set Cover is APX-hard for $k \geq 4$.*

Proof: Recall that Minimum k-Set Cover is APX-hard for any $k \geq 3$ (by reduction from Minimum Vertex Cover on graphs of degree at most 3 [9]). Kuhn et al. [186] give a gap-preserving reduction from Minimum Set Cover to Minimum Membership Set Cover. Here one element is added to the universe and to each set. Hence a minimum set cover corresponds to a set cover of minimum membership in the new set system and vice versa. The theorem follows immediately. □

Using these results, we can prove the following theorem in the same spirit as Theorem 9.2.1.

Theorem 10.4.3 *Geometric Membership Set Cover is not approximable to $(1-\epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{O(\log\log n)}$), on the following objects:*

- *convex polygons,*
- *translated copies of a single polygon,*
- *rotated copies of a single convex polygon,*
- *α-fat objects for any $\alpha > 1$.*

Geometric Membership Set Cover is APX-hard on the following objects:

- *convex polygons with r corners, where $r \geq 4$,*
- *α-fat objects of constant description complexity for any $\alpha > 1$.*

We strengthen these APX-hardness results by showing a lower bound on the approximation ratio that any polynomial-time approximation algorithm for Geometric Membership Set Cover can attain on unit disks and on unit squares.

Theorem 10.4.4 *There is no polynomial-time approximation algorithm attaining an approximation ratio smaller than 2 for Minimum Membership Set Cover on unit disks or unit squares, unless P=NP.*

Proof: We claim that for Geometric Membership Set Cover on unit disks or on unit squares, it is NP-hard to decide if a solution with maximum membership 1 exists or not. It is clear that the theorem statement follows from this claim.

We give a reduction from the NP-complete problem of checking whether a planar graph G of maximum degree 4 is 3-colorable [116]. We create an instance of Minimum Membership Set Cover on unit disks as follows (the construction on unit squares is very similar). First, we compute a rectilinear embedding of G in the plane [255], which determines the layout of the components of the construction. For each vertex v, we construct a vertex gadget as shown in Figure 10.2. In order to cover the point p_v once, a solution must choose exactly one of the three disks containing p_v, and this corresponds to assigning a color to v. Depending on the choice, either 0, 1, or 2 points among the triple of points on the right are already covered.

The next gadget is a transport gadget, which allows transporting a chosen color along a chain of disks. The gadget for transporting information from left to right is shown in Figure 10.3.

Depending on which of the three disks on the left is in the solution, the triple of points uniquely determines which one of the next three disks needs to be chosen to achieve maximum membership equal to 1, and so on.

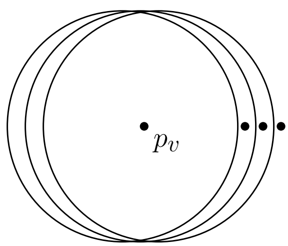

Figure 10.2: The vertex gadget.

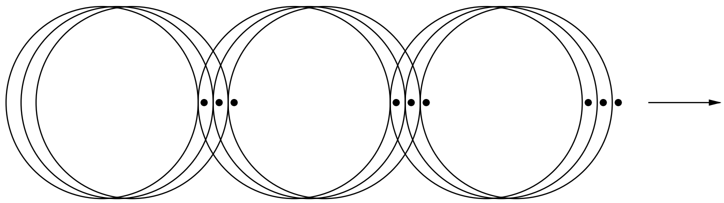

Figure 10.3: The transport gadget.

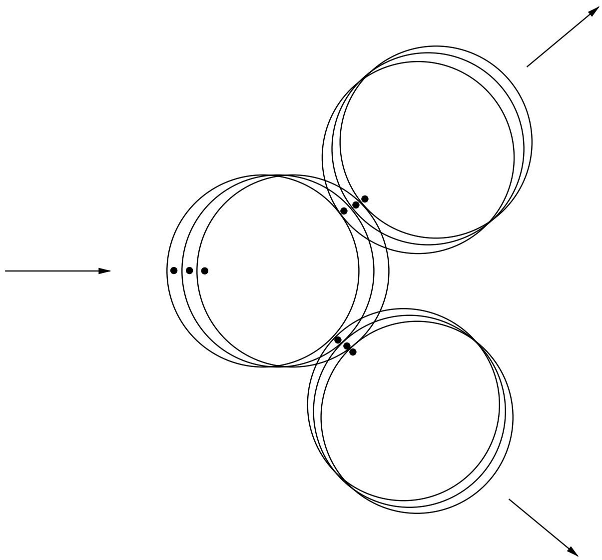

Figure 10.4: The copy gadget.

It is also easy to duplicate information. The copy gadget shown in Figure 10.4 demonstrates how this is accomplished.

Finally, we need a checking gadget to check whether two vertices u and v that are adjacent in G have indeed been assigned different colors. The gadget is shown in Figure 10.5, assuming that a chain transporting the color of u arrives from the left and a chain transporting the color of v arrives from the right. The triple of points on the left (of which 0, 1, or 2 are already covered by the chain transporting u's color) forces a unique choice of exactly one of the three solid disks on the left to be included in the solution, and similarly for the triple of points on the right. The solid disk d_ℓ chosen on the left contains exactly one of a, b, and c, and the same holds for the solid disk d_r chosen on the right. If u and v have received the same color, then d_ℓ and d_r cover the same point among a, b and c, a contradiction to the maximum membership being 1. If u and v have different colors, then d_ℓ and d_r cover two different points among a, b, and c, and the third point can be covered by one of the three dashed disks.

In summary, we have that the constructed instance of Minimum Membership Set Cover on unit disks has a solution with maximum membership 1 if and only if the given planar graph is 3-colorable.

A similar construction is possible for unit squares, only the checking gadget is slightly more complicated (see Figure 10.6). Each of the three points in the middle needs to be covered by a dashed square, but each of the dashed squares contains also the point to its left or to its right (or both). If both vertices have received the same color, the row for that color will have the point to the left and to the right of the middle point already covered, and it is impossible to cover only the point in the middle with a dashed square. □

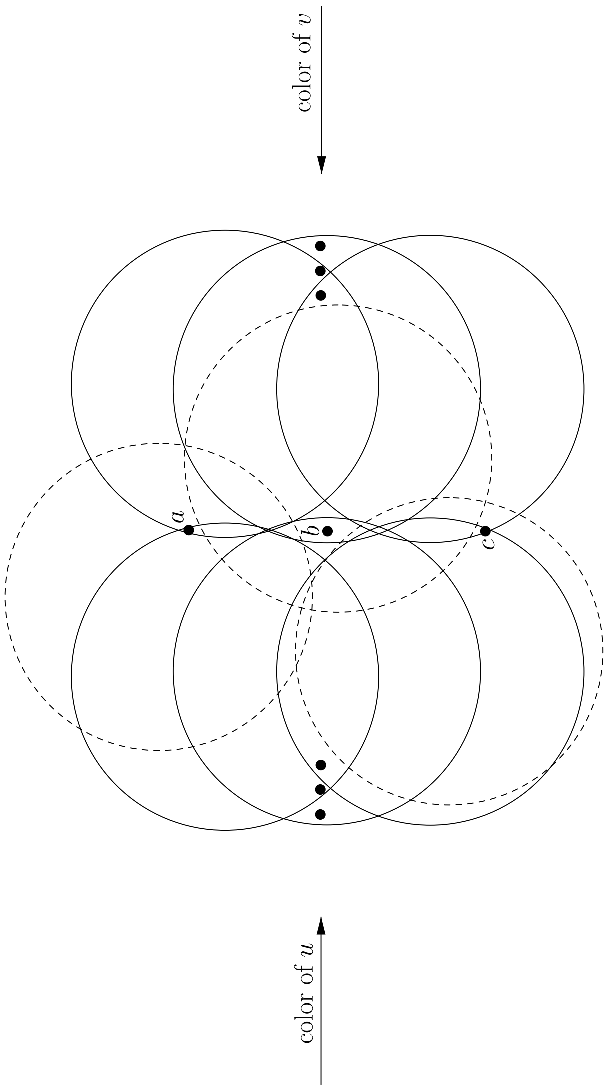

Figure 10.5: The checking gadget.

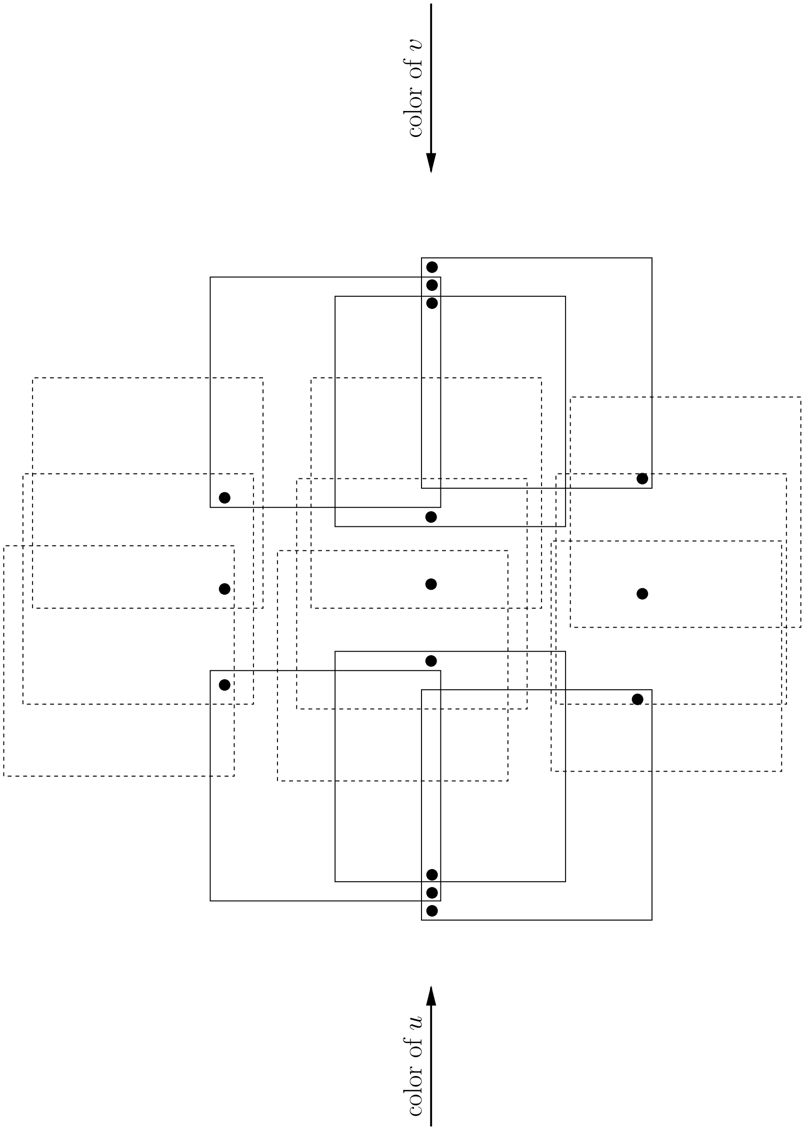

Figure 10.6: The checking gadget for unit squares.

Chapter 11

Conclusion

In this thesis, we studied the approximability of hard combinatorial optimization problems on (intersection graphs of) systems of geometric objects. We gave both positive results, in the form of approximation algorithms, and negative results, in the form of hardness of approximation statements. Here we look back at these results and provide directions for further research.

First, it is clear that having (the intersection graph of) a system of geometric objects as the underlying structure of an optimization problem helps to approximate it better than would be possible in the general case. We provided new, improved approximation schemes for Maximum Independent Set and Minimum Vertex Cover on intersection graphs of fat objects in any constant dimension and in the weighted case. We presented the first approximation algorithms for Minimum (Connected) Dominating Set on intersection graphs of objects of arbitrary size and improved approximation schemes on unit fat objects. We gave the first approximation scheme for Geometric Set Cover on unit squares and the first approximation algorithms for Geometric Unique Coverage on unit disks and on unit squares.

Secondly, we demonstrated that the type of the objects in the system of geometric objects has a significant impact on the approximability of optimization problems on such systems. Maximum Independent Set and Minimum Vertex Cover on intersection graphs of three-dimensional convex polygons of ply 1 and Minimum Dominating Set, Geometric Set Cover, Geometric Unique Coverage, and Geometric Membership Set Cover on planar fat objects are as hard on systems of these objects as on general set systems. Minimum Dominating Set and Geometric Set Cover are APX-hard on arbitrary rectangles. This is in sharp contrast to the positive results of the previous paragraph.

The influence of the object type is nowhere more visible than with Minimum Dominating Set on intersection graphs of objects of arbitrary size. The immediate open question there is whether Minimum Dominating Set admits a constant-factor approximation algorithm or even a ptas on disk graphs of arbitrary ply. The hardness results of Section 8.5 show that on objects whose boundaries can intersect an arbitrary number of times, Minimum Dominating Set is very hard to approximate. On the contrary, if object boundaries intersect at most twice (i.e. the objects are pseudo-disks), a linear-size ϵ-net exists and at least for cases such as r-polygons with constant r or rectangles with bounded aspect-ratio, we get constant-factor approximation algorithms. An intriguing question is whether Minimum Dominating Set on disk graphs is harder to approximate than on other intersection graph classes such as intersection graphs of squares, or whether the algorithmic ideas of Chapter 8 can be extended to disks or maybe even to arbitrary pseudo-disks (the linear bound on the size of an ϵ-net starts to fail 'naturally' beyond pseudo-disks).

Similarly, in Section 8.4 we used the shifting technique in a $(3+\epsilon)$-approximation algorithm for Minimum Dominating Set on disk graphs of bounded ply. However, we do not know if the shifting technique can be used to give a constant-factor approximation algorithm (or even a ptas) for Minimum Dominating Set on disk graphs of arbitrary ply. We can point to two reasons for this. First, there is no upper bound on the number of 'large' disks intersecting a j-square that are in the dominating set. Secondly, we cannot track which j-square is 'responsible' for dominating a disk intersecting more than one j-square at its level. The algorithms in Section 8.4 got around the first problem by assuming that the ply is bounded and around the second problem by considering $\preceq_{\mathsf{Leb}}$-dominating sets (Theorem 8.4.19), or by disregarding the domination of disks intersecting a boundary on their level and combining three result sets (Theorem 8.4.29).

The weighted case of Minimum Dominating Set also poses interesting new questions. We presented the first ptas for this problem on intersection graphs of two-dimensional geometric objects, namely on unit square graphs. On unit disk graphs however, the best result so far is a $(5+\epsilon)$-approximation algorithm [74]. Can one improve to a ptas in this case?

Another problem where we do not yet have a clear picture of its approximability is Geometric Set Cover and its variants. For Geometric Set Cover on unit squares we discovered a ptas, the first approximation scheme for this problem. The scheme even extends to the weighted case. Recently, a ptas for Geometric Set Cover on arbitrary disks was announced [217], but it does not extend to the weighted case. Does the weighted case have a ptas? Hopefully the ideas behind the scheme on unit squares can be used for a ptas on systems of unit disks or other objects.

With respect to variants of Geometric Set Cover, we considered Geometric Unique Coverage and gave constant-factor approximation algorithms on unit disks and unit squares. There is as yet however no reason to suspect that the attained constants are optimal. In particular, one wonders whether the ptas for Geometric Set Cover on unit squares can be adapted to this problem. The geometric version of Minimum Membership Set Cover presents even bigger challenges. We proposed a constant-factor approximation algorithm that runs in polynomial time if the optimum is constant. Can this restriction be lifted?

In summary, this thesis answered many question surrounding the approximability of optimization problems on systems of geometric objects. These answers in turn lead to new questions, that will be challenging to answer...

Bibliography

[1] Abramson, N., “Development of the ALOHANET”, *IEEE Transactions on Information Theory* IT-31:2 (March 1985), pp. 119–123. Cited on p. 3

[2] Agarwal, P.K., Mustafa, N.H., “Independent set of intersection graphs of convex objects in 2D”, *Computational Geometry* 34:2 (May 2006), pp. 83–95. Cited on p. 44, 45

[3] Agarwal, P.K., Procopiuc, C.M., “Exact and Approximation Algorithms for Clustering”, *Algorithmica* 33:2 (December 2002), pp. 201–226. Cited on p. 163

[4] Agarwal, P.K., van Kreveld, M.J., Suri, S., “Label Placement by Maximum Independent Set in Rectangles”, *Computational Geometry* 11:3–4 (December 1998), pp. 209–218. Cited on p. 2, 4, 45

[5] Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E., “Wireless sensor networks: a survey”, *Computer Networks* 38:4 (March 2002), pp. 393–422. Cited on p. 3

[6] Alber, J., Bodlaender, H.L., Fernau, H., Kloks, T., Niedermeier, R., “Fixed Parameter Algorithms for DOMINATING SET and Related Problems on Planar Graphs”, *Algorithmica* 33:4 (August 2002), pp. 461–493. Cited on p. 50, 57, 59

[7] Alber, J., Fiala, J., “Geometric separation and exact solutions for the parameterized independent set problem on disk graphs”, *Journal of Algorithms* 52:2 (August 2004), pp. 134–151. Cited on p. 70

[8] Alber, J., Niedermeier, R., “Improved Tree Decomposition Based Algorithms for Domination-like Problems” in Rajsbaum, S. (ed.) *LATIN 2002: Theoretical Informatics, 5th Latin American Symposium, Cancun, Mexico, April 3-6, 2002, Proceedings*, Lecture Notes in Computer Science 2286, Springer-Verlag, Berlin, 2002, pp. 613–628. Cited on p. 50, 57, 59

[9] Alimonti, P., Kann, V., “Some APX-completeness results for cubic graphs”, *Theoretical Computer Science* 237:1–2 (April 2000), pp. 123–134. Cited on p. 148, 150, 200

[10] Alon, N., Katchalski, M., Pulleyblank, W.R., "Cutting Disjoint Disks by Straight Lines", *Discrete and Computational Geometry* 4:1 (December 1989), pp. 239–243. Cited on p. 69

[11] Alt, H., Arkin, E.M., Brönnimann, H., Erickson, J., Fekete, S.P., Knauer, C., Lenchner, J., Mitchell, J.S.B., Whittlesey, K., "Minimum-Cost Coverage of Point Sets by Disks" in Amenta, N., Cheong, O. (eds.) *Proceedings of the 22nd ACM Symposium on Computational Geometry, Sedona, Arizona, USA, June 5-7, 2006*, Association for Computing Machinery, 2006, pp. 449–458. Cited on p. 163

[12] Alzoubi, K.M., Wan, P.J., Frieder, O., "Distributed Heuristics for Connected Dominating Sets in Wireless Ad Hoc Networks", *Journal of Communications and Networks* 4:1 (March 2002), pp. 22–29. Cited on p. 47

[13] Ambühl, C., Erlebach, T., Mihalák, M., Nunkesser, M., "Constant-Factor Approximation for Minimum-Weight (Connected) Dominating Sets in Unit Disk Graphs" in Díaz, J., Jansen, K., Rolim, J.D.P., Zwick, U. (eds.) *Approximation, Randomization, and Combinatorial Optimization, Algorithms and Techniques, 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2006, and 10th International Workshop on Randomization and Computation, RANDOM 2006, Barcelona, Spain, August 28-30 2006, Proceedings*, Lecture Notes in Computer Science 4110, Springer-Verlag, Berlin, 2006, pp. 3–14. Cited on p. 47, 83, 161, 183

[14] Ambühl, C., Wagner, U., "The Clique Problem in Intersection Graphs of Ellipses and Triangles", *Theory of Computing Systems* 38:3 (May 2005), pp. 279–292. Cited on p. 48

[15] Armitage, P., "An Overlap Problem Arising in Particle Counting", *Biometrika* 36:3/4 (December 1949), pp. 257–266. Cited on p. 4

[16] Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M., "Proof Verification and the Hardness of Approximation Problems", *Journal of the ACM* 45:3 (May 1998), pp. 501–555. Cited on p. 148

[17] Asano, T., "Difficulty of the Maximum Independent Set Problem on Intersection Graphs of Geometric Objects" in Alavi, Y., Chartrand, G., Oellermann, O.R., Schwenk, A.J. (eds.) *Proceedings of the Sixth Quadrennial International Conference on the Theory and Applications of Graphs, Western Michigan University, Kalamazoo, Michigan, May 30-June 3, 1988*, Graph Theory, Combinatorics, and Applications 4, John Wiley & Sons, New York, 1991, pp. 9–18. Cited on p. 2, 44, 68

[18] Åslund, N., "Informationsteoriens Fundamentalsatser, II", *Nordisk Matematisk Tidskrift* 9:3 (1961), pp. 97–108. Cited on p. 72

[19] Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Protasi, M., *Complexity and Approximation – Combinatorial Optimization Problems and Their Approximability Properties*, Springer-Verlag, Berlin, 1999. Cited on p. 11, 12, 44

[20] Bădoiu, M., Har-Peled, S., Indyk, P., "Approximate clustering via core-sets" in *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, May 19 - 21, 2002, Montreal, Quebec, Canada*, Association for Computing Machinery, 2002, pp. 250–257. Cited on p. 163

[21] Bafna, V., Narayanan, B., Ravi, R., "Nonoverlapping local alignments (weighted independent sets of axis-parallel rectangles)", *Discrete Applied Mathematics* 71:1–3 (December 1996), pp. 41–53. Cited on p. 44

[22] Baker, B.S., "Approximation Algorithms for NP-Complete Problems on Planar Graphs", *Journal of the ACM* 41:1 (January 1994), pp. 153–180. Cited on p. 45, 71, 88, 90, 91, 128

[23] Bang-Jensen, J., Reed, B., Schacht, M., Šámal, R., Toft, B., Wagner, U., "Topics in Discrete Mathematics", Algorithms and Combinatorics 26, Springer-Verlag, Berlin, 2006, pp. 613–627. Cited on p. 48

[24] Bar-Yehuda, R., Even, S., "A Linear-Time Approximation Algorithm for the Weighted Vertex Cover Problem", *Journal of Algorithms* 2:2 (June 1981), pp. 198–203. Cited on p. 128

[25] Bar-Yehuda, R., Even, S., "A Local-Ratio Theorem for Approximating the Weighted Vertex Cover Problem" in Ausiello, G., Lucertini, M. (eds.) *Analysis and Design of Algorithms for Combinatorial Problems*, Annals of Discrete Mathematics 25, North-Holland, Amsterdam, 1985, pp. 27–46. Cited on p. 46

[26] Bazgan, C., *Schémas d'approximation et complexité paramétrée*, Rapport de stage de DEA d'Informatique, Université Paris-Sud, Orsay, 1995. Cited on p. 13, 14, 87

[27] Becker, H.W., "Rooks and Rhymes", *Mathematics Magazine* 22:1 (September - October 1948), pp. 23–26. Cited on p. 63

[28] Becker, H.W., "Planar Rhyme Schemes", *Bulletin of the American Mathematical Society* 58 (1952), p. 39. Cited on p. 63

[29] Bell, E.T., "Exponential Polynomials", *Annals of Mathematics* 35:2 (April 1934), pp. 258–277. Cited on p. 62

[30] Bell, E.T., "The Iterated Exponential Integers", *Annals of Mathematics* 39:3 (July 1938), pp. 539–557. Cited on p. 62

[31] Berman, P., "A $d/2$ Approximation for Maximum Weight Independent Set in d-Claw Free Graphs" in Halldórsson, M.M. (ed.) *Algorithm Theory - SWAT 2000, 7th Scandinavian Workshop on Algorithm Theory, Bergen, Norway, July 5-7, 2000, Proceedings*, Lecture Notes in Computer Science 1851, Springer-Verlag, Berlin, 2000, pp. 214–219. Cited on p. 44

[32] Berman, P., DasGupta, B., Muthukrishnan, S., Ramaswami, S., "Efficient Approximation Algorithms for Tiling and Packing Problems with Rectangles", *Journal of Algorithms* 41:2 (November 2001), pp. 443–470. Cited on p. 46

[33] Besicovitch, A.S., "On Crum's problem", *Journal of the London Mathematical Society* 22:4 (October 1947), pp. 285–287. Cited on p. 25

[34] Bilò, V., Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., "Geometric Clustering to Minimize the Sum of Cluster Sizes" in Brodal, G.S., Leonardi, S. (eds.) *Algorithms - ESA 2005, 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, Lecture Notes in Computer Science 3669, Springer-Verlag, Berlin, 2005, pp. 460–471. Cited on p. 163

[35] Blumer, A., Ehrenfeucht, A., Haussler, D., Warmuth, M.K., "Learnability and the Vapnik-Chervonenkis dimension", *Journal of the ACM* 36:4 (October 1989), pp. 929–965. Cited on p. 114

[36] Bodlaender, H.L., "A partial k-arboretum of graphs with bounded treewidth", *Theoretical Computer Science* 209:1–2 (December 1998), pp. 1–45. Cited on p. 50, 70

[37] Bodlaender, H.L., "Treewidth: Characterizations, Applications, and Computations" in Fomin, F.V. (ed.) *Graph-Theoretic Concepts in Computer Science, 32nd International Workshop, WG 2006, Bergen, Norway, June 22-24, 2006, Revised Papers*, Lecture Notes in Computer Science 4271, Springer-Verlag, Berlin, 2006, pp. 1–14. Cited on p. 50

[38] Booth, K.S., Lueker, G.S., "Testing for the Consecutive Ones Property, Interval Graphs, and Graph Planarity Using PQ-Tree Algorithms", *Journal of Computer and System Sciences* 13:3 (December 1976), pp. 335–379. Cited on p. 17, 19

[39] Boppana, R.B., Halldórsson, M.M., "Approximating Maximum Independent Sets by Excluding Subgraphs", *BIT* 32:2 (1992), pp. 180–196. Cited on p. 44

[40] Breu, H., *Algorithmic Aspects of Constrained Unit Disk Graphs*, PhD thesis, The University of British Columbia, 1996. Cited on p. 20

[41] Breu, H., Kirkpatrick, D.G., "On the Complexity of Recognizing Intersection and Touching Graphs of Disks" in Brandenburg, F.J. (ed.) *Graph Drawing, Symposium on Graph Drawing, GD '95, Passau, Germany, September 20-22, 1995, Proceedings*, Lecture Notes in Computer Science 1027, Springer-Verlag, Berlin, 1995, pp. 88–98. Cited on p. 17, 21, 23, 26

[42] Breu, H., Kirkpatrick, D.G., "Unit disk graph recognition is NP-hard", *Computational Geometry* 9:1–2 (January 1998), pp. 3–24. Cited on p. 17, 21, 22, 23

[43] Brightwell, G.R., Scheinerman, E.R., "Representations of Planar Graphs", *SIAM Journal on Discrete Mathematics* 6:2 (1993), pp. 214–229. Cited on p. 27

[44] Brönnimann, H., "Towards Faster Linear-Sized Nets for Axis-Aligned Boxes in the Plane" in Akiyama, J., Kano, M., Tan, X. (eds.) *Discrete and Computational Geometry, Japanese Conference, JCDCG 2004, Tokyo, Japan, October 8-11, 2004, Revised Selected Papers*, Lecture Notes in Computer Science 3742, Springer-Verlag, Berlin, 2005, pp. 54–61. Cited on p. 114

[45] Brönnimann, H., Goodrich, M.T., "Almost Optimal Set Covers in Finite VC-Dimension", *Discrete and Computational Geometry* 14:1 (December 1995), pp. 463–479. Cited on p. 116, 118, 161

[46] Buneman, P., "A characterisation of rigid circuit graphs", *Discrete Mathematics* 9:3 (September 1974), pp. 205–212. Cited on p. 18, 26

[47] Butman, A., Hermelin, D., Lewenstein, M., Rawitz, D., "Optimization problems in multiple-interval graphs" in Bansal, N., Pruhs, K., Stein, C. (eds.) *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, Association for Computing Machinery, 2007, pp. 268–277. Cited on p. 125

[48] Cai, L., Fellows, M.R., Juedes, D.W., Rosamond, F.A., "The Complexity of Polynomial-Time Approximation", *Theory of Computing Systems* 41:3 (October 2007), pp. 459–477. Cited on p. 14

[49] Călinescu, G., Mandoiu, I.I., Wan, P.J., Zelikovsky, A., "Selecting forwarding neighbors in wireless ad hoc networks" in *Proceedings of the 5th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M 2001), Rome, Italy, July 21, 2001*, Association for Computing Machinery, 2001, pp. 34–43. Cited on p. 161

[50] Cannon, A.H., Cowen, L.J., "Approximation Algorithms for the Class Cover Problem", *Annals of Mathematics and Artificial Intelligence* 40:3–4 (March 2004), pp. 215–223. Cited on p. 163

[51] Canny, J.F., "Some Algebraic and Geometric Computations in PSPACE" in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing, 2-4 May 1988, Chicago, Illinois, USA*, Association for Computing Machinery, 1988, pp. 460–467. Cited on p. 23

[52] Carmi, P., Katz, M.J., Lev-Tov, N., "Covering Points by Unit Disks of Fixed Location" in Tokuyama, T. (ed.) *Algorithms and Computation, 18th International Symposium, ISAAC 2007, Sendai, Japan, December 17-19, 2007, Proceedings*, Lecture Notes in Computer Science 4835, Springer-Verlag, Berlin, 2007, pp. 644–655. Cited on p. 161

[53] Cesati, M., Trevisan, L., "On the efficiency of polynomial time approximation schemes", *Information Processing Letters* 64:4 (November 1997), pp. 165–171. Cited on p. 13, 14, 87

[54] Chalermsook, P., Chuzhoy, J., "Maximum Independent Set of Rectangles" in Mathieu, C. (ed.) *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '09)*, Association for Computing Machinery, 2009. Cited on p. 46

[55] Chalopin, J., Gonçalves, D., "Every planar graph is the intersection graph of segments in the plane" in *Proceedings of the Fourty-First Annual ACM Symposium on Theory of Computing, May 31 - June 2, 2009, Bethesda, Maryland, USA*, 2009. Cited on p. 27

[56] Chalopin, J., Gonçalves, D., Ochem, P., "Planar Graphs are in 1-STRING" in Bansal, N., Pruhs, K., Stein, C. (eds.) *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, Association for Computing Machinery, 2007, pp. 609–617. Cited on p. 27

[57] Chan, T.M., "Polynomial-time Approximation Schemes for Packing and Piercing Fat Objects", *Journal of Algorithms* 46:2 (February 2003), pp. 178–189. Cited on p. 45, 92, 104, 164

[58] Chan, T.M., "A Note on Maximum Independent Sets in Rectangle Intersection Graphs", *Information Processing Letters* 89:1 (January 2004), pp. 19–23. Cited on p. 45, 46

[59] Chan, T.M., Har-Peled, S., "Approximation Algorithms for Maximum Independent Set of Pseudo-Disks" in *Proceedings of the 25th ACM Symposium on Computational Geometry, Aarhus, Denmark, June 8-10, 2009*, 2009. Cited on p. 45, 46

[60] Chandra, A.K., Hirschberg, D.S., Wong, C.K., "Approximate Algorithms for Some Generalized Knapsack Problems", *Theoretical Computer Science* 3:3 (December 1976), pp. 293–304. Cited on p. 185

[61] Chang, M.S., "Efficient Algorithms for the Domination Problems on Interval and Circular-Arc Graphs", *SIAM Journal on Computing* 27:6 (1998), pp. 1671–1694. Cited on p. 125

[62] Chazelle, B., Friedman, J., "A Deterministic View of Random Sampling and Its Use in Geometry", *Combinatorica* 10:3 (September 1990), pp. 229–249. Cited on p. 114, 115

[63] Chen, Z., Fu, B., Tang, Y., Zhu, B., "A PTAS for a Disc Covering Problem using Width-Bounded Separators", *Journal of Combinatorial Optimization* 11:2 (March 2006), pp. 203–217. Cited on p. 164

[64] Cheng, X., Huang, X., Li, D., Wu, W., Du, D.Z., "A Polynomial-Time Approximation Scheme for the Minimum-Connected Dominating Set in Ad Hoc Wireless Networks", *Networks* 42:4 (December 2003), pp. 202–208. Cited on p. 47, 82

[65] Chlebík, M., Chlebíková, J., "The Complexity of Combinatorial Optimization Problems on d-Dimensional Boxes", *SIAM Journal on Discrete Mathematics* 21:1 (2007), pp. 158–169. Cited on p. 44, 46, 47, 68, 114, 146, 150

[66] Chvátal, V., "A Greedy Heuristic for the Set-Covering Problem", *Mathematics of Operations Research* 4:3 (1979), pp. 233–235. Cited on p. 46, 113, 119, 129, 161

[67] Clark, B.N., Colbourn, C.J., Johnson, D.S., "Unit Disk Graphs", *Discrete Mathematics* 86:1–3 (December 1990), pp. 165–177. Cited on p. 2, 44, 48, 68, 77, 161, 181

[68] Clarkson, K.L., Varadarajan, K.R., "Improved Approximation Algorithms for Geometric Set Cover", *Discrete and Computational Geometry* 37:1 (January 2007), pp. 43–58. Cited on p. 2, 114, 115, 161

[69] Clementi, A.E.F., Monti, A., Pasquale, F., Silvestri, R., "Optimal Gossiping in Directed Geometric Radio Networks in Presence of Dynamical Faults" in Kucera, L., Kucera, A. (eds.) *Mathematical Foundations of Computer Science 2007, 32nd International Symposium, MFCS 2007, Ceský Krumlov, Czech Republic, August 26-31, 2007, Proceedings*, Lecture Notes in Computer Science 4708, Springer-Verlag, Berlin, 2007, pp. 430–441. Cited on p. 24

[70] Corneil, D.G., Kim, H., Natarajan, S., Olariu, S., Sprague, A.P., "Simple linear time recognition of unit interval graphs", *Information Processing Letters* 55:2 (July 1995), pp. 99–104. Cited on p. 20, 30, 36

[71] Cozzens, M.B., *Higher and multi-dimensional analogues of interval graphs*, PhD thesis, Rutgers University, New Brunswick, New Jersey, 1981. Cited on p. 20

[72] Czyzowicz, J., Dobrev, S., Fevens, T., González-Aguilar, H., Kranakis, E., Opatrny, J., Urrutia, J., "Local Algorithms for Dominating and Connected Dominating Sets of Unit Disk Graphs with Location Aware Nodes" in Laber, E.S., Bornstein, C.F., Nogueira, L.T., Faria, L. (eds.) *LATIN 2008: Theoretical Informatics, 8th Latin American Symposium, Búzios, Brazil, April 7-11, 2008, Proceedings*, Lecture Notes in Computer Science 4957, Springer-Verlag, Berlin, 2008, pp. 158–169. Cited on p. 47

[73] Czyzowicz, J., Kranakis, E., Krizanc, D., Urrutia, J., "Discrete Realizations of Contact and Intersection Graphs" in Battista, G.D. (ed.) *Graph Drawing, 5th International Symposium, GD '97, Rome, Italy, September 18-20, 1997, Proceedings*, Lecture Notes in Computer Science 1353, Springer-Verlag, Berlin, 1997, pp. 359–370. Cited on p. 20, 30, 33, 36

[74] Dai, D., Yu, C., "A $5+\epsilon$-approximation algorithm for minimum weighted dominating set in unit disk graph", *Theoretical Computer Science* 410:8–10 (March 2009), pp. 756–765. Cited on p. 2, 47, 83, 208

[75] Damian, M., Pemmaraju, S.V., "APX-hardness of domination problems in circle graphs", *Information Processing Letters* 97:6 (March 2006), pp. 231–237. Cited on p. 146, 148

[76] Damian-Iordache, M., Pemmaraju, S.V., "A $(2+\epsilon)$-Approximation Scheme for Minimum Domination on Circle Graphs", *Journal of Algorithms* 42:2 (February 2002), pp. 255–276. Cited on p. 146

[77] de Castro, N., Cobos, F.J., Dana, J.C., Márquez, A., Noy, M., "Triangle-Free Planar Graphs and Segment Intersection Graphs", *Journal of Graph Algorithms and Applications* 6:1 (2002), pp. 7–26. Cited on p. 27

[78] de Fraysseix, H., de Mendez, P.O., "Representations by Contact and Intersection of Segments", *Algorithmica* 47:4 (April 2007), pp. 453–463. Cited on p. 27

[79] de Fraysseix, H., de Mendez, P.O., Pach, J., "Representation of planar graphs by segments" in Böröczky, K., Tóth, G.F. (eds.) *Intuitive Geometry: International Conference held by the János Bolyai Mathematical Society, Szeged, Hungary, September 2-7, 1991*, Colloquia Mathematica Societatis János Bolyai 63, János Bolyai Mathematical Society, Budapest, 1994, pp. 109–117. Cited on p. 27, 28

[80] de Fraysseix, H., de Mendez, P.O., Rosenstiehl, P., "On Triangle Contact Graphs", *Combinatorics, Probability and Computing* 3:2 (June 1994), pp. 233–246. Cited on p. 17, 21, 28

[81] de Groot, C., Peikert, R., Würtz, D., *The optimal packing of ten equal circles in a square*, Technical Report IPS Research Report No. 90-12, ETH Zürich, Zürich, Switzerland, 1990. Cited on p. 106

[82] de Morais Cordeiro, C., Agrawal, D.P., *Ad Hoc & Sensor Networks: Theory And Applications*, World Scientific Publishing Company, 2006. Cited on p. 3

[83] Demaine, E.D., Feige, U., Hajiaghayi, M., Salavatipour, M.R., "Combination Can Be Hard: Approximability of the Unique Coverage Problem", *SIAM Journal on Computing* 38:4 (2008), pp. 1464–1483. Cited on p. 159, 160, 162, 199

[84] Demaine, E.D., Hajiaghayi, M.T., "Bidimensionality: new connections between FPT algorithms and PTASs" in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23-25, 2005*, Association for Computing Machinery, 2005, pp. 590–601. Cited on p. 50, 63, 88

[85] Demaine, E.D., Hajiaghayi, M., "Equivalence of Local Treewidth and Linear Local Treewidth and its Algorithmic Applications" in *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '04)*, Association for Computing Machinery, 2004, pp. 840–849. Cited on p. 88, 89

[86] DeWitt, H.K., Krieger, M.M., "Expected Structure of Euclidean Graphs" in Traub, J.F. (ed.) *Proceedings of a Symposium on New Directions and Recent Results in Algorithms and Complexity, held by the Computer Science Department Carnegie-Mellon University, April 7-9, 1976*, Academic Press, New York, 1976, p. 451. Cited on p. 22

[87] Dinur, I., Safra, S., "The importance of being biased" in *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, May 19 - 21, 2002, Montreal, Quebec, Canada*, Association for Computing Machinery, 2002, pp. 33–42. Cited on p. 46

[88] Dobiński, G., "Summirung der Reihe $\sum \frac{n^m}{n!}$ für $m = 1, 2, 3, 4, 5, \ldots$", *Archiv der Mathematik und Physik* 61 (1877), pp. 333–336. Cited on p. 62

[89] Dorn, F., "Dynamic Programming and Fast Matrix Multiplication" in Azar, Y., Erlebach, T. (eds.) *Algorithms - ESA 2006, 14th Annual European Symposium, Zurich, Switzerland, September 11-13, 2006, Proceedings*, Lecture Notes in Computer Science 4168, Springer-Verlag, Berlin, 2006, pp. 280–291. Cited on p. 50, 90

[90] Dorn, F., Fomin, F.V., Thilikos, D.M., "Subexponential Parameterized Algorithms" in Arge, L., Cachin, C., Jurdzinski, T., Tarlecki, A. (eds.)

Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007, Proceedings, Lecture Notes in Computer Science 4596, Springer-Verlag, Berlin, 2007, pp. 15–27. Cited on p. 63

[91] Dorn, F., Fomin, F.V., Thilikos, D.M., "Catalan structures and dynamic programming in H-minor-free graphs" in Teng, S.H. (ed.) *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, Association for Computing Machinery, 2008, pp. 631–640. Cited on p. 63, 89

[92] Downey, R.G., Fellows, M.R., *Parameterized Complexity*, Springer, New York, 1999. Cited on p. 14

[93] Du, D.Z., Thai, M.T., Li, Y., Liu, D., Zhu, S., "Strongly Connected Dominating Sets in Wireless Sensor Networks with Unidirectional Links" in Zhou, X., Li, J., Shen, H.T., Kitsuregawa, M., Zhang, Y. (eds.) *Frontiers of WWW Research and Development - APWeb 2006, 8th Asia-Pacific Web Conference, Harbin, China, January 16-18, 2006, Proceedings*, Lecture Notes in Computer Science 3841, Springer-Verlag, Berlin, 2006, pp. 13–24. Cited on p. 24

[94] Du, H., Jia, X., Li, D., Wu, W., "Coloring of Double Disk Graphs", *Journal of Global Optimization* 28:2 (January 2004), pp. 115–119. Cited on p. 24

[95] Duchet, P., "Classical perfect graphs" in Berge, C., Chvátal, V. (eds.) *Topics on Perfect Graphs*, Annals of Discrete Mathematics 21, North-Holland, Amsterdam, 1984, pp. 67–96. Cited on p. 26

[96] Efrat, A., Hoffmann, F., Knauer, C., Kriegel, K., Rote, G., Wenk, C., "Covering with Ellipses", *Algorithmica* 38:1 (October 2003), pp. 145–160. Cited on p. 163

[97] Efrat, A., Sharir, M., "The Complexity of the Union of Fat Objects in the Plane", *Discrete and Computational Geometry* 23:2 (February 2000), pp. 171–189. Cited on p. 83, 104, 149

[98] Ehrlich, G., Even, S., Tarjan, R.E., "Intersection Graphs of Curves in the Plane", *Journal of Combinatorial Theory, Series B* 21:1 (August 1976), pp. 8–20. Cited on p. 20, 25

[99] Eppstein, D., "Diameter and Treewidth in Minor-Closed Graph Families", *Algorithmica* 27:3 (June 2000), pp. 275–291. Cited on p. 88, 89, 90

[100] Eppstein, D., Miller, G.L., Teng, S.H., "A Deterministic Linear Time Algorithm for Geometric Separators and its Applications", *Fundamenta Informaticae* 22:4 (April 1995), pp. 309–329. Cited on p. 99

[101] Erlebach, T., Fiala, J., "Independence and Coloring Problems on Intersection Graphs of Disks" in Erlebach, T., Kaklamanis, C. (eds.) *Approximation and Online Algorithms, 4th International Workshop, WAOA 2006, Zurich, Switzerland, September 14-15, 2006, Revised Papers*, Lecture Notes in Computer Science 4368, Springer-Verlag, Berlin, 2006, pp. 135–155. Cited on p. 48

[102] Erlebach, T., Hagerup, T., Jansen, K., Minzlaff, M., Wolff, A., "Trimming of Graphs, with Application to Point Labeling" in Albers, S., Weil, P. (eds.) *STACS 2008, 25th Annual Symposium on Theoretical Aspects of Computer Science, Bordeaux, France, February 21-23, 2008, Proceedings*, Dagstuhl Seminar Series 08001, Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2008, pp. 265–276. Cited on p. 4

[103] Erlebach, T., Jansen, K., Seidel, E., "Polynomial-time Approximation Schemes for Geometric Intersection Graphs", *SIAM Journal on Computing* 34:6 (2005), pp. 1302–1323. Cited on p. 2, 45, 46, 91, 92, 99, 104, 113

[104] Erlebach, T., van Leeuwen, E.J., "Approximating Geometric Coverage Problems" in Teng, S.H. (ed.) *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, Association for Computing Machinery, 2008, pp. 1267–1276. Cited on p. 7

[105] Erlebach, T., van Leeuwen, E.J., "Domination in Geometric Intersection Graphs" in Laber, E.S., Bornstein, C.F., Nogueira, L.T., Faria, L. (eds.) *LATIN 2008: Theoretical Informatics, 8th Latin American Symposium, Búzios, Brazil, April 7-11, 2008, Proceedings*, Lecture Notes in Computer Science 4957, Springer-Verlag, Berlin, 2008, pp. 747–758. Cited on p. 6

[106] Even, G., Lotker, Z., Ron, D., Smorodinsky, S., "Conflict-Free Colorings of Simple Geometric Regions with Applications to Frequency Assignment in Cellular Networks", *SIAM Journal on Computing* 33:1 (2003), pp. 94–136. Cited on p. 48

[107] Even, G., Rawitz, D., Shahar, S.M., "Hitting Sets When the VC-Dimension is Small", *Information Processing Letters* 95:2 (July 2005), pp. 358–362. Cited on p. 118

[108] Feige, U., "A Threshold of $\ln n$ for Approximating Set Cover", *Journal of the ACM* 45:4 (July 1998), pp. 634–652. Cited on p. 46, 113, 146, 149, 161

[109] Fishburn, P.C., "On the sphericity and cubicity of graphs", *Journal of Combinatorial Theory, Series B* 35:3 (December 1983), pp. 309–318. Cited on p. 22

[110] Fowler, R.J., Paterson, M.S., Tanimoto, S.L., "Optimal Packing and Covering in the Plane are NP-Complete", *Information Processing Letters* 12:3 (June 1981), pp. 133–137. Cited on p. 2, 161, 163

[111] Frank, A., Gyárfás, A., "How to orient the edges of a graph?" in Hajnal, A., Sós, V.T. (eds.) *Combinatorics, Proceedings of the 5th Hungarian Combinatorial Colloquium, Keszthely, Hungary, June 28-July 3, 1976*, Colloquia Mathematica Societatis János Bolyai 18, North-Holland, Amsterdam, 1978, pp. 353–364. Cited on p. 129

[112] Fu, B., "Theory and Application of Width Bounded Geometric Seperator" in Durand, B., Thomas, W. (eds.) *STACS 2006, 23rd Annual Symposium on Theoretical Aspects of Computer Science, Marseille, France, February 23-25, 2006, Proceedings*, Lecture Notes in Computer Science 3884, Springer-Verlag, Berlin, 2006, pp. 277–288. Cited on p. 70

[113] Gandhi, R., Khuller, S., Srinivasan, A., "Approximation algorithms for partial covering problems", *Journal of Algorithms* 53:1 (October 2004), pp. 55–84. Cited on p. 163

[114] Garey, M.R., Johnson, D.S., "The Rectilinear Steiner Tree Problem is NP-Complete", *SIAM Journal on Applied Mathematics* 32:4 (1977), pp. 826–834. Cited on p. 181

[115] Garey, M.R., Johnson, D.S., *Computers and Intractability - A Guide to the Theory of NP-Completeness*, W.H. Freeman and Company, San Francisco, 1979. Cited on p. 12, 43, 44

[116] Garey, M.R., Johnson, D.S., Stockmeyer, L.J., "Some Simplified NP-Complete Graph Problems", *Theoretical Computer Science* 1:3 (February 1976), pp. 237–267. Cited on p. 201

[117] Garwood, F., "The Variance of the Overlap of Geometrical Figures with Reference to a Bombing Problem", *Biometrika* 34:1/2 (January 1947), pp. 1–17. Cited on p. 4

[118] Gavril, F., "The intersection graphs of subtrees in trees are exactly the chordal graphs", *Journal of Combinatorial Theory, Series B* 16:1 (February 1974), pp. 47–56. Cited on p. 18, 26

[119] Geelen, J., Guo, A., McKinnon, D., "Straight line embeddings of cubic planar graphs with integer edge lengths", *Journal of Graph Theory* 58:3 (July 2008), pp. 270–274. Cited on p. 27

[120] Gens, G.V., Levner, E.V., "Computational Complexity of Approximation Algorithms for Combinatorial Problems" in Becvár, J. (ed.) *Mathematical Foundations of Computer Science 1979, Proceedings, 8th Symposium, Olomouc, Czechoslovakia, September 3-7, 1979*, Lecture Notes in Computer Science 74, Springer-Verlag, Berlin, 1979, pp. 292–300. Cited on p. 185

[121] Ghosh, A., Das, S.K., "Coverage and connectivity issues in wireless sensor networks: A survey", *Pervasive and Mobile Computing* 4:3 (June 2008), pp. 303–334. Cited on p. 3

[122] Gilbert, E.N., "Random Plane Networks", *Journal of the Society for Industrial and Applied Mathematics* 9:4 (December 1961), pp. 533–543. Cited on p. 2

[123] Gilmore, P.C., Hoffman, A.J., "A characterization of comparability graphs and of interval graphs", *Canadian Journal of Mathematics* 16 (1964), pp. 539–548. Cited on p. 19

[124] Glaßer, C., Reith, S., Vollmer, H., "The complexity of base station positioning in cellular networks", *Discrete Applied Mathematics* 148:1 (April 2005), pp. 1–12. Cited on p. 3, 164

[125] Glaßer, C., Reitwießner, C., Schmitz, H., "Multiobjective Disk Cover Admits a PTAS" in Hong, S.H., Nagamochi, H., Fukunaga, T. (eds.) *Algorithms and Computation, 19th International Symposium, ISAAC 2008, Gold Coast, Australia, December 15-17, 2008, Proceedings*, Lecture Notes in Computer Science 5369, Springer-Verlag, Berlin, 2008, pp. 40–51. Cited on p. 162, 163

[126] Golumbic, M.C., *Algorithmic Graph Theory and Perfect Graphs*, Second Edition, 57, Elsevier, 2004. Cited on p. 19

[127] Golumbic, M.C., Trenk, A.N., *Tolerance Graphs*, Cambridge University Press, Cambridge, 2004. Cited on p. 19

[128] Gonzalez, T.F., "Covering a Set of Points in Multidimensional Space", *Information Processing Letters* 40:4 (November 1991), pp. 181–188. Cited on p. 163

[129] Gräf, A., Stumpf, M., Weißenfels, G., *On Coloring Unit Disk Graphs*, Technical Report Musikinformatik & Medientechnik 17/94, Johannes Gutenberg-Universität Mainz, Mainz, Germany, 1994. Cited on p. 24

[130] Gräf, A., Stumpf, M., Weißenfels, G., "On Coloring Unit Disk Graphs", *Algorithmica* 20:3 (March 1998), pp. 277–293. Cited on p. 48

[131] Grohe, M., "Local Tree-Width, Excluded Minors, and Approximation Algorithms", *Combinatorica* 23:4 (December 2003), pp. 613–632. Cited on p. 88, 90

[132] Guha, S., Khuller, S., "Approximation Algorithms for Connected Dominating Sets", *Algorithmica* 20:4 (April 1998), pp. 374–387. Cited on p. 46

[133] Guo, S., Yang, O.W.W., "Energy-aware multicasting in wireless ad hoc networks: A survey and discussion", *Computer Communications* 30:9 (June 2007), pp. 2129–2148. Cited on p. 3

[134] Guruswami, V., Trevisan, L., "The Complexity of Making Unique Choices: Approximating 1-in-k SAT" in Chekuri, C., Jansen, K., Rolim, J.D.P., Trevisan, L. (eds.) *Approximation, Randomization and Combinatorial Optimization, Algorithms and Techniques, 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2005 and 9th International Workshop on Randomization and Computation, RANDOM 2005, Berkeley, CA, USA, August 22-24, 2005, Proceedings*, Lecture Notes in Computer Science 3624, Springer-Verlag, Berlin, 2005, pp. 99–110. Cited on p. 162

[135] Gyárfás, A., Lehel, J., "A Helly-type problem in trees" in Erdös, P., Rényi, A., Sós, V.T. (eds.) *Combinatorial Theory and its Applications II*, Colloquia Mathematica Societatis János Bolyai 4, North-Holland, Amsterdam, 1970, pp. 571–584. Cited on p. 18, 26

[136] Hale, W.K., "Frequency Assignment: Theory and Applications", *Proceedings of the IEEE* 68:12 (December 1980), pp. 1497–1514. Cited on p. 3

[137] Halldórsson, M.M., "Approximating Discrete Collections via Local Improvements" in *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1995, San Francisco, California, USA, January 22-24, 1995*, Association for Computing Machinery, 1995, pp. 160–169. Cited on p. 44

[138] Halperin, E., "Improved Approximation Algorithms for the Vertex Cover Problem in Graphs and Hypergraphs", *SIAM Journal on Computing* 31:5 (2002), pp. 1608–1623. Cited on p. 46

[139] Hansen, L.J., "On the Rodin and Sullivan Ring Lemma", *Complex Variables: Theory and Application* 10 (1988), pp. 23–30. Cited on p. 26

[140] Har-Peled, S., Smorodinsky, S., "Conflict-Free Coloring of Points and Simple Regions in the Plane", *Discrete and Computational Geometry* 34:1 (July 2005), pp. 47–70. Cited on p. 48

[141] Hartman, I.B.A., Newman, I., Ziv, R., "On grid intersection graphs", *Discrete Mathematics* 87:1 (January 1991), pp. 41–52. Cited on p. 27, 28

[142] Håstad, J., "Clique is hard to approximate within $n^{1-\epsilon}$", *Acta Mathematica* 182:1 (March 1999), pp. 105–142. Cited on p. 44

[143] Haussler, D., Welzl, E., "ϵ-Nets and Simplex Range Queries", *Discrete and Computational Geometry* 2:1 (December 1987), pp. 127–151. Cited on p. 114

[144] Havel, T.F., *The Combinatorial Distance Geometry Approach to the Calculation of Molecular Conformation*, PhD thesis, University of California, Berkeley, 1982. Cited on p. 22

[145] Havel, T.F., Kuntz, I.D., Crippen, G.M., "The Combinatorial Distance Geometry Method for the Calculation of Molecular Conformation I: A New Approach to an Old Problem", *Journal of Theoretical Biology* 104:3 (October 1983), pp. 359–381. Cited on p. 22

[146] Hliněný, P., "Touching Graphs of Unit Balls" in Battista, G.D. (ed.) *Graph Drawing, 5th International Symposium, GD '97, Rome, Italy, September 18-20, 1997, Proceedings*, Lecture Notes in Computer Science 1353, Springer-Verlag, Berlin, 1997, pp. 350–358. Cited on p. 23

[147] Hliněný, P., Kratochvíl, J., "Representing graphs by disks and balls (a survey of recognition-complexity results)", *Discrete Mathematics* 229:1-3 (February 2001), pp. 101–124. Cited on p. 23, 83

[148] Hochbaum, D.S., "Approximation Algorithms for the Set Covering and Vertex Cover Problems", *SIAM Journal on Computing* 11:3 (1982), pp. 555–556. Cited on p. 127

[149] Hochbaum, D.S., "Approximating Covering and Packing Problems: Set Cover, Vertex Cover, Independent Set, and Related Problems" in Hochbaum, D.S. (ed.) *Approximation Algorithms for NP-hard Problems*, PWS Publishing Company, Boston, 1997, pp. 46–93. Cited on p. 119, 129, 161, 162

[150] Hochbaum, D.S., Maass, W., "Approximation Schemes for Covering and Packing Problems in Image Processing and VLSI", *Journal of the ACM* 32:1 (January 1985), pp. 130–136. Cited on p. 2, 4, 45, 71, 92, 163, 195

[151] Hochbaum, D.S., Pathria, A., *Analysis of the greedy approach in covering problems*, unpublished manuscript, 1994. Cited on p. 162

[152] Hopcroft, J., Tarjan, R.E., "Efficient Planarity Testing", *Journal of the ACM* 21:4 (October 1974), pp. 549–568. Cited on p. 23

[153] Huang, Y., Gao, X., Zhang, Z., Wu, W., "A Better Constant-Factor Approximation for Weighted Dominating Set in Unit Disk Graph", *Journal of Combinatorial Optimization* (2008), doi: 10.1007/s10878-008-9146-0. Cited on p. 47, 83

[154] Hunt III, H.B., Marathe, M.V., Radhakrisnan, V., Ravi, S.S., Rosenkrantz, D.J., Stearns, R.E., "NC-Approximation Schemes for NP- and PSPACE-Hard Problems for Geometric Graphs", *Journal of Algorithms* 26:2 (February 1998), pp. 238–274. Cited on p. 2, 45, 46, 47, 68, 71, 74, 76, 78, 79, 84

[155] Imai, H., Asano, T., "Finding the Connected Components and a Maximum Clique of an Intersection Graph of Rectangles in the Plane", *Journal of Algorithms* 4:4 (December 1983), pp. 310–323. Cited on p. 48

[156] Johnson, D.S., "Approximation algorithms for combinatorial problems", *Journal of Computer and System Sciences* 9:3 (December 1974), pp. 256–278. Cited on p. 46, 113, 119, 129, 161

[157] Johnson, D.S., "The NP-Completeness Column: An Ongoing Guide", *Journal of Algorithms* 3:2 (June 1982), pp. 182–195. Cited on p. 2, 161, 163

[158] Kalinin, V.B., "A Problem of Berge", *Mathematical Notes* 34:1 (July 1983), pp. 551–552, (translated from Matematicheskie Zametki, 34:1 (July 1983), pp. 131–133). Cited on p. 25

[159] Karmarkar, N., "A new polynomial-time algorithm for linear programming", *Combinatorica* 4:4 (December 1984), pp. 373–396. Cited on p. 118

[160] Kaufmann, M., Kratochvíl, J., Lehmann, K.A., Subramanian, A.R., "Max-tolerance graphs as intersection graphs: cliques, cycles, and recognition" in *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, Association for Computing Machinery, 2006, pp. 832–841. Cited on p. 4, 21, 48, 122

[161] Kedem, K., Livne, R., Pach, J., Sharir, M., "On the Union of Jordan Regions and Collision-Free Translational Motion Amidst Polygonal Obstacles", *Discrete and Computational Geometry* 1:1 (December 1986), pp. 59–70. Cited on p. 115

[162] Kelmans, A.K., "On convex embeddings of planar 3-connected graphs", *Journal of Graph Theory* 33:2 (February 2000), pp. 120–124. Cited on p. 27

[163] Khachiyan, L.G., “A Polynomial Algorithm in Linear Programming”, *Soviet Mathematics Doklady* 20:1 (1979), pp. 191–194, (translated from Doklady Akademii Nauk SSSR 244 (1979), pp. 1093–1096). Cited on p. 118

[164] Khanna, S., Muthukrishnan, S., Paterson, M., “On Approximating Rectangle Tiling and Packing” in *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1998, San Francisco, California, United States, January 25-27, 1998*, Association for Computing Machinery, 1998, pp. 384–393. Cited on p. 46

[165] Khuller, S., Moss, A., Naor, J.S., “The Budgeted Maximum Coverage Problem”, *Information Processing Letters* 70:1 (April 1999), pp. 39–45. Cited on p. 162, 179

[166] Kiess, W., Mauve, M., “A survey on real-world implementations of mobile ad-hoc networks”, *Ad Hoc Networks* 5:3 (April 2007), pp. 324–339. Cited on p. 3

[167] Kim, S.J., Kostochka, A.V., Nakprasit, K., “On the Chromatic Number of Intersection Graphs of Convex Sets in the Plane”, *The Electronic Journal of Combinatorics* 11:1 (August 2004), p. R52. Cited on p. 48, 124

[168] Kloks, T., *Treewidth: Computation and Approximation*, Lecture Notes in Computer Science 842, Springer-Verlag, Berlin, 1994. Cited on p. 52, 56, 59

[169] Koebe, P., “Kontaktprobleme der konformen Abbildung”, *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften Leipzig, Mathematische-Physische Klasse* Sitzung vom 11. Februar 1935, 88 (1936), pp. 141–164. Cited on p. 17, 23, 26, 91, 128

[170] Komlós, J., Pach, J., Woeginger, G.J., “Almost Tight Bounds for ϵ-Nets”, *Discrete and Computational Geometry* 7:2 (December 1992), pp. 163–173. Cited on p. 114

[171] Kozyrev, V.P., Yushmanov, S.V., “Representations of graphs and networks (coding, layouts and embeddings)”, *Journal of Mathematical Sciences* 61:3 (September 1992), pp. 2152–2194, (translated from Itogi Nauki i Tekhniki, Seriya Teoriya Veroyatnostei, Matematicheskaya Statistika, Teoreticheskaya Kibernetika, 27 (1990), pp. 129–196). Cited on p. 18, 19

[172] Kratochvíl, J., “String graphs. II. recognizing string graphs is NP-hard”, *Journal of Combinatorial Theory, Series B* 52:1 (May 1991), pp. 67–78. Cited on p. 20

[173] Kratochvíl, J., "A Special Planar Satisfiability Problem and a Consequence of Its NP-completeness", *Discrete Applied Mathematics* 52:3 (August 1994), pp. 233–252. Cited on p. 19, 20

[174] Kratochvíl, J., "Intersection Graphs of Noncrossing Arc-Connected Sets in the Plane" in North, S.C. (ed.) *Graph Drawing, Symposium on Graph Drawing, GD '96, Berkeley, California, USA, September 18-20, Proceedings*, Lecture Notes in Computer Science 1190, Springer-Verlag, Berlin, 1996, pp. 257–270. Cited on p. 17, 21, 23, 27

[175] Kratochvíl, J., *Geometric representations of graphs*, Course notes, Graph Theory Course at UPC Barcelona, April 5-21, 2005. Cited on p. 20

[176] Kratochvíl, J., "Geometric Representations of Graphs" in *New Directions in Algorithms, Combinatorics, and Optimization, A Conference Honoring the 65th Birthday of William T. Trotter, Atlanta, Georgia, USA, May 5-9, 2008*, 2008. Cited on p. 21, 28

[177] Kratochvíl, J., Kuběna, A., "On intersection representations of co-planar graphs", *Discrete Mathematics* 178:1–3 (January 1998), pp. 251–255. Cited on p. 28

[178] Kratochvíl, J., Matoušek, J., "NP-Hardness Results for Intersection Graphs", *Commentationes Mathematicae Universitatis Carolinae* 30:4 (1989), pp. 761–773. Cited on p. 21

[179] Kratochvíl, J., Matoušek, J., "Intersection Graphs of Segments", *Journal of Combinatorial Theory, Series B* 62:2 (November 1994), pp. 289–315. Cited on p. 19, 20, 23

[180] Kratochvíl, J., Pergel, M., "Two Results on Intersection Graphs of Polygons" in Liotta, G. (ed.) *Graph Drawing, 11th International Symposium, GD 2003, Perugia, Italy, September 21-24, 2003, Revised Papers*, Lecture Notes in Computer Science 2912, Springer-Verlag, Berlin, 2003, pp. 59–70. Cited on p. 21

[181] Kratochvíl, J., Pergel, M., "Geometric Intersection Graphs: Do Short Cycles Help?" in Lin, G. (ed.) *Computing and Combinatorics, 13th Annual International Conference, COCOON 2007, Banff, Canada, July 16-19, 2007, Proceedings*, Lecture Notes in Computer Science 4598, Springer-Verlag, Berlin, 2007, pp. 118–128. Cited on p. 21

[182] Kratochvíl, J., Pergel, M., "Intersection graphs of homothetic polygons", *Electronic Notes in Discrete Mathematics* 31 (August 2008), pp. 277–280. Cited on p. 21

[183] Kreweras, G., "Sur les partitions non croisees d'un cycle", *Discrete Mathematics* 1:4 (February 1972), pp. 333–350. Cited on p. 63

[184] Kuhn, F., Moscibroda, T., Wattenhofer, R., "Unit Disk Graph Approximation" in Basagni, S., Phillips, C.A. (eds.) *Proceedings of the 2004 Joint Workshop on Foundations of Mobile Computing, DIALM-POMC '04, Philadelphia, PA, USA, 2004*, Association for Computing Machinery, 2004, pp. 17–23. Cited on p. 17, 22, 24

[185] Kuhn, F., Nieberg, T., Moscibroda, T., Wattenhofer, R., "Local Approximation Schemes for Ad Hoc and Sensor Networks" in *Proceedings of the 2005 Joint Workshop on Foundations of Mobile Computing, DIALM-POMC '05, Cologne, Germany, 2005*, Association for Computing Machinery, 2005, pp. 97–103. Cited on p. 47

[186] Kuhn, F., von Rickenbach, P., Wattenhofer, R., Welzl, E., Zollinger, A., "Interference in Cellular Networks: The Minimum Membership Set Cover Problem" in Wang, L. (ed.) *Computing and Combinatorics, 11th Annual International Conference, COCOON 2005, Kunming, China, August 16-29, 2005, Proceedings*, Lecture Notes in Computer Science 3595, Springer-Verlag, Berlin, 2005, pp. 188–198. Cited on p. 163, 200

[187] Kuhn, F., Wattenhofer, R., Zollinger, A., "Ad hoc networks beyond unit disk graphs", *Wireless Networks* 14:5 (October 2008), pp. 715–729. Cited on p. 24

[188] Laue, S., "Geometric Set Cover and Hitting Sets for Polytopes in $\mathbb{R}^3$" in Albers, S., Weil, P. (eds.) *STACS 2008, 25th Annual Symposium on Theoretical Aspects of Computer Science, Bordeaux, France, February 21-23, 2008, Proceedings*, Dagstuhl Seminar Series 08001, Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2008, pp. 479–490. Cited on p. 114, 115, 116, 126, 162, 179

[189] Lekkerkerker, C.G., Boland, J.C., "Representation of a finite graph by a set of intervals on the real line", *Fundamenta Informaticae* 51 (1962), pp. 45–62. Cited on p. 18

[190] Lev-Tov, N., Peleg, D., "Exact Algorithms and Approximation Schemes for Base Station Placement Problems" in Penttonen, M., Schmidt, E.M. (eds.) *Algorithm Theory - SWAT 2002, 8th Scandinavian Workshop on Algorithm Theory, Turku, Finland, July 3-5, 2002 Proceedings*, Lecture Notes in Computer Science 2368, Springer-Verlag, Berlin, 2002, pp. 90–99. Cited on p. 163

[191] Lev-Tov, N., Peleg, D., "Polynomial time approximation schemes for base station coverage with minimum total radii", *Computer Networks* 47:4 (March 2005), pp. 489–501. Cited on p. 161, 163

[192] Li, X.Y., Wang, Y., "Simple approximation algorithms and PTASs for various problems in wireless ad hoc networks", *Journal of Parallel and*

Distributed Computing 66:4 (April 2006), pp. 515–530. Cited on p. 45, 46

[193] Liao, C., Hu, S., "Polynomial time approximation schemes for minimum disk cover problems", *Journal of Global Optimization* (2009), doi: 10.1007/s10878-009-9216-y. Cited on p. 161

[194] Lichtenstein, D., "Planar Formulae and Their Uses", *SIAM Journal on Computing* 11:2 (May 1982), pp. 329–343. Cited on p. 68

[195] Lin, M.C., Szwarcfiter, J.L., "Unit Circular-Arc Graph Representations and Feasible Circulations", *SIAM Journal on Discrete Mathematics* 22:1 (2008), pp. 409–423. Cited on p. 19

[196] Lingas, A., Wahlen, M., "A note on maximum independent set and related problems on box graphs", *Information Processing Letters* 93:4 (February 2005), pp. 169–171. Cited on p. 20, 30

[197] Lovász, L., "On the ratio of optimal integral and fractional covers", *Discrete Mathematics* 13:4 (1975), pp. 383–390. Cited on p. 46, 113, 119, 129, 161

[198] Maehara, H., "Space graphs and sphericity", *Discrete Applied Mathematics* 7:1 (January 1984), pp. 55–64. Cited on p. 22, 83

[199] Malesińska, E., *Graph-Theoretical Models for Frequency Assignment Problems*, PhD thesis, Technical University of Berlin, Berlin, Germany, 1997. Cited on p. 22, 24, 27, 46, 48

[200] Malitz, S.M., Papakostas, A., "On the Angular Resolution of Planar Graphs", *SIAM Journal on Discrete Mathematics* 7:2 (1994), pp. 173–183. Cited on p. 26

[201] Marathe, M.V., Breu, H., Hunt III, H.B., Ravi, S.S., Rosenkrantz, D.J., "Simple Heuristics for Unit Disk Graphs", *Networks* 25 (1995), pp. 59–68. Cited on p. 2, 44, 46, 48

[202] Marx, D., "Efficient Approximation Schemes for Geometric Problems?" in Brodal, G.S., Leonardi, S. (eds.) *Algorithms - ESA 2005, 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, Lecture Notes in Computer Science 3669, Springer-Verlag, Berlin, 2005, pp. 448–459. Cited on p. 46, 70, 76, 77, 87, 88, 99

[203] Marx, D., "Parameterized Complexity of Independence and Domination Problems on Geometric Graphs" in Bodlaender, H.L., Langston, M.A. (eds.) *Parameterized and Exact Computation, Second International Workshop, IWPEC 2006, Zürich, Switzerland, September 13-15, 2006, Proceedings*, Lecture Notes in Computer Science 4169, Springer-Verlag, Berlin, 2006, pp. 154–165. Cited on p. 70, 87, 88

[204] Marx, D., "On the Optimality of Planar and Geometric Approximation Schemes" in *48th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2007, October 20-23, 2007, Providence, RI, USA, Proceedings*, IEEE Computer Society, 2007, pp. 338–348. Cited on p. 84, 85, 86, 88, 90

[205] Matoušek, J., Seidel, R., Welzl, E., "How to Net a Lot with Little: Small ϵ-Nets for Disks and Halfspaces" in *Proceedings of the 6th Annual Symposium on Computational Geometry, Berkley, California, USA, 1990*, Association for Computing Machinery, 1990. Cited on p. 114

[206] Matsui, T., "Approximation Algorithms for Maximum Independent Set Problems and Fractional Coloring Problems on Unit Disk Graph" in Akiyama, J., Kano, M., Urabe, M. (eds.) *Discrete and Computational Geometry, Japanese Conference, JCDCG'98, Tokyo, Japan, December 9-12, 1998, Revised Papers*, Lecture Notes in Computer Science 1763, Springer-Verlag, Berlin, 1998, pp. 194–200. Cited on p. 45

[207] McConnell, R.M., "Linear-Time Recognition of Circular-Arc Graphs", *Algorithmica* 37:2 (October 2003), pp. 93–147. Cited on p. 19

[208] McKee, T.A., McMorris, F.R., *Topics in Intersection Graph Theory*, Society for Industrial and Applied Mathematics, Philadelphia, 1999. Cited on p. 18, 19

[209] Mihalák, M., *Optimization Problems in Communication Networks*, PhD thesis, University of Leicester, Leicester, UK, 2006. Cited on p. 161, 165, 177, 187

[210] Miller, G.L., Teng, S.H., Thurston, W.P., Vavasis, S.A., "Separators for Sphere-Packings and Nearest Neighbor Graphs", *Journal of the ACM* 44:1 (January 1992), pp. 1–29. Cited on p. 91, 127, 128

[211] Mohar, B., "A polynomial time circle packing algorithm", *Discrete Mathematics* 117:1–3 (July 1993), pp. 257–263. Cited on p. 27

[212] Mohar, B., "Circle Packings of Maps in Polynomial Time", *European Journal of Combinatorics* 18:7 (October 1997), pp. 785–805. Cited on p. 27

[213] Mohar, B., Thomassen, C., *Graphs on Surfaces*, The John Hopkins University Press, Baltimore, 2001. Cited on p. 27

[214] Monien, B., Speckenmeyer, E., "Ramsey Numbers and an Approximation Algorithm for the Vertex Cover Problem", *Acta Informatica* 22:1 (April 1985), pp. 115–123. Cited on p. 46

[215] Moscibroda, T., O'Dell, R., Wattenhofer, M., Wattenhofer, R., "Virtual Coordinates for Ad Hoc and Sensor Networks" in Basagni, S., Phillips, C.A. (eds.) *Proceedings of the 2004 Joint Workshop on Foundations of Mobile Computing, DIALM-POMC '04, Philadelphia, PA, USA, 2004*, Association for Computing Machinery, 2004, pp. 8–16. Cited on p. 22

[216] Moser, H., Raman, V., Sikdar, S., "The Parameterized Complexity of the Unique Coverage Problem" in Tokuyama, T. (ed.) *Algorithms and Computation, 18th International Symposium, ISAAC 2007, Sendai, Japan, December 17-19, 2007, Proceedings*, Lecture Notes in Computer Science 4835, Springer-Verlag, Berlin, 2007, pp. 621–631. Cited on p. 162

[217] Mustafa, N., Ray, S., "PTAS for Geometric Hitting Set Problems via Local Search" in *Proceedings of the 25th ACM Symposium on Computational Geometry, Aarhus, Denmark, June 8-10, 2009*, 2009. Cited on p. 161, 165, 208

[218] Narayanappa, S., Vojtěchovský, P., "An Improved Approximation Factor For The Unit Disk Covering Problem" in *Proceedings of the 18th Annual Canadian Conference on Computational Geometry, CCCG 2006, August 14-16, 2006, Queen's University, Ontario, Canada*, 2006. Cited on p. 161

[219] Nieberg, T., Hurink, J., Kern, W., "Approximation schemes for wireless networks", *ACM Transactions on Algorithms* 4:4 (2008), pp. 49:1–49:17. Cited on p. 24, 45, 46, 47, 84

[220] Papadimitriou, C.H., Yannakakis, M., "Optimization, Approximation, and Complexity Classes", *Journal of Computer and System Sciences* 43:3 (December 1991), pp. 425–440. Cited on p. 150

[221] Paz, A., Moran, S., "Non deterministic polynomial optimization problems and their approximations", *Theoretical Computer Science* 15:3 (1981), pp. 251–277. Cited on p. 12

[222] Peeters, R., *On Coloring j-Unit Sphere Graphs*, Technical Report FEW 512, Department of Economics, Tilburg University, Tilburg, The Netherlands, 1991. Cited on p. 48

[223] Peikert, R., Würtz, D., Monagan, M., de Groot, C., "Packing Circles in a Square: A Review and New Results" in Kall, P. (ed.) *System Modelling and Optimization, Proceedings of the 15th IFIP Conference Zurich, Switzerland, September 2-6, 1991*, Lecture Notes in Control and Information Sciences 180, Springer-Verlag, Berlin, 1991, pp. 45–54. Cited on p. 106

[224] Pergel, M., "Recognition of Polygon-Circle Graphs and Graphs of Interval Filaments Is NP-Complete" in Brandstädt, A., Kratsch, D., Müller,

H. (eds.) *Graph-Theoretic Concepts in Computer Science, 33rd International Workshop, WG 2007, Dornburg, Germany, June 21-23, 2007, Revised Papers*, Lecture Notes in Computer Science 4769, Springer-Verlag, Berlin, 2007, pp. 238–247. Cited on p. 21

[225] Perkins, C.E., *Ad Hoc Networking*, Addison Wesley, Reading, Massachussets, 2001. Cited on p. 3

[226] Pyrga, E., Ray, S., "New Existence Proofs for ϵ-nets" in Teillaud, M. (ed.) *Proceedings of the 24th ACM Symposium on Computational Geometry, College Park, MD, USA, June 9-11, 2008*, Association for Computing Machinery, 2008, pp. 199–207. Cited on p. 23, 114, 115

[227] Rado, R., "A sequence of polyhedra having intersections of specified dimensions", *Journal of the London Mathematical Society* 22:4 (October 1947), pp. 287–289. Cited on p. 25

[228] Roberts, F.S., "On the boxicity and cubicity of a graph" in Tutte, W.T. (ed.) *Recent Progress in Combinatorics: Proceedings of the Third Waterloo Conference on Combinatorics*, Academic Press, New York, 1969, pp. 301–310. Cited on p. 20

[229] Robertson, N., Seymour, P.D., "Graph Minors. I. Excluding a Forest", *Journal of Combinatorial Theory, Series B* 35:1 (August 1983), pp. 39–61. Cited on p. 50

[230] Robertson, N., Seymour, P.D., "Graph minors. X. Obstructions to tree-decomposition", *Journal of Combinatorial Theory, Series B* 52:2 (July 1991), pp. 153–190. Cited on p. 50

[231] Sachs, H., "Coin graphs, polyhedra, and conformal mapping", *Discrete Mathematics* 134:1–3 (November 1994), pp. 133–138. Cited on p. 23, 26

[232] Santi, P., "Topology control in wireless ad hoc and sensor networks", *ACM Computing Surveys* 37:2 (June 2005), pp. 164–194. Cited on p. 3

[233] Schaefer, M., Sedgwick, E., Štefankovič, D., "Recognizing string graphs in NP", *Journal of Computer and System Sciences* 67:2 (September 2003), pp. 365–380. Cited on p. 20

[234] Scheinerman, E.R., *Intersection classes and multiple intersection parameters of graphs*, PhD thesis, Princeton University, 1984. Cited on p. 27

[235] Scheinerman, E.R., West, D.B., "The interval number of a planar graph: Three intervals suffice", *Journal of Combinatorial Theory, Series B* 35:3 (December 1983), pp. 224–239. Cited on p. 26

[236] Schmid, S., Wattenhofer, R., "Algorithmic Models for Sensor Networks" in *20th International Parallel and Distributed Processing Symposium (IPDPS 2006), Proceedings, 25-29 April 2006, Rhodes Island, Greece*, IEEE Computer Society, 2006. Cited on p. 24

[237] Schramm, O., *Packing two-dimensional bodies with prescribed combinatorics and applications to the construction of conformal and quasiconformal mappings*, PhD thesis, Princeton, 1990. Cited on p. 17, 28

[238] Schramm, O., "Conformal Uniformization and Packings", *Israel Journal of Mathematics* 93:1 (December 1996), pp. 399–428. Cited on p. 28

[239] Seese, D., "Tree-partite graphs and the complexity of algorithms" in Budach, L. (ed.) *Fundamentals of Computation Theory, FCT '85, Cottbus, GDR, September 9-13, 1985*, Lecture Notes in Computer Science 199, Springer-Verlag, Berlin, 1985, pp. 412–421. Cited on p. 52

[240] Simion, R., "Noncrossing partitions", *Discrete Mathematics* 217:1–3 (April 2000), pp. 367–409. Cited on p. 63

[241] Sinden, F.W., "Topology of Thin Film RC-Circuits", *The Bell System Technical Journal* 45 (November 1966), pp. 1639–1662. Cited on p. 20, 25

[242] Sinden, F.W., "Topology of Thin Film RC-Circuits" in Rosenstiehl, P. (ed.) *Theory of Graphs, International Symposium, International Computation Centre, Rome, July 1966*, Gordon and Breach, New York, 1967, pp. 389–393. Cited on p. 20, 25

[243] Smith, W.D., *Accurate Circle Configurations and Numerical Conformal Mapping in Polynomial Time*, unpublished, 1991. Cited on p. 27

[244] Smith, W.D., Wormald, N.C., "Geometric Separator Theorem & Applications" in *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, IEEE Computer Society, 1998, pp. 232–243. Cited on p. 70

[245] Smorodinsky, S., "On the Chromatic Number of Geometric Hypergraphs", *SIAM Journal on Discrete Mathematics* 21:3 (2007), pp. 676–687. Cited on p. 48

[246] Spinrad, J.P., *Efficient Graph Representations*, Field Institute Monographs 19, American Mathematical Society, Providence, Rhode Island, 2003. Cited on p. 18

[247] Stein, S.K., "Convex Maps", *Proceedings of the American Mathematical Society* 2:3 (June 1951), pp. 464–466. Cited on p. 27

[248] Szpilrajn-Marczewski, E., "Sur deux propriétés des classes d'ensembles", *Fundamenta Informaticae* 33 (1945), pp. 303–307. Cited on p. 18

[249] Telle, J.A., Proskurowski, A., "Practical Algorithms on Partial k-Trees with an Application to Domination-like Problems" in Dehne, F.K.H.A., Sack, J.R., Santoro, N., Whitesides, S. (eds.) *Algorithms and Data Structures, Third Workshop, WADS '93, Montréal, Canada, August 11-13, 1993, Proceedings*, Springer-Verlag, Berlin 709, Springer-Verlag, Berlin, 1993, pp. 610–621. Cited on p. 50, 55

[250] Thomassen, C., "Planarity and duality of finite and infinite graphs", *Journal of Combinatorial Theory, Series B* 29:2 (October 1980), pp. 244–271. Cited on p. 28

[251] Thomassen, C., "Interval Representations of Planar Graphs", *Journal of Combinatorial Theory, Series B* 40:1 (February 1986), pp. 9–20. Cited on p. 17, 28

[252] Toh, C.K., *Ad Hoc Mobile Wireless Networks: Protocols and Systems*, Prentice-Hall, 2002. Cited on p. 3

[253] Tucker, A., "Structure theorems for some circular-arc graphs", *Discrete Mathematics* 7:1–2 (1974), pp. 167–195. Cited on p. 19

[254] Tutte, W.T., "Convex Representations of Graphs", *Journal of the London Mathematical Society* 10:1 (1960), pp. 304–320. Cited on p. 27

[255] Valiant, L.G., "Universality Considerations in VLSI Circuits", *IEEE Transactions on Computers* 30:2 (1981), pp. 135–140. Cited on p. 182, 201

[256] van Leeuwen, E.J., *A Proof of the Clarkson-Varadarajan Result*, unpublished. Cited on p. 115

[257] van Leeuwen, E.J., *Approximation Algorithms for Unit Disk Graphs*, Technical Report UU-CS-2004-066, Institute of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands, 2004. Cited on p. 22

[258] van Leeuwen, E.J., *Optimization Problems on Mobile Ad Hoc Networks – Algorithms for Disk Graphs*, Master's thesis INF/SCR-04-32, Institute of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands, 2004. Cited on p. 53, 68, 70, 92

[259] van Leeuwen, E.J., "Approximation Algorithms for Unit Disk Graphs" in Kratsch, D. (ed.) *Graph-Theoretic Concepts in Computer Science, 31st International Workshop, WG 2005, Metz, France, June 23-25, 2005, Revised Selected Papers*, Lecture Notes in Computer Science 3787, Springer-Verlag, Berlin, 2005, pp. 351–361. Cited on p. 6

[260] van Leeuwen, E.J., “Better Approximation Schemes for Disk Graphs” in Arge, L., Freivalds, R. (eds.) *Algorithm Theory - SWAT 2006, 10th Scandinavian Workshop on Algorithm Theory, Riga, Latvia, July 6-8, 2006, Proceedings*, Lecture Notes in Computer Science 4059, Springer-Verlag, Berlin, 2006, pp. 316–327. Cited on p. 6

[261] van Leeuwen, E.J., van Leeuwen, J., *Convex Polygon Intersection Graphs*, submitted. Cited on p. 21, 30, 36

[262] van Leeuwen, E.J., van Leeuwen, J., *On the Representation of Disk Graphs*, Technical Report UU-CS-2006-037, Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands, 2006. Cited on p. 5, 38

[263] van Leeuwen, E.J., van Leeuwen, J., *Structure of Polynomial-Time Approximation*, manuscript, 2008. Cited on p. 5

[264] Vohra, R., Hall, N.G., “A Probabilistic Analysis of the Maximal Covering Location Problem”, *Discrete Applied Mathematics* 43:2 (May 1993), pp. 175–183. Cited on p. 162

[265] Walter, J.R., “Representations of chordal graphs as subtrees of a tree”, *Journal of Graph Theory* 2:3 (1978), pp. 265–267. Cited on p. 18, 26

[266] Wan, P.J., Alzoubi, K.M., Frieder, O., “Distributed Construction of Connected Dominating Set in Wireless Ad Hoc Networks”, *Mobile Networks and Applications* 9:2 (April 2004), pp. 141–149. Cited on p. 47

[267] Wang, D.W., Kuo, Y.S., “A Study on Two Geometric Location Problems”, *Information Processing Letters* 28:6 (August 1988), pp. 281–286. Cited on p. 2, 44, 68

[268] Wegner, G., *Eigenschaften der Nerven homologisch-einfacher Familien im* $\mathbb{R}^n$, PhD thesis, Georg-August-Universität, Göttingen, 1967. Cited on p. 17, 19, 25, 27

[269] West, D.B., Shmoys, D.B., “Recognizing graphs with fixed interval number is NP-complete”, *Discrete Applied Mathematics* 8:3 (July 1984), pp. 295–305. Cited on p. 19

[270] Wiese, A., Kranakis, E., “Local PTAS for Independent Set and Vertex Cover in Location Aware Unit Disk Graphs” in Nikoletseas, S.E., Chlebus, B.S., Johnson, D.B., Krishnamachari, B. (eds.) *Distributed Computing in Sensor Systems, 4th IEEE International Conference, DCOSS 2008, Santorini Island, Greece, June 11-14, 2008, Proceedings*, Lecture Notes in Computer Science 5067, Springer-Verlag, Berlin, 2008, pp. 415–431. Cited on p. 47, 75

[271] Wilf, H.S., *generatingfunctionology*, Second Edition, Academic Press, New York, 1994. Cited on p. 62

[272] Xu, J., Berger, B., "Fast and Accurate Algorithms for Protein Side-Chain Packing", *Journal of the ACM* 53:4 (July 2006), pp. 533–557. Cited on p. 4, 45

[273] Yannakakis, M., "The Complexity of the Partial Order Dimension Problem", *SIAM Journal on Algebraic and Discrete Methods* 3:3 (September 1982), pp. 351–358. Cited on p. 20

[274] Yick, J., Mukherjee, B., Ghosal, D., "Wireless sensor network survey", *Computer Networks* 52:12 (August 2008), pp. 2292–2330. Cited on p. 3

[275] Young, N.E., "Sequential and Parallel Algorithms for Mixed Packing and Covering" in *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, IEEE Computer Society, 2001, pp. 538–546. Cited on p. 118

[276] Yu, G., Goldschmidt, O., "On Locally Optimal Independent Sets and Vertex Covers", *Naval Research Logistics* 43:5 (August 1996), pp. 737–748. Cited on p. 44

[277] Yu, G., Kouvelis, P., Luo, S., "Weighted Vertex Packing Problem for Specially Structured Geometric Graphs", *Naval Research Logistics* 42:1 (February 1995), pp. 81–102. Cited on p. 44

[278] Zhang, Z., Gao, X., Wu, W., Du, D.Z., "PTAS for Minimum Connected Dominating Set in Unit Ball Graph" in Li, Y., Huynh, D.T., Das, S.K., Du, D.Z. (eds.) *Wireless Algorithms, Systems, and Applications, Third International Conference, WASA 2008, Dallas, TX, USA, October 26-28, 2008, Proceedings*, Lecture Notes in Computer Science 5258, Springer-Verlag, Berlin, 2008, pp. 154–161. Cited on p. 47, 82

[279] Zou, F., Li, X., Kim, D., Wu, W., "Two Constant Approximation Algorithms for Node-Weighted Steiner Tree in Unit Disk Graphs" in Yang, B., Du, D.Z., Wang, C.A. (eds.) *Combinatorial Optimization and Applications, Second International Conference, COCOA 2008, St. John's, NL, Canada, August 21-24, 2008, Proceedings*, Lecture Notes in Computer Science 5165, Springer-Verlag, Berlin, 2008, pp. 278–285. Cited on p. 2, 47

Author Index

For each author that (co)-wrote a paper referenced in the bibliography, this index lists the pages on which a paper of this author was cited as well as (between []) a list of the papers this author was involved in.

Index

Samenvatting

Praktische problemen in bijvoorbeeld draadloze netwerken, computationele biologie of cartografie kunnen vaak gemodelleerd worden door een optimaliseringsprobleem op een systeem van geometrische objecten te definiëren. Veel optimaliseringsproblemen zijn NP-moeilijk om exact op te lossen en soms ook bewijsbaar moeilijk te benaderen. Echter, als de onderliggende structuur van het optimaliseringsprobleem een systeem van geometrische objecten is, dan blijkt het probleem meestal makkelijk exact op te lossen of goed te benaderen. Dit proefschrift onderzoekt de benaderbaarheid van moeilijke optimaliseringsproblemen op systemen van geometrische objecten. In het bijzonder bekijkt het wat de invloed van de vorm van de objecten op de benaderheid is.

De belangrijkste structuur over systemen van geometrische objecten die we tijdens dit onderzoek beschouwen is de intersectiegraaf van de gegeven objecten, een geometrische intersectiegraaf genoemd. Iedere knoop in deze graaf correspondeert met een object en er is een kant tussen twee knopen dan en slechts dan als de corresponderende objecten een niet-lege doorsnede hebben. Het bekendste voorbeeld hiervan zijn schijfintersectiegrafen (disk graphs), die vaak als model voor draadloze netwerken worden gebruikt. Verscheidene klassieke optimaliseringsproblemen, zoals Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set, zijn relevant in deze context. Wat betreft de benaderbaarheid van deze problemen op geometrische intersectiegrafen concluderen we het volgende.

Als een verzameling van n schijven van gelijke grootte (eenheidsschijven) gegeven is, waarvan de dichtheid (density) d is, dan kan voor ieder van de bestudeerde problemen in $d^{O(1/\epsilon)}n^{O(1)}$ tijd een $(1+\epsilon)$-benadering van het optimum gevonden worden. Dit resulteert in een eptas als $d = d(n) = n^{o(1)}$ en een ptas in het algemeen. Voor Minimum Vertex Cover kunnen we het algoritme zelfs versterken tot een eptas in het algemene geval. Deze schema's zijn uitbreidbaar tot intersectiegrafen van gelijke grote dikke objecten in constante dimensie. Voorts kan worden aangetoond dat er, op constantes na, geen sneller algoritme is om een $(1+\epsilon)$-benadering te vinden, tenzij de exponential time hypothesis onwaar is. Behalve voor Minimum Vertex Cover is er ook geen eptas voor de problemen als $d = d(n) = n^{\alpha}$ voor een willekeurige constante $\alpha > 0$, tenzij FPT = W[1]. Dit is een sterke indicatie dat de ontworpen schema's optimaal zijn.

De benaderingsschema's zijn voor Maximum Independent Set and Minimum Vertex Cover uit te breiden tot intersectiegrafen van willekeurig grote schijven. We krijgen een eptas als de level density $d = d(n) = n^{o(1)}$ is, een ptas in het algemeen en voor Minimum Vertex Cover zelfs een eptas in het al-

gemene geval. De schema's werken ook op intersectiegrafen van dikke objecten in constante dimensie. We kunnen net als voor eenheidsschijfintersectiegrafen aantonen dat deze schema's optimaal zijn op constantes na.

Bekijken we echter Minimum (Connected) Dominating Set, dan blijkt dit probleem een stuk lastiger te worden op intersectiegrafen van objecten van willekeurige grootte. Zijn de objecten willekeurig geschaalde en getransleerde kopieën van een convexe veelhoek, dan is er een constante-factor-benaderingsalgoritme, het eerste benaderingsalgoritme voor dit probleem dat beter is dan de $\ln n$-benadering van het greedy algoritme. Uitbreiding naar schijfintersectiegrafen is niet direct mogelijk, omdat de constante in de benaderingsfactor van het algoritme afhangt van de complexiteit van de veelhoek. Wordt ieder punt van het vlak echter maar door een begrensd aantal schijven overlapt, dan kan in polynomiale tijd een $(3 + \epsilon)$-benadering gevonden worden (voor vaste $\epsilon > 0$). Dit werkt ook voor intersectiegrafen van dikke objecten in constante dimensie. Is het aantal objecten dat een bepaald punt overlapt niet begrensd en wijkt de vorm van de objecten maar iets af van de vorm van een schijf (maar is nog steeds dik), dan kan aangetoond worden dat het probleem niet beter dan een $\ln n$-benaderingsalgoritme heeft, tenzij NP $\subset$ DTIME$(n^{O(\log \log n)})$.

Een ander belangrijk probleem om te beschouwen is de geometrische versie van het bekende minimum set cover probleem en enkele van z'n varianten. Zij gegeven een verzameling geometrische objecten en een verzameling punten in het vlak, dan dient bij dit probleem een zo klein mogelijke deelverzameling van de objecten gevonden te worden die alle punten overdekken. Zijn de objecten vierkanten van gelijke grootte, dan geven wij een ptas voor dit probleem, een van de eerste ptas-en voor dit probleem in twee dimensies en de eerste die ook in het gewogen geval werkt.

Als we een deelverzameling van de objecten willen waarbij zoveel mogelijk gegeven punten door precies één object overdekt wordt, dan hebben we een algoritme dat een 1/18-benadering geeft als de objecten eenheidsschijven zijn en een 1/2-benadering als de objecten eenheidsvierkanten zijn. Als de objecten verschillende grootte hebben en dik zijn, dan hebben is er een eptas als ieder punt in het vlak maar door een begrensd aantal objecten overlapt wordt. Is dit aantal niet begrensd, dan is het wederom zo dat het probleem niet beter dan een $\ln n$-benaderingsalgoritme heeft.

Alle drie de bovenstaande algoritmen zijn uitbreidbaar is tot het geval waarbij ieder object een prijs heeft, ieder punt een opbrengst en we de totale opbrengst van (enkelvoudig) overdekte punten willen maximaliseren zodanig dat de totale kosten binnen een bepaald budget vallen.

De voornaamste conclusie van het onderzoek is dat bekende optimaliseringsproblemen op systemen van geometrische objecten vaak beter benaderbaar zijn dan op algemene systemen. We hebben echter laten zien dat het type van de objecten van grote invloed is op de benaderbaarheid. In het bijzonder lijken problemen als Minimum Dominating Set, Minimum Set Cover en enkele van z'n varianten lastiger op een systeem van schijven dan op een systeem van (zeg) vierkanten. Verder onderzoek naar dit fenomeen vormt een nieuwe uitdaging.

Summary

Practical problems in for example wireless networks, computational biology, or cartography can often be modeled by defining an optimization problem on a system of geometric objects. Many optimization problems are NP-hard to solve and sometimes even provably hard to approximate. However, if the underlying structure of the optimization problem is a system of geometric objects, then the problem is usually easy to solve or can be approximated well. This PhD thesis investigates the approximability of hard optimization problems on systems of geometric objects. In particular it considers how the object type influences the approximability of these problems.

The most important structure on systems of geometric objects that we consider during this investigation is the intersection graph of the given objects, called a geometric intersection graph. Every vertex of this graph corresponds to an object and there is an edge between two vertices if and only if the corresponding objects intersect. A famous example of this is the disk graph, which is frequently used to model wireless networks. Several classic optimization problems, such as Maximum Independent Set, Minimum Vertex Cover, and Minimum (Connected) Dominating Set, are relevant in this context. With respect to the approximability of these problems, we conclude the following.

If a collection of n disks of equal size (unit disks) is given of which the density is d, then we can give a $(1+\epsilon)$-approximation for each of the studied problems in $d^{O(1/\epsilon)}n^{O(1)}$ time. This yields an eptas if $d = d(n) = n^{o(1)}$ and a ptas in general. For Minimum Vertex Cover we can even strengthen the algorithm to an eptas in the general case. These schemes can be extended to intersection graphs of unit fat objects in constant dimension. Moreover we can show that, up to constants, there is no faster algorithm to find a $(1+\epsilon)$-approximation for these problems, unless the exponential time hypothesis fails. Except for Minimum Vertex Cover, we also prove that there is no eptas for the studied problems if $d = d(n) = n^{\alpha}$ for any constant $\alpha > 0$, unless FPT = W[1]. This gives a strong indication that the given schemes are optimal.

The approximation schemes for Maximum Independent Set and Minimum Vertex Cover can be extended to intersection graphs of arbitrarily sized disks. We obtain an eptas if the level density is $d = d(n) = n^{o(1)}$, a ptas in general, and for Minimum Vertex Cover we even obtain an eptas in the general case. The schemes also apply to intersection graphs of fat objects in constant dimension. As in the case of unit disk graphs, we can show that these schemes are optimal, up to constants.

If however we consider Minimum (Connected) Dominating Set, then it turns out that this problem becomes a lot harder on intersection graphs of ob-

jects of arbitrary size. If the objects are arbitrarily scaled and translated copies of a convex polygon, then there is a constant-factor approximation algorithm. This is the first approximation algorithm for this problem that beats the $\ln n$-approximation given by the greedy algorithm. Extending to disk graphs is not directly possible, because the constant in the approximation factor depends on the complexity of the polygon. If however each point of the plane is overlapped by a bounded number of disks, then a $(3 + \epsilon)$-approximation can be found in polynomial time (for fixed $\epsilon > 0$). This also applies to intersection graphs of fat objects in constant dimension. If the number of objects overlapping some point is not bounded and the shape of the objects is only slightly different from the shape of a disk (but still is fat), then we can show that the $\ln n$-approximation algorithm gives the best possible approximation factor, unless $\mathrm{NP} \subset \mathrm{DTIME}(n^{\mathrm{O}(\log \log n)})$.

Another important problem to study is the geometric version of the well-known minimum set cover problem and several of its variants. Given a set of geometric objects and a set of points in the plane, one needs to find a smallest subset of the objects that cover all points. If the objects are squares of equal size, then we give a ptas for this problem, one of the first approximation schemes for this problem in two dimensions and the first that also applies to the weighted case.

If we want to find a subset of the objects such that a maximum number of the given points is overlapped by precisely one object, then we give a 1/18-approximation algorithm if the objects are unit disks and a 1/2-approximation algorithm if the objects are unit squares. If the objects have arbitrary size and are fat, then we give an eptas if each point in the plane is overlapped by a bounded number of objects. If this number is not bounded however, then we can again prove that the $\ln n$-approximation algorithm gives the best possible approximation factor.

The above three algorithms can be extended to the case where each object has a cost, each point has a profit, and we aim to maximize the total profit of (uniquely) covered points such that the total cost is within a given budget.

The foremost conclusion of this thesis is that well-known optimization problems are often better approximable on systems of geometric objects than on general systems. We have shown however that the object type has strong influence on the approximability of these problems. In particular, problems such as Minimum Dominating Set, Minimum Set Cover, and several of its variants seem more difficult on a system of disks than on a system of (say) squares. Further investigation of this phenomenon poses a new challenge.

Acknowledgments

A PhD thesis is not just a collection of definitions, lemmas, and theorems. It marks the end of a four-year research period and the start of something new. It is a waypoint on one's path through science. My journey along this path has so far been a fun and challenging experience, but one that I could not have made without the support of many people, whom I would like to thank below.

First and foremost, I would like to thank Lex Schrijver, who has been a great supervisor. Lex gave me a lot of freedom to pursue my own research, even when this led to a topic further away from his current line of research. I am very grateful for his support in this. Besides our mathematical discussions, I greatly enjoyed our talks about science in general.

A crucial ingredient in this thesis is the research I did in collaboration with Thomas Erlebach. This has been a great experience for me and I want to thank him for hosting me in Leicester several times. Thanks also to Mark de Berg for giving me the opportunity to assist in teaching algorithms in Eindhoven.

The research in this thesis has been carried out in the PNA1 group at the Centrum Wiskunde & Informatica in Amsterdam. This has been a wonderful and fruitful environment that allowed me to learn about a lot of topics not necessarily always directly related to the subject of my thesis, for instance in the SDP seminar and the Graph Theory reading group. A lot of people contributed to this environment, among them Benny Mounits, Bert Gerards, Gabor Maroti, Hartwig Bosse, Jarek Byrka, Karen Aardal, Leen Stougie, Monique Laurant, Nebojša Gvozdenović, Nicole Immorlica, Peter Korteweg, Steven Kelk, and Vangelis Markakis.

Time at CWI would certainly have gone a lot slower had it not been for my roommates over the years: Willem-Jan van Hoeve, Dion Gijswijt, Jules van Kempen, and during the last three years, Frank Vallentin and Fernando de Oliveira Filho. I want to thank them for their support, feedback, and friendship. In particular, I would like to thank Fernando, especially for the talks we had about movies and for the trip to Milan.

Navigating CWI would have been a lot less easy without the support staff. I want to thank Susanne van Dam and Nada Mitrovic in particular for helping me find my way. The library staff was also more than helpful and always managed to find that elusive book or paper.

My favorite hypothesis is that the less you think about a problem, the more ideas you have on solving it. A good way of achieving this is to do sports, in my case volleyball. My thanks for this useful distraction go to the people I met on Thursday's at training and on Tuesday's at Van Slag, especially to the members of the two volleyball teams I played in: Los! and Doorgewinterd.

Last year, some Van Slaggers banded together to form a team to play in the regular volleyball competition. Playing with Tovo Heren 4 has been a blast and becoming champions in our league was a wonderful experience.

At volleyball I met many people and made a lot of friends. Thank you all for making volleyball my favorite pastime! I want to thank Erik Brave, Gijs Scheenstra, Ivo Becker, Walter Bruins, and in particular Mark Bouts for many memorable evenings, talking about the meaningful (and sometimes the less meaningful) things in life.

Last, but certainly not least, I want to thank my parents. It is a wonderful feeling to have such unwavering support to help me achieve my dreams.