

EPTAS and Subexponential Algorithm for Maximum Clique on Disk and Unit Ball Graphs

MARTHE BONAMY, CNRS, LaBRI, Université de Bordeaux, France

ÉDOUARD BONNET, Université de Lyon (COMUE), CNRS, ENS de Lyon, Université Claude-Bernard Lyon 1, LIP, France

NICOLAS BOUSQUET, CNRS, G-SCOP laboratory, France

PIERRE CHARBIT, Université Paris Diderot - IRIF, France and Université de Lyon (COMUE), CNRS, ENS de Lyon, Université Claude-Bernard Lyon 1, LIP

PANOS GIANNOPOULOS, giCentre, Department of Computer Science, City University of London, United Kingdom

EUN JUNG KIM, Université Paris-Dauphine, PSL University, CNRS UMR, LAMSADE, France

PAWEŁ RZAŻEWSKI, Faculty of Mathematics and Information Science, Warsaw University of Technology, Poland and Institute of Informatics, University of Warsaw, Poland

FLORIAN SIKORA, Université Paris-Dauphine, PSL University, CNRS UMR, LAMSADE, France

STÉPHAN THOMASSÉ, Université de Lyon (COMUE), CNRS, ENS de Lyon, Université Claude-Bernard Lyon 1, LIP, France and Institut Universitaire de France

A (unit) disk graph is the intersection graph of closed (unit) disks in the plane. Almost three decades ago, an elegant polynomial-time algorithm was found for MAXIMUM CLIQUE on unit disk graphs [Clark, Colbourn, Johnson; *Discrete Mathematics* '90]. Since then, it has been an intriguing open question whether or not tractability can be extended to general disk graphs. We show that the disjoint union of two odd cycles

Research partially supported by EPSRC grant FptGeom (EP/N029143/1), ANR grant ESIGMA (ANR-17-CE23-0010), the LABEX MILYON (ANR-10- LABX-0070) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR), the ANR Project DISTANCIA (ANR-17-CE40-0015) operated by the French National Research Agency (ANR), and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme Grant Agreement 714704.

Extended abstracts of this article appeared in References [17] and [15].

Authors’ addresses: M. Bonamy, CNRS, LaBRI, Université de Bordeaux, Bordeaux, France; email: marthe.bonamy@u-bordeaux.fr; É. Bonnet, Université de Lyon (COMUE), CNRS, ENS de Lyon, Université Claude-Bernard Lyon 1, LIP, Lyon, France; email: edouard.bonnet@dauphine.fr; N. Bousquet, CNRS, G-SCOP laboratory, Grenoble-INP, France; email: nicolas.bousquet@grenoble-inp.fr; P. Charbit, Université Paris Diderot - IRIF, Paris, France, Université de Lyon (COMUE), CNRS, ENS de Lyon, Université Claude-Bernard Lyon 1, LIP; email: charbit@irif.fr; P. Giannopoulos, giCentre, Department of Computer Science, City University of London, London, United Kingdom; email: panos.giannopoulos@city.ac.uk; E. J. Kim, Université Paris-Dauphine, PSL University, CNRS UMR, LAMSADE, Paris, 75016, France; email: eun-jung.kim@dauphine.fr; P. Rzażewski, Faculty of Mathematics and Information Science, Warsaw University of Technology, Warsaw, Poland, Institute of Informatics, University of Warsaw, Warsaw, Poland; email: p.rzazewski@mini.pw.edu.pl; F. Sikora, Université Paris-Dauphine, PSL University, CNRS UMR, LAMSADE, Paris, 75016, France; email: florian.sikora@dauphine.fr; S. Thomassé, Université de Lyon (COMUE), CNRS, ENS de Lyon, Université Claude-Bernard Lyon 1, LIP, Lyon, France, Institut Universitaire de France; email: stephan.thomasse@ens-lyon.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

0004-5411/2020/01-ART9 \$15.00

<https://doi.org/10.1145/3433160>

is never the complement of a disk graph nor of a unit (3-dimensional) ball graph. From that fact and existing results, we derive a simple QPTAS and a subexponential algorithm running in time $2^{\tilde{O}(n^{2/3})}$ for MAXIMUM CLIQUE on disk and unit ball graphs. We then obtain a randomized EPTAS for computing the independence number on graphs having no disjoint union of two odd cycles as an induced subgraph, bounded VC-dimension, and linear independence number. This, in combination with our structural results, yields a randomized EPTAS for MAX CLIQUE on disk and unit ball graphs. MAX CLIQUE on unit ball graphs is equivalent to finding, given a collection of points in \mathbb{R}^3 , a maximum subset of points with diameter at most some fixed value.

In stark contrast, MAXIMUM CLIQUE on ball graphs and unit 4-dimensional ball graphs, as well as intersection graphs of filled ellipses (even close to unit disks) or filled triangles is unlikely to have such algorithms. Indeed, we show that, for all those problems, there is a constant ratio of approximation that cannot be attained even in time $2^{n^{1-\epsilon}}$, unless the Exponential Time Hypothesis fails.

CCS Concepts: • **Mathematics of computing** → **Graph theory; Approximation algorithms**; • **Theory of computation** → **Computational geometry**;

Additional Key Words and Phrases: Disk graph, maximum clique, computational complexity, approximation algorithms, subexponential algorithms

ACM Reference format:

Marthe Bonamy, Édouard Bonnet, Nicolas Bousquet, Pierre Charbit, Panos Giannopoulos, Eun Jung Kim, Paweł Rzażewski, Florian Sikora, and Stéphan Thomassé. 2020. EPTAS and Subexponential Algorithm for Maximum Clique on Disk and Unit Ball Graphs. *J. ACM* 68, 2, Article 9 (January 2021), 38 pages.

<https://doi.org/10.1145/3433160>

1 INTRODUCTION

In an *intersection graph*, the vertices are geometric objects with an edge between any pair of intersecting objects. Intersection graphs have been studied for many different families of objects due to their practical applications and their rich structural properties [18, 50]. Among the most studied ones are *disk graphs*, which are intersection graphs of closed disks in the plane, and their special case, *unit disk graphs*, where all the radii are equal. Their applications range from sensor networks to map labeling [34], and many standard optimization problems have been studied on disk graphs; see for example Reference [60] and references therein. Most of the hard optimization and decision problems remain NP-hard on disk graphs and even unit disk graphs. For instance, disk graphs contain planar graphs [44] on which several of those problems are intractable. However, shifting techniques and separator theorems may often lead to subexponential classical or parameterized algorithms [5, 12, 49, 56]. Many approximation algorithms have been designed specifically on (unit) disk graphs, or more generally on geometric intersection graphs; see for instance References [3, 26, 27, 33, 36, 53, 54, 59] to cite only a few. Besides ad hoc techniques, local search and VC-dimension play an important role in the approximability of problems on (unit) disk graphs. For the main packing and covering problems (MAXIMUM INDEPENDENT SET, MIN VERTEX COVER, MINIMUM DOMINATING SET, MINIMUM HITTING SET, and their weighted variants) at least a PTAS is known.

However, all the mentioned techniques are only amenable to packing and covering problems. The MAXIMUM CLIQUE problem is arguably the most prominent problem that does not fall into those categories. For example, anything along the lines of exploiting a small separator cannot work for MAXIMUM CLIQUE, where the densest instances are the hardest. Therefore, it seems that new ideas are necessary to get improved approximate or exact algorithms for this problem. In this article, we focus on solving MAXIMUM CLIQUE on (unit) disk graphs in dimension 2 or higher.

Previous results. In 1990, Clark et al. [29] gave an elegant polynomial-time algorithm for MAXIMUM CLIQUE on unit disk graphs when the input is a geometric representation of the graph. It goes

as follows: Guess in quadratic time the two more distant centers of disks in a maximum clique (at distance at most 2), remove all the centers that would contradict this maximality, observe that the resulting graph is co-bipartite. Hence, one can find an optimum solution in polynomial time by looking for a maximum independent set in the complement graph, which is bipartite. However, recognizing unit disk graphs is NP-hard [19], and even $\exists\mathbb{R}$ -complete [41]. In particular, if the input is the mere unit disk graph, one cannot expect to efficiently compute a geometric representation to run the previous algorithm. Raghavan and Spinrad showed how to overcome this issue and suggested a polynomial-time algorithm that does not require the geometric representation [55]. Their algorithm is a subtle *blind* reinterpretation of the algorithm by Clark et al. It solves MAXIMUM CLIQUE on a superclass of the unit disk graphs or correctly claims that the input is not a unit disk graph. Hence, it cannot be used to efficiently recognize unit disk graphs.

The complexity of MAXIMUM CLIQUE on general disk graphs is a notorious open question in computational geometry. On the one hand, no polynomial-time algorithm is known, even when the geometric representation is given. On the other hand, the NP-hardness of the problem has not been established, even when only the graph is given as input.

The piercing number of a collection of geometric objects is the minimum number of points that hit all the objects. It is known since the '50s (although the first published records of that result came later in the '80s) that the piercing number of pairwise intersecting disks is 4 [30, 57], meaning that 4 points are always sufficient and sometimes necessary. An account of this story can be found in a recent paper by Har-Peled et al. [38]. The same paper gives a linear-time algorithm to find 5 points piercing any collection of pairwise intersecting disks, as well as a configuration with 13 pairwise intersecting disks for which checking that 4 points are necessary to hit all the disks is quite simple. Carmi et al. [24] improved the former result by presenting a linear-time algorithm that finds a hitting set of at most 4 points.

Ambühl and Wagner observed that this yields a 2-approximation for MAXIMUM CLIQUE [7]. Indeed, after guessing in polynomial time four points hitting a maximum clique and removing every disk not hit by those points, the instance is partitioned into four cliques; or equivalently, two co-bipartite graphs. One can then solve optimally each instance formed by one co-bipartite graph and return the larger solution of the two. This cannot give a solution more than twice smaller than the optimum. Since then, the problem has proved to be elusive with no new positive or negative results. The question on the complexity and further approximability of MAXIMUM CLIQUE on general disk graphs is considered as folklore [11], but was also explicitly mentioned as an open problem by Fishkin [34] and Ambühl and Wagner [7]. Cabello even asked if there is a 1.99-approximation for disk graphs with two sizes of radii [21, 22].

Our results. We fully characterize the disk graphs whose complements have maximum degree 2, i.e., are unions of cycles and paths: A complement of a disjoint union of paths and cycles is a disk graph if and only if the number of odd cycles in this union is at most one. For algorithmic purposes, the interesting part of that statement is:

THEOREM 1.1. *A complement of a disk graph cannot have a disjoint union of two odd cycles as an induced subgraph. In other words, if G is a disk graph, then $\text{iocp}(\overline{G}) \leq 1$.*

In the previous statement iocp denotes the *induced odd cycle packing number* of a graph, i.e., the maximum number of odd cycles as a disjoint union in an induced subgraph.

We show the same forbidden induced subgraphs for unit ball graphs as for disk graphs. The proofs for disk graphs and unit ball graphs are quite different and the classes are incomparable.

THEOREM 1.2. *A complement of a unit ball graph cannot have a disjoint union of two odd cycles as an induced subgraph. In other words, if G is a unit ball graph, then $\text{iocp}(\overline{G}) \leq 1$.*

We then present a randomized EPTAS (Efficient Polynomial-Time Approximation Scheme, that is, a PTAS in time $f(\varepsilon)n^{O(1)}$) for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, \beta, 1)$. The class $\mathcal{X}(d, \beta, 1)$ denotes the class of graphs whose neighborhood hypergraph has VC-dimension at most d , independence number at least βn , and no disjoint union of two odd cycles as an induced subgraph (for formal definitions see Section 2).

THEOREM 1.3. *For any constants $d \in \mathbb{N}$, $0 < \beta \leq 1$, for every $0 < \varepsilon < 1$, there is a randomized $(1 - \varepsilon)$ -approximation algorithm running in time $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$, and a deterministic PTAS running in time $n^{\tilde{O}(1/\varepsilon^3)}$ for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, \beta, 1)$ with n vertices.*

It is then easy to reduce MAXIMUM CLIQUE on disk graphs to MAXIMUM INDEPENDENT SET on $\mathcal{X}(4, \beta, 1)$ for some constant β . We therefore obtain, due to Theorem 1.3 and Theorem 1.1, a randomized EPTAS (and a deterministic PTAS) for MAXIMUM CLIQUE on disk graphs, settling almost¹ completely the approximability of this problem.

THEOREM 1.4. *There is a randomized EPTAS for MAXIMUM CLIQUE on disk graphs, even if a geometric representation is not given. Its running time is $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$ for a $(1 - \varepsilon)$ -approximation on a graph with n vertices.*

Similarly, Theorem 1.3 and Theorem 1.2 naturally lead to:

THEOREM 1.5. *There is a randomized EPTAS in time $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$ for MAXIMUM CLIQUE on unit ball graphs, even if a geometric representation is not given.*

Before that result, the best approximation factor was 2.553, due to Afshani and Chan [1]. In particular, even getting a 2-approximation algorithm (as for disk graphs) was open.

Finally, we show that such an approximation scheme, even in subexponential time, is unlikely for ball graphs (i.e., intersection graphs of 3-dimensional balls with arbitrary radii), unit 4-dimensional ball graphs, as well as intersection graphs of (filled) triangles and (filled) ellipses. Our lower bounds also imply NP-hardness. To the best of our knowledge, the NP-hardness of MAXIMUM CLIQUE on unit d -dimensional ball graphs was only known when d is superconstant ($d = \Omega(\log n)$) [2].

In the following paragraphs, we sketch the principal lines of our three main contributions.

The complement of the union of two odd cycles is not a disk graph. Let H be the complement of two cycles $C_1 \uplus C_2$. We focus on the possible configurations for the four centers associated to two non-edges of H : one edge of C_1 and one edge of C_2 . It can be shown that, if those four centers are in convex position, then the non-edges of H should *cross*. In other words, there is a forbidden configuration: four centers in convex position with the two non-edges being sides of the quadrangle. For each edge e of C_1 (respectively, C_2), we define three quantities counting how many edges of C_2 (respectively, C_1) are *crossed* by e (for three different meanings of *crossing*). The forbidden configuration and parity arguments on sums of those quantities bring the desired result that the two cycles cannot be both of odd length.

The complement of the union of two odd cycles is not a unit ball graph. Given a needle in \mathbb{R}^3 whose middle-point is attached to the origin, one can apply a continuous motion to turn it around (a motion à la Kakeya, henceforth *Kakeya motion*). A Kakeya motion can be seen as a closed antipodal curve on the 2-sphere. If we now consider two needles, each with a Kakeya motion, then the two needles have to go through a same position. This simply follows from the fact that two antipodal curves on the 2-sphere intersect. The second main result of this article is a translation of this Jordan-type theorem in terms of intersection graphs: The complement of a unit ball graph does not contain the disjoint union of two odd cycles. The proof can really be seen as two Kakeya

¹The NP-hardness, ruling out a 1-approximation, is still to be shown.

motions, each one along the two odd cycles, leading to a contradiction when the needles achieve parallel directions.

Together with the next result, it implies a randomized EPTAS for the following problem: Given a set S of points in \mathbb{R}^3 , find a largest subset of S of diameter at most 1.

EPTAS for MAXIMUM INDEPENDENT SET on $\mathcal{X}(d, \beta, 1)$. We show that if a graph G satisfies that every two odd cycles are joined by an edge, the Vapnik-Chervonenkis dimension (VC-dimension for short) of the hypergraph of the neighborhoods of G is bounded, and $\alpha(G)$ is at least a constant fraction of $|V(G)|$, then $\alpha(G)$ can be approximated in polynomial time at any given precision. More precisely, we present in that case a randomized EPTAS running in time $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$ and a deterministic PTAS.

Our algorithm works as follows: We start by sampling a small subset of vertices. Hoping that this small subset is entirely contained in a fixed optimum solution I , we include the selected vertices to our solution and remove their neighborhood from the graph. Due to the classic result of Haussler and Welzl [39] on ε -nets of size $O(d/\varepsilon \log 1/\varepsilon)$ (where d is the VC-dimension), this sampling lowers the degree in I of the remaining vertices. We compute a shortest odd cycle. If this cycle is short, we can remove its neighborhood from the graph and solve optimally the problem in the resulting graph, which is bipartite by assumption. If this cycle is long, we can efficiently find a small odd-cycle transversal. This is shown by a careful analysis on the successive neighborhoods of the cycle, and the recurrent fact that this cycle is a shortest among the ones of odd length.

Organization. The rest of the article is organized as follows: In Section 2, we recall some relevant notations for graphs and elementary geometry, the definitions of VC-dimension, disk graphs, and approximation schemes. We finish this section by introducing a class $\mathcal{X}(d, \beta, i)$ of graphs parameterized by three constants: d , upper-bounding the VC-dimension, β , lower-bounding the ratio $\alpha(G)/|V(G)|$ (*independence number divided by number of vertices*), and i , upper-bounding the maximum number of odd cycles that can be found as a disjoint union in an induced subgraph. In Section 3, we characterize disk graphs whose complements have maximum degree 2. The most useful result of this section is that the complement of the disjoint union of two odd cycles is not a disk graph. In Section 4, we show that unit ball graphs have the same forbidden induced subgraphs. In Section 5, we show that those forbidden induced subgraphs, together with existing results and algorithms, readily give a subexponential algorithm and a QPTAS for MAXIMUM CLIQUE on disk and unit ball graphs. In Section 6, we design a randomized EPTAS for MAXIMUM INDEPENDENT SET on the class $\mathcal{X}(d, \beta, i)$. We then show how this yields a randomized EPTAS for MAXIMUM CLIQUE on disk and unit ball graphs. This is tight in two directions: having different values of radii, and the dimension of the ambient space. Indeed, in Section 7, we complement those positive results by showing that MAXIMUM CLIQUE is unlikely to have a QPTAS (even a SUBEXPAS) on ball graphs where all the radii are arbitrarily close to 1, and on 4-dimensional unit ball graphs. We also show the same lower bounds for intersection graphs of simple objects in the plane, namely, filled triangles and filled ellipses. In Section 8, we make some observations about the EPTAS and propose some lines of thoughts on how to tackle the computational complexity of MAXIMUM CLIQUE in disk and unit ball graphs.

2 PRELIMINARIES

For two integers $i \leq j$, we denote by $[i, j]$ the set of integers $\{i, i + 1, \dots, j - 1, j\}$. For a positive integer i , we denote by $[i]$ the set of integers $[1, i]$.

The *Exponential Time Hypothesis* (ETH) is a conjecture by Impagliazzo et al. asserting that there is no $2^{o(n)}$ -time algorithm for 3-SAT on instances with n variables [40]. The ETH, together with the sparsification lemma [40], even implies that there is no $2^{o(n+m)}$ -time algorithm solving 3-SAT.

Graph notations. Let G be a simple graph. We denote by \overline{G} its complement, i.e., the graph obtained by making every non-edge an edge and vice versa. $V(G)$ and $E(G)$ represent its set of vertices and its set of edges, respectively. We denote by $\alpha(G)$ the *independence number* of G , i.e., the size of a maximum independent set (or stable set), and by $\omega(G)$ the *clique number* of G , i.e., the size of a maximum clique. For $S \subseteq V(G)$, its open neighborhood, denoted by $N_G(S)$, is the set of vertices that are not in S and have a neighbor in S , and its closed neighborhood is defined by $N_G[S] = S \cup N_G(S)$. We omit the subscript G if the graph is obvious from the context and we write $N_G(x)$ instead of $N_G(\{x\})$.

The *odd cycle packing number* of G , denoted by $\text{ocp}(G)$, is defined as the maximum number of vertex-disjoint odd cycles and the *induced odd cycle packing number* of G , denoted by $\text{iocp}(G)$, is the maximum number of vertex-disjoint odd cycles with no edge between any two of them.

The *2-subdivision* of a graph G is the graph H obtained by subdividing each edge of G exactly twice. If G has n vertices and m edges, then H has $n + 2m$ vertices and $3m$ edges. The *co-2-subdivision* of G is the complement of H . Hence, it has $n + 2m$ vertices and $\binom{n+2m}{2} - 3m$ edges. The *co-degree* of a graph is the maximum degree of its complement. A *co-disk* is a graph that is the complement of a disk graph.

VC-dimension. VC-dimension has been introduced by Vapnik and Chervonenkis in the seminal paper [61]. Let $H = (V, E)$ be a hypergraph. A set $X \subseteq V$ of vertices of H is *shattered* if for every subset Y of X there exists a hyperedge $e \in E$ such that $e \cap X = Y$. An intersection between X and a hyperedge e of E is called a *trace* (on X). Equivalently, a set X is shattered if all its $2^{|X|}$ traces exist. The *VC-dimension* of a hypergraph is the maximum size of a shattered set. As an abuse of language, we call VC-dimension of a graph G , denoted by $\text{VCdim}(G)$, the VC-dimension of the neighborhood hypergraph $(V(G), \{N_G(v) \mid v \in V(G)\})$. The *geometric VC-dimension* of a collection of objects is the VC-dimension of the (uncountable infinite) hypergraph where vertices are all the points of the ambient space and the hyperedges are the objects. If, in general, the geometric VC-dimension does not necessarily coincide with the VC-dimension of the intersection graph, it does coincide for unit balls in any dimension. Indeed, shattering a set S of points with balls of radius 1 is equivalent to shattering the balls of radius $1/2$ centered at S with balls of radius $1/2$.

Geometric notations. For a positive integer d , we denote by \mathbb{R}^d the d -dimensional Euclidean space. We denote by $d(x, y)$ the Euclidean distance between x and y . If x and y are two points of \mathbb{R}^d , $\text{seg}(x, y)$, or simply xy , is the straight-line segment whose endpoints are x and y . If x and y are distinct, $\ell(x, y)$ is the unique line going through x and y . If s is a segment with positive length, then we denote by $\ell(s)$ the unique line containing s .

We will often define disks and elliptical disks by their boundary, i.e., circles and ellipses, and also use the following basic facts: There are exactly two circles that pass through a given point with a given tangent at this point and a given radius; one if we further specify on which side of the tangent the circle is. There is exactly one circle that passes through two points with a given tangent at one of the two points, provided the other point is *not* on this tangent. Finally, there exists one (not necessarily unique) ellipse that passes through two given points with two given tangents at those points.

A *d -dimensional closed ball* is defined from a center $x \in \mathbb{R}^d$ and a radius $r \in \mathbb{R}^+$, as the set of points $\{y \in \mathbb{R}^d \mid d(x, y) \leq r\}$, i.e., at distance at most r from x . The *diameter* of a subset $S \subseteq \mathbb{R}^d$ is defined as $\sup_{x, y \in S} d(x, y)$. The *piercing number* (also called *hitting set* or *transversal*) of a collection \mathcal{O} of geometric objects in \mathbb{R}^d is the minimum number of points of \mathbb{R}^d that *pierce* (or *hit*) all the objects of \mathcal{O} , i.e., each object contains at least one of these points.

Disk graphs and their forbidden induced subgraphs. A *d -dimensional ball graph* is the intersection graph of d -dimensional closed balls of \mathbb{R}^d . We shorten 2-dimensional ball graph in *disk graph*, and

3-dimensional ball graph in *ball graph*. A d -dimensional unit ball graph is the intersection graph of unit d -dimensional closed balls of \mathbb{R}^d , that is, balls with radius 1. Unit d -dimensional ball graphs can be thought of only with points: vertices are points (at the center of the ball) and two points are adjacent if they are at distance at most 2. In particular, solving MAXIMUM CLIQUE on those graphs is equivalent to finding a maximum sub-collection of points whose diameter is at most a fixed value.

Approximation schemes. A PTAS (Polynomial-Time Approximation Scheme) for a minimization (respectively, maximization) problem is an approximation algorithm that takes an additional parameter $\varepsilon > 0$ and outputs in time $n^{f(\varepsilon)}$ a solution of value at most $(1 + \varepsilon)\text{OPT}$ (respectively, at least $(1 - \varepsilon)\text{OPT}$) where OPT is the optimum value. Observe that from now on, we consider that approximation ratios of maximization problems are smaller than 1, unlike the convention we used in the introduction. An EPTAS (Efficient PTAS) is the same with running time $f(\varepsilon)n^{O(1)}$, an FPTAS (Fully PTAS) has running time $(1/\varepsilon)^{O(1)}n^{O(1)}$, a QPTAS (Quasi PTAS) has running time $n^{\text{polylog } n}$ for every $\varepsilon > 0$. Finally, and this is less standard, we call SUBEXPAS (subexponential AS) an approximation scheme with running time 2^{n^γ} for some $\gamma < 1$ not depending on ε , for every $\varepsilon > 0$. All those approximation schemes can come deterministic or randomized.

The class $\mathcal{X}(d, \beta, i)$. In Section 6, we present a randomized EPTAS and a deterministic PTAS for approximating the independence number α on graphs with constant VC-dimension, linear independence number, and the induced odd cycle packing number equal to 1.

Actually, we extend these algorithms to the case $\text{iocp}(G) = i$, for any constant i . Let $\mathcal{X}(d, \beta, i)$ be the class of simple graphs G satisfying:

- $\text{VCdim}(G) \leq d$,
- $\alpha(G) \geq \beta|V(G)|$, and
- $\text{iocp}(G) \leq i$.

For any positive constants $d, \beta < 1, i$, we get a deterministic PTAS and a randomized EPTAS for MAXIMUM INDEPENDENT SET on $\mathcal{X}(d, \beta, i)$.

3 DISK GRAPHS WITH CO-DEGREE AT MOST 2

In this section, we fully characterize the degree-2 complements of disk graphs. We show the following:

THEOREM 3.1. *A disjoint union of paths and cycles is the complement of a disk graph if and only if the number of odd cycles is at most one.*

We split this theorem into two parts. In the first one, Section 3.1, we show that the union of two disjoint odd cycles is not the complement of a disk graph. This is the part that will be algorithmically useful. As disk graphs are closed under taking induced subgraphs, it implies that in the complement of a disk graph two vertex-disjoint odd cycles have to be linked by at least one edge. This will turn out useful when solving MAXIMUM INDEPENDENT SET on the complement of the graph (to solve MAXIMUM CLIQUE on the original graph). In the second part, Section 3.2, we show how to represent the complement of the disjoint union of even cycles and exactly one odd cycle. Although this result is not needed for the forthcoming algorithmic section, it nicely highlights the singular role that parity plays and exposes the complete set of disk graphs of co-degree at most 2.

3.1 The Disjoint Union of Two Odd Cycles Is Not Co-disk

We call *positive distance* between two non-intersecting disks the minimum of $d(x, y)$ where x is in one disk and y is in the other. If the disks are centered at c_1 and c_2 with radius r_1 and r_2 , respectively,




Fig. 1. Disk realization of a $K_{2,2}$. As the centers are positioned, it is impossible that the two non-edges are between the disks 2 and 3, and between the disks 1 and 4 (or between the disks 1 and 3, and between the disks 2 and 4).

then this value is $d(c_1, c_2) - r_1 - r_2$. We call *negative distance* between two intersecting disks the length of the straight-line segment defined as the intersection of three objects: the two disks and the line joining their centers. This value is $r_1 + r_2 - d(c_1, c_2)$, which is positive.

We call *proper representation* a disk representation where every edge is witnessed by a proper intersection of the two corresponding disks, i.e., the interiors of the two disks intersect. It is easy to transform a disk representation into a proper representation (of the same graph).

LEMMA 3.2. *If a graph has a disk representation, then it has a proper representation.*

PROOF. If two disks intersect non-properly, we increase the radius of one of them by $\varepsilon/2$ where ε is the smallest positive distance between any two disks in the representation. \square

To not have to discuss about the special case of three aligned centers in a disk representation, we show that such a configuration is never needed to represent a disk graph.

LEMMA 3.3. *If a graph has a disk representation, it has a proper representation where no three centers are aligned.*

PROOF. By Lemma 3.2, we have or obtain a proper representation. Let ε be the minimum between the smallest positive distance and the smallest negative distance. As the representation is proper, $\varepsilon > 0$. If three centers are aligned, we move one of them to any point that is not lying in a line defined by two centers in a ball of radius $\varepsilon/2$ centered at it. This decreases the number of triples of aligned centers by at least one and can be repeated until no three centers are aligned. \square

From now on, we assume that every disk representation is proper and without three aligned centers. We show the folklore result that in a representation of $K_{2,2}$ that sets the four centers in convex position, both non-edges have to be *diagonal*.

LEMMA 3.4. *In a disk representation of $K_{2,2}$ with the four centers in convex position, the non-edges are between vertices corresponding to opposite centers in the quadrangle.*

PROOF. Let c_1 and c_2 be the centers of one non-edge, and c_3 and c_4 the centers of the other non-edge. Let r_i be the radius associated to center c_i for $i \in [4]$. It should be that $d(c_1, c_2) > r_1 + r_2$ and $d(c_3, c_4) > r_3 + r_4$ (see Figure 1). Assume c_1 and c_2 are consecutive on the convex hull formed by $\{c_1, c_2, c_3, c_4\}$, and say, without loss of generality, that the order is c_1, c_2, c_3, c_4 . Let c be the intersection of $\text{seg}(c_1, c_3)$ and $\text{seg}(c_2, c_4)$. It holds that $d(c_1, c_3) + d(c_2, c_4) = d(c_1, c) + d(c, c_3) + d(c_2, c) + d(c, c_4) = (d(c_1, c) + d(c, c_2)) + (d(c_3, c) + d(c, c_4)) > d(c_1, c_2) + d(c_3, c_4) > r_1 + r_2 + r_3 + r_4 = (r_1 + r_3) + (r_2 + r_4)$. Which implies that $d(c_1, c_3) > r_1 + r_3$ or $d(c_2, c_4) > r_2 + r_4$; a contradiction. \square

We derive a useful consequence of the previous lemma, phrased in terms of intersections of lines and segments.

COROLLARY 3.5. *In any disk representation of $K_{2,2}$ with centers c_1, c_2, c_3, c_4 with the two non-edges between the vertices corresponding to c_1 and c_2 , and between c_3 and c_4 , it should be that $\ell(c_1, c_2)$ intersects $\text{seg}(c_3, c_4)$ or $\ell(c_3, c_4)$ intersects $\text{seg}(c_1, c_2)$.*

PROOF. First, assume that the disk representation has the four centers in convex position. Then, by Lemma 3.4, $\text{seg}(c_1, c_2)$ and $\text{seg}(c_3, c_4)$ are the diagonals of a convex quadrangle. Hence, they intersect, and *a fortiori*, $\ell(c_1, c_2)$ intersects $\text{seg}(c_3, c_4)$ ($\ell(c_3, c_4)$ intersects $\text{seg}(c_1, c_2)$, too).

In the other case, the disk representation has one center, say without loss of generality, c_1 , in the interior of the triangle formed by the other three centers. In this case, $\ell(c_1, c_2)$ intersects $\text{seg}(c_3, c_4)$. If instead a center in $\{c_3, c_4\}$ is in the interior of the triangle formed by the other centers, then $\ell(c_3, c_4)$ intersects $\text{seg}(c_1, c_2)$. \square

We can now prove the main result of this section, thanks to the previous corollary, parity arguments, and *some elementary properties of closed plane curves*, namely, Property I and Property III of the eponymous paper [58].

THEOREM 1.1. *A complement of a disk graph cannot have a disjoint union of two odd cycles as an induced subgraph. In other words, if G is a disk graph, then $\text{iocp}(\overline{G}) \leq 1$.*

PROOF. Let s and t be two positive integers and $G = \overline{C_s + C_t}$ the complement of the disjoint union of a cycle of length s and a cycle of length t . Assume that G is a disk graph. Let C_1 (respectively, C_2) be the cycle embedded in the plane formed by s (respectively, t) straight-line segments joining the consecutive centers of disks along the first (respectively, second) cycle. Observe that the segments of those two cycles correspond to the non-edges of G . We number the segments of C_1 from S_1 to S_s , and the segments of C_2 , from S'_1 to S'_t .

For the i th segment S_i of C_1 , let a_i be the number of segments of C_2 intersected by the line $\ell(S_i)$ prolonging S_i , let b_i be the number of segments S'_j of C_2 such that the prolonging line $\ell(S'_j)$ intersects S_i , and let c_i be the number of segments of C_2 intersecting S_i . For the second cycle, we define similarly a'_j, b'_j, c'_j . The quantity $a_i + b_i - c_i$ counts the number of segments of C_2 that can possibly represent a $K_{2,2}$ with S_i according to Corollary 3.5. As we assumed that G is a disk graph, $a_i + b_i - c_i = t$ for every $i \in [s]$. Otherwise, there would be at least one segment S'_j of C_2 such that $\ell(S_i)$ does not intersect S'_j and $\ell(S'_j)$ does not intersect S_i .


By summing, it holds that $\sum_{i=1}^s (a_i + b_i - c_i) = st$. Observe that a_i is an even integer, since C_2 is a closed curve. This implies that st and $\sum_{i=1}^s (b_i - c_i)$ have the same parity. Note that $\sum_{i=1}^s c_i$ counts the number of intersections of the two closed curves C_1 and C_2 , and is therefore even. Finally, observe that $\sum_{i=1}^s b_i = \sum_{j=1}^t a'_j$ by reordering and reinterpreting the sum from the point of view of the segments of C_2 . Since the a'_j 's are all even, $\sum_{i=1}^s b_i$ is also even.

This means then st is even. Therefore, s and t cannot be both odd integers. \square

3.2 The Disjoint Union of Cycles with at Most One Odd is Co-disk

We only show the following part of Theorem 3.1 to emphasize that, rather unexpectedly, parity plays a crucial role in disk graphs of co-degree at most 2. It is also amusing that the complement of any odd cycle is a *unit* disk graph while the complement of any even cycle of length at least 8 is not [10]. Here, the situation is somewhat reversed: Complements of even cycles are *easier* to represent than complements of odd cycles.

THEOREM 3.6. *The complement of the disjoint union of even cycles and one odd cycle is a disk graph.*



(a) Three important disks with the same size $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{2s}$.

(b) Zoom where \mathcal{D}_1 almost touches \mathcal{D}_2 and \mathcal{D}_{2s} .

Fig. 2. The disks $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_{2s}$ and the convex chain p_1, p_2, \dots, p_s . The curvature of the boundary of \mathcal{D}_1 is exaggerated in the zoom for the sake of clarity.

PROOF. We start with a disk representation of the complement of one even cycle C_{2s} . Recall that this construction is not possible with *unit* disks for even cycles of length at least 8. We assume that the vertices of the cycle C_{2s} are $1, 2, \dots, 2s$ in this order. For each $i \in [2s]$, the disk \mathcal{D}_i encodes the vertex i . We start by fixing the disks $\mathcal{D}_1, \mathcal{D}_2$, and \mathcal{D}_{2s} . Those three disks have the same radius. We place \mathcal{D}_2 and \mathcal{D}_{2s} side by side: Their centers have the same y -coordinate. They intersect and the distance between their centers is $\varepsilon > 0$. We define \mathcal{D}_1 as the disk above \mathcal{D}_2 and \mathcal{D}_{2s} tangent to those two disks and sharing the same radius. We denote by p_1 its intersection with \mathcal{D}_2 and by p_s its intersection with \mathcal{D}_{2s} . We then slightly shift \mathcal{D}_1 upward so it does not touch (nor does it intersect) \mathcal{D}_2 and \mathcal{D}_{2s} anymore. While we do this translation, we imagine that the points p_1 and p_s remain fixed at the boundary of \mathcal{D}_2 and \mathcal{D}_{2s} , respectively (see Figure 2(a)). Let p_2, p_3, \dots, p_{s-1} points in the interior of \mathcal{D}_1 and below the line $\ell(p_1, p_s)$ such that $p_1, p_2, \dots, p_{s-1}, p_s$ form an x -monotone convex chain (see Figure 2(b)).

Now, we define the disks $\mathcal{D}_4, \mathcal{D}_6, \dots, \mathcal{D}_{2s-2}$. For each $i \in \{4, 6, \dots, 2s - 2\}$, let \mathcal{D}_i be the unique disk with the same radius as \mathcal{D}_2 and such that the boundary of \mathcal{D}_i crosses $p_{i/2}$ and is below its tangent $\tau_{i/2}$ at this point, which has the direction of $\ell(p_{i/2-1}, p_{i/2+1})$.

It should be observed that the only disk with even index i that contains $p_{i/2}$ is \mathcal{D}_i . We can further choose the convex chain $\{p_i\}_{i \in [s]}$ such that one co-tangent $\tau_{i,i+1}$ to \mathcal{D}_{2i} and \mathcal{D}_{2i+2} has a slope between the slopes of τ_i and τ_{i+1} . Finally, we define the disks $\mathcal{D}_3, \mathcal{D}_5, \dots, \mathcal{D}_{2s-1}$. For each $i \in \{3, 5, \dots, 2s - 1\}$, let \mathcal{D}_i be tangent to $\tau_{i,i+1}$ at the point of x -coordinate the mean between the x -coordinates of $p_{\frac{i-1}{2}}$ and $p_{\frac{i+1}{2}}$. Moreover, \mathcal{D}_i is above $\tau_{i,i+1}$ and has a radius sufficiently large to intersect every disk with even index that are not \mathcal{D}_{i-1} and \mathcal{D}_{i+1} . It is easy to see that the disks \mathcal{D}_i with even index (respectively, odd index) form a clique. By construction, the disk \mathcal{D}_i with odd index greater than 3 intersects every disk with even index except \mathcal{D}_{i-1} and \mathcal{D}_{i+1} , since \mathcal{D}_i is on the other side of $\tau_{i,i+1}$ than those two disks. As the line $\tau_{i,i+1}$ intersects every other disk with even index, there is a sufficiently large radius so \mathcal{D}_i does so, too. The particular case of \mathcal{D}_1 has been settled at the beginning of the construction. This disk avoids \mathcal{D}_2 and \mathcal{D}_{2s} and contains p_2, p_3, \dots, p_{s-1} , so intersects all the other disks with even index.

We now explain how to *stack* even cycles. We make the distance ε between the center of \mathcal{D}_2 and \mathcal{D}_{2s} a thousandth of their common radius. Note that this distance does not depend on the value of s . We identify the small region (point) where the disk \mathcal{D}_1 intersects with the disks of even index, between two different complements of cycles. We then rotate from this point one representation by a small angle (see Figure 3 for multiple complements of even cycles stacked).

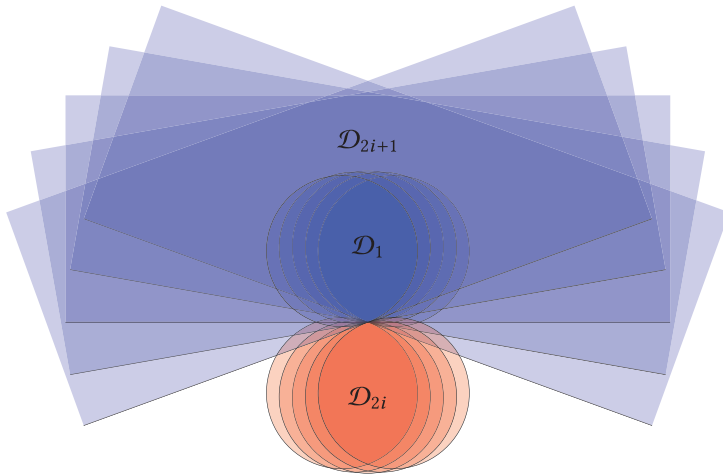
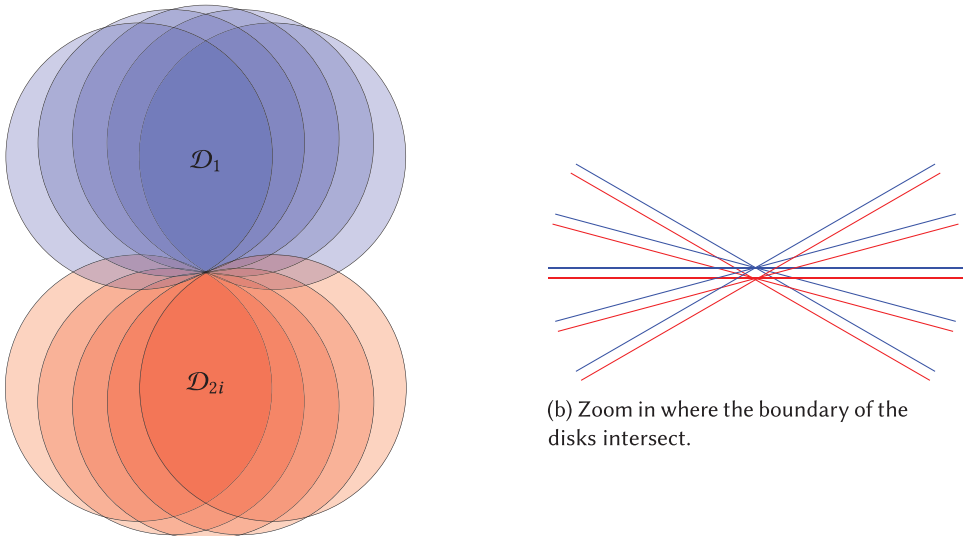


Fig. 3. A disk realization of the complement of the disjoint union of an arbitrary number of even cycles.



(a) The only potential non-edges are between two disks represented almost tangent.

(b) Zoom in where the boundary of the disks intersect.

Fig. 4. Zoom in where the disk \mathcal{D}_1 of the several complements of even cycles intersects all the \mathcal{D}_{2i} of the other cycles.

The reason why there are indeed all the edges between two complements of cycles is intuitive and depicted in Figure 4 and more specifically Figure 4(b). We superimpose all the complements of even cycles in a way that the maximum rotation angle between two complements of cycles is small (see for instance Figure 7).

Finally, we need to add one disjoint odd cycle in the complement. There is a nice representation of a complement of an odd cycle by unit disks in the paper of Atminas and Zamaraev [10] (see Figure 5).




Fig. 5. A disk realization of the complement of an odd cycle with unit disks as described by Atminas and Zamaraev [10]. Unfortunately, we cannot use this representation.




Fig. 6. A disk realization of the complement of an odd cycle of length $2s + 1$.

However, we will use a different and non-unit representation for the next step to work. Let $2s + 1$ be the length of the cycle. We use a similar construction as for the complement of an even cycle. We denote the disks $\mathcal{D}'_1, \mathcal{D}'_2, \dots, \mathcal{D}'_{2s+1}$. The difference is that we separate \mathcal{D}'_1 away from \mathcal{D}'_2 but not from \mathcal{D}'_{2s} . Then, we represent all the disks with odd index but \mathcal{D}'_{2s+1} as before. The disk \mathcal{D}'_{2s+1} is chosen as being cotangent to \mathcal{D}'_1 and \mathcal{D}'_{2s} and to the left of them. Then, we very slightly move \mathcal{D}'_{2s+1} to the left so it does not intersect those two disks anymore. The disk \mathcal{D}'_{2s} has the rightmost center among the disks with even index. Therefore, \mathcal{D}'_{2s+1} still intersects all the other disks of even index.

Moreover, the disks with even index form a clique and the disks with odd index form a clique minus an edge between the vertex 1 and the vertex $2s + 1$. Hence, the intersection graph of those disks is indeed the complement of C_{2s+1} (see Figure 6).

This representation of $\overline{C_{2s+1}}$ can now be put on top of complements of even cycles. We identify the small region (point) where the disk \mathcal{D}_1 intersects the disks of even index (in complements of even cycles) with the small region (point) where the disk \mathcal{D}'_1 intersects the disks of even index (in the one complement of odd cycle). We make the disk \mathcal{D}'_1 significantly smaller than \mathcal{D}_1 and rotate the representation of $\overline{C_{2s+1}}$ by a sizable angle, say, 60 degrees (see Figure 7).

It is easy to see that the disks of the complement of the odd cycle intersect all the disks of the complements of even cycles. A good sanity check is to observe why we cannot stack representations of complements of odd cycles with the same rotation scheme. In Figure 8, the rotation of two




Fig. 7. Placing the complement of odd cycle on top of the complements of even cycles.




Fig. 8. Sanity check: trying to stack the complements of two odd cycles fails. The disks \mathcal{D}'_1 and $\mathcal{D}''_{2s'+1}$ do not intersect.

representations of the complement of an odd cycle leaves disks \mathcal{D}'_1 and $\mathcal{D}''_{2s'+1}$ far apart when they should intersect. \square

Theorem 1.1 and Theorem 3.6, together with the fact that disk graphs are closed by taking induced subgraphs prove Theorem 3.1.

4 SAME OBSTRUCTION FOR UNIT BALL GRAPHS

In this section, we show that unit ball graphs, like disk graphs, do not contain the complement of two disjoint odd cycles. In other words, for any unit ball graph G , $\text{iocp}(\overline{G}) \leq 1$. It is interesting

to note that the proofs we could find for disk graphs and unit ball graphs turn out to be quite different.

A *closed polygonal chain* C in \mathbb{R}^d is defined by a sequence of points (or vertices) $x_1, x_2, \dots, x_p \in \mathbb{R}^d$ as the straight-edge segments $x_1x_2, x_2x_3, \dots, x_{p-1}x_p, x_px_1$. We call *direction* of a non-zero vector its equivalence class by the relation $\vec{u} \sim \vec{v} \Leftrightarrow \exists \lambda \in \mathbb{R}^+, \vec{u} = \lambda\vec{v}$. We denote the direction of \vec{u} by $\text{dir}(\vec{u})$. We define the set of directions

$$\text{Needle}(C) := \bigcup_{1 \leq i \leq p} [x_{i-1} \leftarrow x_i \rightarrow x_{i+1}] \cup [x_{i-1} \rightarrow x_i \leftarrow x_{i+1}],$$

where the indices are taken modulo p , $[x_{i-1} \leftarrow x_i \rightarrow x_{i+1}] := \{\text{dir}(\overrightarrow{x_ix}) \mid x \in x_{i-1}x_{i+1}\}$, and $[x_{i-1} \rightarrow x_i \leftarrow x_{i+1}] := \{\text{dir}(\overrightarrow{xx_i}) \mid x \in x_{i-1}x_{i+1}\}$.

LEMMA 4.1. *Let C_1 and C_2 be two closed polygonal chains of \mathbb{R}^3 on an odd number of vertices each. Then, $\text{Needle}(C_1) \cap \text{Needle}(C_2)$ is non-empty.*

PROOF. We want to establish the existence of a direction that is common to $\text{Needle}(C_1)$ and $\text{Needle}(C_2)$. We identify $\text{Needle}(C_i)$ ($i \in \{1, 2\}$) to its trace on the 2-sphere. Indeed, the set of directions in \mathbb{R}^3 is isomorphic to the set of points on the 2-sphere.

Let x_1, x_2, \dots, x_p be the vertices of C_1 . We show that $\text{Needle}(C_1)$ is path-connected (due to p being odd). To do so, we continuously modify (within $\text{Needle}(C_1)$) the initial vector $\text{dir}(\overrightarrow{x_1x_2})$ into any other vector of $\text{Needle}(C_1)$. We start with a in x_1 and b in x_2 . We continuously move a from x_1 to x_3 (on the straight-edge segment x_1x_3) while b stays fixed at x_2 (we are moving in $[x_1 \leftarrow x_2 \rightarrow x_3]$). For the next step, a is fixed at x_3 and b continuously moves from x_2 to x_4 . (we are moving in $[x_2 \rightarrow x_3 \leftarrow x_4]$). And in general, we move the point with index $i - 1$ from x_{i-1} to x_{i+1} while the other point stays fixed at x_i (where the indices are modulo p). Since p is odd, we reach the situation where b is set to x_1 and a is set to x_2 , when we have completed once the walk on the closed polygonal chain. We repeat a walk on the chain once again, so a is back to x_1 and b is back to x_2 , and we stop. This process spans $\text{Needle}(C_1)$.

Therefore, C_1 is a closed curve on the 2-sphere, since we are finally back to $\text{dir}(\overrightarrow{x_1x_2})$, that is, from where we started. Furthermore, C_1 is *antipodal*, i.e., closed by taking antipodal points. Indeed, for each direction attained in a $[x_{i-1} \leftarrow x_i \rightarrow x_{i+1}]$, we reach the opposite direction in $[x_{i-1} \rightarrow x_i \leftarrow x_{i+1}]$. Similarly $\text{Needle}(C_2)$ draws a closed antipodal curve C_2 on the 2-sphere. The curves C_1 and C_2 intersect, since they are closed and antipodal. An intersection point corresponds to a direction shared by $\text{Needle}(C_1)$ and $\text{Needle}(C_2)$. \square

We will apply this lemma on the closed polygonal chains C_1 and C_2 formed by the centers of unit balls realizing two odd cycle complements. The contradiction will come from the fact that not all the pairs of centers $x \in C_1$ and $y \in C_2$ can be at distance at most 2.

THEOREM 1.2. *A complement of a unit ball graph cannot have a disjoint union of two odd cycles as an induced subgraph. In other words, if G is a unit ball graph, then $\text{iocp}(\overline{G}) \leq 1$.*

PROOF. Let $x_1, x_2, \dots, x_p \in \mathbb{R}^3$ (respectively, $y_1, y_2, \dots, y_q \in \mathbb{R}^3$) be the centers of unit balls representing the complement of an odd cycle of length p (respectively, q), such that x_i and x_{i+1} (respectively, y_i and y_{i+1}) encode the non-adjacent pairs. Let C_1 (respectively, C_2) be the closed polygonal chain obtained from the centers x_1, x_2, \dots, x_p (respectively, y_1, y_2, \dots, y_q). By Lemma 4.1, there are two collinear vectors $\overrightarrow{x_ix}$ and $\overrightarrow{y_jy}$ with x on the straight-line segment $x_{i-1}x_{i+1}$ and y on the straight-line segment $y_{j-1}y_{j+1}$.




Fig. 9. The two cases for the collinear vectors $\overrightarrow{x_i x}$ and $\overrightarrow{y_j y}$.

Let us suppose that $\overrightarrow{x_i x}$ and $\overrightarrow{y_j y}$ have the same direction (Figure 9, left). In the plane² containing $x_i, x, y_j,$ and y , those four points are in convex position and the convex hull is cyclically ordered x_i, x, y, y_j . We obtain a contradiction by showing that the sum of the diagonals $d(x_i, y) + d(y_j, x)$ is strictly smaller than the sum of two opposite sides $d(x_i, x) + d(y_j, y)$.

Considering edges and non-edges, for every i , the only vertices at distance greater than 2 from y_i in $\cup\{x_j, y_j\}$ are y_{i-1} and y_{i+1} . Then, we have $d(x_i, y_{j-1}) \leq 2 < d(y_j, y_{j-1})$ and $d(x_i, y_{j+1}) \leq 2 < d(y_j, y_{j+1})$. The points strictly closer to x_i than to y_j form an open half-space. In particular, they form a convex set and all the points in the segment $y_{j-1}y_{j+1}$ are therefore closer to x_i than to y_j . Hence, $d(x_i, y) < d(y_j, y)$. Symmetrically, $d(y_j, x) < d(x_i, x)$. So $d(x_i, y) + d(y_j, x) < d(y_j, y) + d(x_i, x)$, a contradiction.

Let us now assume that $\overrightarrow{x_i x}$ and $\overrightarrow{y_j y}$ have opposite direction (Figure 9, right). In that case, the four coplanar points x_i, x, y_j, y are in convex position in their plane and the convex hull is cyclically ordered x_i, x, y_j, y . We will attain the similar contradiction that $d(x, y) + d(x_i, y_j) < d(x_i, x) + d(y_j, y)$. As previously, $d(y_{j-1}, x_{i-1}) \leq 2 < d(x_i, x_{i-1})$ and $d(y_{j-1}, x_{i+1}) \leq 2 < d(x_i, x_{i+1})$. Hence, by convexity of the set of points closer to y_{j-1} than to x_i , we have $d(y_{j-1}, x) < d(x_i, x)$. Similarly, $d(y_{j+1}, x) < d(x_i, x)$. We obtained that point x is closer to y_{j-1} and y_{j+1} than to x_i . Therefore, applying again the convexity argument, we get that $d(x, y) < d(x, x_i)$. Besides, $d(x_i, y_j) \leq 2 < d(y, y_j)$. So, we arrive at the contradiction $d(x, y) + d(x_i, y_j) < d(x, x_i) + d(y, y_j)$. \square

5 DIRECT ALGORITHMIC CONSEQUENCES

Now, we show how to use the structural results from Section 3 and Section 4 to obtain algorithms for MAXIMUM CLIQUE in disk and unit ball graphs. A clique in a graph G is an independent set in \overline{G} . So, leveraging the result from Theorem 1.1 and Theorem 1.2, we will focus on solving MAXIMUM INDEPENDENT SET in graphs without two vertex-disjoint odd cycles as an induced subgraph.

5.1 QPTAS

Here, we present a simple argument to get a QPTAS for MAXIMUM CLIQUE in disk and unit ball graphs due to a known approximation algorithm for MAXIMUM INDEPENDENT SET on graphs with relatively small odd cycle packing number. In Section 6, we will improve the QPTAS to a randomized EPTAS using a different approach. We recall that the odd cycle packing number $\text{ocp}(H)$ of a graph H is the maximum number of vertex-disjoint odd cycles in H . Unfortunately, the condition that \overline{G} does not contain two vertex-disjoint odd cycles as an induced subgraph is not quite the same as saying that the odd cycle packing number of \overline{G} is 1. Otherwise, we would immediately get a PTAS by the following result of Bock et al. [14]:

²Or a plane if it is not unique.

THEOREM 5.1 (BOCK ET AL. [14]). *For every fixed $\varepsilon > 0$ there is a polynomial $(1 - \varepsilon)$ -approximation algorithm for MAXIMUM INDEPENDENT SET for graphs H with n vertices and $\text{ocp}(H) = o(n/\log n)$.*

The algorithm by Bock et al. [14] works in polynomial time if $\text{ocp}(H) = o(n/\log n)$, but it does not need the odd cycle packing explicitly given as an input. This is important, since finding a maximum odd cycle packing is NP-hard [42]. We start by proving a structural lemma, which spares us having to determine the odd cycle packing number.

LEMMA 5.2. *Let H be a graph with n vertices, whose complement is a disk graph (respectively, unit ball graph). If $\text{ocp}(H) > n/\log^2 n$, then H has a vertex of degree at least $n/\log^4 n$.*

PROOF. Consider a maximum odd cycle packing \mathcal{P} . By assumption, it contains more than $n/\log^2 n$ vertex-disjoint cycles. Hence, a shortest cycle C in \mathcal{P} has size at most $\log^2 n$. Now, by Theorem 1.1 (respectively, by Theorem 1.2), H has no two vertex-disjoint odd cycles with no edges between them. Therefore, there must be an edge from C to every other cycle of \mathcal{P} , which constitutes at least $n/\log^2 n$ edges. Let v be a vertex of C with the maximum number of edges to other cycles in \mathcal{P} . By the pigeonhole principle, its degree is at least $n/\log^4 n$. \square

We are ready to suggest a QPTAS for MAXIMUM CLIQUE in disk and unit ball graphs.

THEOREM 5.3. *For any $\varepsilon > 0$, MAXIMUM CLIQUE can be $(1 - \varepsilon)$ -approximated in time $2^{O(\log^5 n)}$, when the input is a disk graph (respectively, unit ball graph) with n vertices.*

PROOF. Let G be the input disk graph (respectively, unit ball graph) and let \overline{G} be its complement, we want to find a $(1 - \varepsilon)$ -approximation for MAXIMUM INDEPENDENT SET in \overline{G} . We consider two cases. If \overline{G} has no vertex of degree at least $n/\log^4 n$, then, by Lemma 5.2, we know that $\text{ocp}(\overline{G}) \leq n/\log^2 n = o(n/\log n)$. In this case, we run the PTAS of Theorem 5.1 and we are done.

In the other case, \overline{G} has a vertex v of degree at least $n/\log^4 n$ (note that it may still be the case that $\text{ocp}(\overline{G}) = o(n/\log n)$). We branch on v : Either we include v in our solution and remove it and all its neighbors, or we discard v . The complexity of this step is described by the recursion $F(n) \leq F(n-1) + F(n - n/\log^4 n) \leq F(n-2) + 2F(n - n/\log^4 n) \leq \dots \leq (n/\log^4 n)F(n - n/\log^4 n)$. Thus, $F(n) \leq n^{O(\log^4 n)} = 2^{O(\log^5 n)}$. Note that this step is exact, i.e., we do not lose any solution. \square

We can actually generalize the QPTAS from graphs with $\text{iocp} \leq 1$ to graphs with $\text{iocp} = O(1)$.

LEMMA 5.4. *Let H be a sufficiently large graph with n vertices, and $\text{iocp}(H) = i$ for some integer i . If $\text{ocp}(H) > n/\log^2 n$, then H has a vertex of degree at least $\frac{n}{2i \log^5 n}$.*

PROOF. Again, let \mathcal{P} be an odd cycle packing ($|\mathcal{P}| \geq n/\log^2 n$). The number of cycles of \mathcal{P} with length at least $\log^3 n$ is at most $n/\log^3 n$. Thus, there are at least $n/\log^2 n - n/\log^3 n > \frac{1}{2}n/\log^2 n$ cycles of \mathcal{P} of length at most $\log^3 n$. We now consider the graph J whose vertices are the cycles of \mathcal{P} of length at most $\log^3 n$, and where there is an edge between C_1 and C_2 if there is at least one edge in H between a vertex of C_1 and a vertex of C_2 . Let $n_J := |V(J)| > \frac{1}{2}n/\log^2 n$.

By assumption, $\alpha(J) \leq i$ (since $\text{iocp}(H) = i$). Hence, by Turan's theorem, J has at most $(1 - \frac{1}{i})\binom{n_J}{2}$ non-edges. So J has at least $\frac{1}{i}\binom{n_J}{2}$ edges, and average degree at least $\frac{n_J-1}{i} \geq n/(2i \log^2 n)$. Let C_ℓ be a vertex of maximum degree in J . By construction, the cycle C_ℓ is adjacent to at least $n/(2i \log^2 n)$ other cycles of \mathcal{P} , and $|C_\ell| \leq \log^3 n$. Hence, the vertex of C_ℓ with the largest degree in H has at least $n/(2i \log^5 n)$ neighbors. \square

From the previous lemma, we get a QPTAS with slightly worse running time, similarly to Theorem 5.3.

THEOREM 5.5. *For any $\varepsilon > 0$ and for every integer d , MAXIMUM INDEPENDENT SET (respectively, MAXIMUM CLIQUE) can be $(1 - \varepsilon)$ -approximated in time $2^{O(i \log^6 n)}$, when the input is a graph G on n vertices with $\text{iocp}(G) = i$ (respectively, $\text{iocp}(\overline{G}) = i$).*

5.2 Subexponential Algorithm

The *odd girth* of a graph is the size of a shortest odd cycle. An *odd cycle cover* is a subset of vertices whose deletion makes the graph bipartite. We will use a result by Györi et al. [37], which says that graphs with large odd girth have small odd cycle cover. Bock et al. [14] turned the non-constructive proof into a polynomial-time algorithm.

THEOREM 5.6 (GYÖRI ET AL. [37], BOCK ET AL. [14]). *Let H be a graph with n vertices and no odd cycle shorter than δn (δ may be a function of n). Then there is an odd cycle cover X of size at most $(48/\delta) \ln(5/\delta)$. Moreover, X can be found in polynomial time.*

Let us show the following three options for an algorithm:

THEOREM 5.7. *Let G be a disk graph (respectively, unit ball graph) with n vertices. Let Δ be the maximum degree of \overline{G} and c the odd girth of \overline{G} (they may be functions of n). MAXIMUM CLIQUE has a branching or can be solved, up to a polynomial factor, in time:*

- (i) $2^{\tilde{O}(n/\Delta)}$ (branching), (ii) $2^{\tilde{O}(n/c)}$ (solved), (iii) $2^{O(c\Delta)}$ (solved).

PROOF. Let G be the input disk graph (respectively, unit ball graph) and let \overline{G} be its complement, we look for a maximum independent set in \overline{G} .

To prove (i), consider a vertex v of degree Δ in \overline{G} . We branch on v : Either we include v in our solution and remove $N[v]$, or discard v . The complexity is described by the recursion $F(n) \leq F(n - 1) + F(n - (\Delta + 1))$ and solving it gives (i). Observe that this does not give an algorithm running in time $2^{\tilde{O}(n/\Delta)}$, since the maximum degree might drop. Therefore, we will do this branching as long as it is *good enough* and then finish with the algorithms corresponding to (ii) and (iii).

For (ii) and (iii), let C be a cycle of length c ; it can be found in polynomial time (see for instance Reference [6]). By application of Theorem 5.6 with $\delta = c/n$, we find an odd cycle cover X in \overline{G} of size $\tilde{O}(n/c)$ in polynomial time. Next, we exhaustively guess in time $2^{\tilde{O}(n/c)}$ the intersection I of an optimum solution with X and finish by finding a maximum independent set in the bipartite graph $\overline{G} - (X \cup N(I))$, which can be done in polynomial time. The total complexity of this case is $2^{\tilde{O}(n/c)}$, which shows (ii).

Finally, observe that the graph $\overline{G} - N[C]$ is bipartite, since otherwise \overline{G} contains two vertex-disjoint odd cycles with no edges between them. Moreover, since every vertex in \overline{G} has degree at most Δ , it holds that $|N[C]| \leq c(\Delta - 1) \leq c\Delta$. Indeed, a vertex of C can only have $c(\Delta - 2)$ neighbors outside C . We can proceed as in the previous step: We exhaustively guess the intersection of the optimal solution with $N[C]$ and finish by finding the maximum independent set in a bipartite graph (a subgraph of $\overline{G} - N[C]$), which can be done in total time $2^{O(c\Delta)}$, which shows (iii). \square

Now, we show how the structure of G affects the bounds in Theorem 5.7.

COROLLARY 5.8. *Let G be a disk graph (respectively, unit ball graph) with n vertices. MAXIMUM CLIQUE can be solved in time:*

- (a) $2^{\tilde{O}(n^{2/3})}$,
 (b) $2^{\tilde{O}(\sqrt{n})}$ if the maximum degree of \overline{G} is constant,
 (c) polynomial, if both the maximum degree and the odd girth of \overline{G} are constant.

PROOF. We use the notation from Theorem 5.7. Both Δ and c can be computed in polynomial time (see, e.g., Reference [6]). Therefore, knowing what is faster among cases (i), (ii), and (iii) is tractable. For case (a), while there is a vertex of degree at least $n^{1/3}$, we branch on it. When this process stops, we do what is more advantageous between cases (ii) and (iii). Note that $\min(n/\Delta, n/c, c\Delta) \leq n^{2/3}$ (the equality is met for $\Delta = c = n^{1/3}$). For case (b), we do what is best between cases (ii) and (iii). Note that $\min(n/c, c) \leq \sqrt{n}$ (the equality is met for $c = \sqrt{n}$). Finally, case (c) follows directly from case (iii) in Theorem 5.7. \square

Observe that case (b) is typically the hardest one for MAXIMUM CLIQUE. Moreover, the win-win strategy of Corollary 5.8 can be directly applied to solve MAXIMUM WEIGHTED CLIQUE, as finding a maximum weighted independent set in a bipartite graph is still polynomial-time solvable. However, this approach cannot be easily adapted to obtain a subexponential algorithm for CLIQUE PARTITION (even CLIQUE p -PARTITION with constant p), since LIST COLORING (even LIST 3-COLORING) has no subexponential algorithm for bipartite graphs, unless the ETH fails (see Reference [45]; the bound can be obtained if we start the reduction from a sparse instance of 1-IN-3-SAT instead of PLANAR 1-IN-3-SAT).

6 EPTAS FOR MAXIMUM INDEPENDENT SET ON $\mathcal{X}(d, \beta, i)$

We start this section by showing that $\mathcal{X}(d, \beta, 1)$ has a randomized EPTAS.

THEOREM 1.3. *For any constants $d \in \mathbb{N}$, $0 < \beta \leq 1$, for every $0 < \varepsilon < 1$, there is a randomized $(1 - \varepsilon)$ -approximation algorithm running in time $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$, and a deterministic PTAS running in time $n^{\tilde{O}(1/\varepsilon^3)}$ for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, \beta, 1)$ with n vertices.*

PROOF. Let H be a graph in $\mathcal{X}(d, \beta, 1)$ with n vertices and I be a maximum independent set of H . In particular, $|I| \geq \beta n$. Since finding a maximum independent set in a bipartite graph can be done in polynomial time, we get the desired $(1 - \varepsilon)$ -approximation if we can find a subset of vertices T such that:

- T is an *odd-cycle transversal*, i.e., its removal yields a bipartite graph, and
- $|T \cap I| \leq \varepsilon |I|$.

At high level, our algorithm will thus select and remove some odd-cycle transversals T , and then apply the bipartite case algorithm. We will do this at least once for a set T that satisfies the second item, with some strong guarantee. Of course the key ingredient in finding a suitable odd-cycle transversal is the fact that $\text{iocp}(H) \leq 1$. Indeed, this implies that for any odd cycle C , the set $N[C]$ is an odd-cycle transversal.

Let $c := 8\left(\frac{1}{(\beta\varepsilon)^2} + \frac{1}{\beta\varepsilon} + 1\right) = O(1/\varepsilon^2)$, $\delta := \frac{\varepsilon}{c} = O(\varepsilon^3)$, and $s := \frac{10d}{\delta} \log \frac{1}{\delta}$. We call *short induced odd cycle* an induced odd cycle of length at most c , and *long induced odd cycle* an induced odd cycle of length more than c . First, we can assume that βn is larger than $2s$, otherwise, we can find an optimum solution by brute-force in time $2^n = 2^{\tilde{O}(1/\varepsilon^3)}$. Hence, $|I| > 2s$. \square

CLAIM 1. *There exists a subset $S \subseteq I$ of size $s = \frac{10d}{\delta} \log \frac{1}{\delta}$ such that $N(S)$ contains all vertices that have more than $\delta|I|$ neighbors in I .*

PROOF. Let A denote the set of vertices v such that $|N(v) \cap I| \geq \delta|I|$. We define the hypergraph $K := (I, \{N(v) \cap I, v \in A\})$. By assumption on H , the hypergraph K has VC-dimension at most d . By the definition of K , all its edges have size at least $\delta|V(K)|$. A celebrated result in VC-dimension theory by Haussler and Welzl [39], later improved by Blumer et al. [13], ensures that every such a hypergraph K admits a hitting set (a set of vertices that intersects every edge) of size at most $\frac{10d}{\delta} \log \frac{1}{\delta}$. \square

Algorithmically, we have two ways of selecting the set S , leading to a deterministic PTAS (see Theorem 6.3) or a randomized EPTAS. Either we run the rest of the algorithm for every subset of $V(H)$ of size $\frac{10d}{\delta} \log \frac{1}{\delta}$ inducing an independent set (which constitutes $n^{f(\epsilon)}$ possible sets), or we use another result proven in Reference [39]: Not only the hitting set exists but a uniform sample of $V(K)$ of size $\frac{10d}{\delta} \log \frac{1}{\delta}$ is a hitting set with high probability. So, we do the following $t := \lceil \frac{\log(10^{-10})}{\log(1 - (\beta/2)^s)} \rceil = 2^{\tilde{O}(1/\epsilon^3)}$ many times: We select uniformly at random a set S of size $s = \frac{10d}{\delta} \log \frac{1}{\delta}$, and continue the rest of the algorithm if S is an independent set. Since $|I| > 2s$, it holds that $\Pr(S \subseteq I) > (\beta/2)^s$. As we try out t samples, at least one satisfies $S \subseteq I$ with probability at least $1 - (1 - (\beta/2)^s)^t \geq 1 - 10^{-10}$.

We now assume that the sample S satisfies the properties of Claim 1. We start by putting in T all the vertices of $N(S)$ (note that no such vertex is in I , since I is an independent set). We define the graph $H' := H - N(S)$. We got rid of the vertices with at least $\delta|I|$ neighbors in I : In H' , there are no such vertices anymore. We want to find an odd-cycle transversal in H' that has few vertices in I .

We now run a polynomial-time algorithm (see for instance Reference [6]) that determines whether the graph is bipartite and, if not, outputs a shortest odd cycle $C_{\text{og}} = v_1 v_2 \dots v_g$ in H' .

- If H' is bipartite, then $T := N_H(S)$ is an odd-cycle transversal of H with $|T \cap I| = 0$.
- If $g = |C_{\text{og}}| \leq c$, that is, if C_{og} is a short induced odd cycle, then $|N_{H'}[C_{\text{og}}] \cap I| \leq c\delta|I| = \epsilon|I|$, and therefore $T := N_H(S) \cup N_{H'}[C_{\text{og}}]$ is an odd-cycle transversal of H with $|T \cap I| \leq \epsilon|I|$.

We can now safely assume that $g > c$, i.e., C_{og} is a long induced odd cycle. We decompose H' into the successive neighborhoods of C_{og} , which we call *layers*. We define the first layer as $L_1 := N_{H'}(C_{\text{og}})$. We define by induction the other layers as the non-empty sets $L_i := \{v \mid \text{there exists } u \in L_{i-1} \text{ with } uv \in E(H') \text{ and } v \notin L_j \text{ for } j < i\}$. Let us denote by λ the index of the last non-empty layer. Before entering into the formal details of the second part of the proof let us briefly explain its structure:

- First, we observe that if there are many layers, then there is one with index at most $\frac{2}{\beta\epsilon}$ that contains at most $\frac{\epsilon\beta}{2}n \leq \frac{\epsilon}{2}|I|$ of the vertices. We can thus delete this layer, and note that connected components that do not contain C_{og} are bipartite. We then focus on the component containing C_{og} , which has only few layers.
- Second, we show that this component admits an odd-cycle transversal of size at most $\frac{\epsilon}{2}|I|$ (informally, the neighborhood at distance up to $O(\frac{1}{\epsilon})$ of $O(\frac{1}{\epsilon})$ consecutive vertices on the cycle C_{og}).

In other words, we can find $\epsilon|I|$ vertices whose deletion yields a bipartite graph (see Figure 10), which together with $N(S)$ form the desired odd-cycle transversal.

If $\lambda > \frac{2}{\beta\epsilon}$, then there is some index $i \leq \lceil \frac{2}{\beta\epsilon} \rceil$ such that L_i is of size at most $\frac{\beta\epsilon}{2}n \leq \frac{\epsilon}{2}|I|$. We remove that layer L_i from the graph. Since $\text{iocp}(H') = 1$, the set $\bigcup_{i < j \leq \lambda} L_j$ induces a bipartite graph. Indeed, it is disjoint from the closed neighborhood of the odd cycle C_{og} . We can easily find a maximum independent set on this part of the graph and focus on the other part, which is $C_{\text{og}} \cup \bigcup_{1 \leq j < i} L_j$. We set $H'' := H'[C_{\text{og}} \cup \bigcup_{1 \leq j < i} L_j]$. If $\lambda \leq \frac{2}{\beta\epsilon}$, then we set $H'' := H'$.

So the graph H'' has at most $\frac{2}{\beta\epsilon}$ layers emanating from C_{og} . We will find an odd-cycle transversal of size at most $\frac{\epsilon}{2}|I|$. We first need some new definitions.

For $1 \leq j \leq g$, let S_j be the set of vertices $w \in V(H')$ such that there is a shortest path from w to C_{og} that ends in v_j , while no shortest path from w to C_{og} ends in v_i with $i < j$ (note that $v_j \in S_j$). We point out that the sets (S_1, \dots, S_g) induce a partition of each layer L_k . This simply follows from




Fig. 10. The layers (columns) and the strata (rows of a column). If the number of successive neighborhoods is large, a small cutset (in blue) is found among the first $\lceil \frac{2}{\beta\epsilon} \rceil$ layers. To the right of this cutset, we know that the graph is bipartite. This brings us back to the case with fewer than $\frac{2}{\beta\epsilon}$ layers, where we can find a small odd cycle transversal (in red).

the fact that for every vertex $w \in L_k$, there is a minimum index $j(w)$ such that there is a shortest path from v to C_{og} ending in $v_{j(w)}$. For each pair (k, ℓ) , we define a *stratum* as $L_k^\ell := S_\ell \cap L_k$. Note that if $L_k^\ell = \emptyset$, then for any $k' > k$, $L_{k'}^\ell = \emptyset$.

Let $z := \lceil \frac{4}{\beta\epsilon} \rceil + 2$ and for any integer γ such that $0 \leq \gamma \leq \frac{g}{z} - 1$, let $S^\gamma := \bigcup_{\gamma z + 1 \leq j \leq (\gamma + 1)z} S_j$. Informally, S^γ consists of the layers emanating from z consecutive vertices of C_{og} . Note that if $\gamma \neq \gamma'$, then S^γ and $S^{\gamma'}$ are disjoint.

CLAIM 2. For any non-negative integer $\gamma \leq \frac{2}{\beta\epsilon}$, the graph $B := H'' - S^\gamma$ is bipartite.

PROOF. Observe that $\frac{g}{z} - 1 \geq \frac{2}{\beta\epsilon}$, so each S^γ of the claim is well-defined. It holds that $C_{og} \cap S^\gamma = \bigcup_{\gamma z + 1 \leq j \leq (\gamma + 1)z} \{v_j\}$. We exhibit a proper 2-coloring of B by coloring its vertices as follows: We start by coloring each vertex of the path $C_{og} \setminus S^\gamma$ in an alternated fashion, i.e., one endpoint of the path gets color 0, its neighbor gets color 1, the next vertex gets color 0, and so on.

For each pair (k, ℓ) such that $1 \leq k < i$, $1 \leq \ell \leq g$, and $\ell \notin \{\gamma z + 1, (\gamma + 1)z\}$, we color all the vertices in the stratum $L_k^\ell = S_\ell \cap L_k$ with the opposite color of the one used for the stratum $L_{k-1}^\ell = S_\ell \cap L_{k-1}$ (with the convention that $L_0 := C_{og}$). This process colors unambiguously all the vertices of B .

Let us prove that the resulting coloring is proper. First note that the vertices of a same stratum form an independent set. Indeed, assume by contradiction that L_k^ℓ contains an edge uw . There is a shortest path P_1 from v_ℓ to u and a path P_2 from w to v_ℓ . Since u and w are in the same layer L_k , P_1 and P_2 have the same length; more precisely, $|P_1| = |P_2| = k < i \leq \frac{4}{\beta\epsilon}$. Thus, P_1, uw, P_2 defines a closed walk with $2k + 1$ edges. An odd closed walk of length $2k + 1$ implies the existence of an odd induced cycle of length at most $2k + 1$. As $2k + 1 < \frac{4}{\beta\epsilon} + 1 < g$, we reach a contradiction on the minimality of C_{og} .

There is no edge between a stratum L_k^ℓ and a stratum $L_{k'}^{\ell'}$ with $|k - k'| \geq 2$, by definition of the layers. Moreover, for $1 \leq \ell < \ell' \leq g$, there is no edge uw with $u \in L_k^\ell$ and $w \in L_{k'}^{\ell'}$ with $\min(\ell' - \ell, \ell + g - \ell') \geq \frac{4}{\beta\epsilon} + 1$, since otherwise it would be possible to construct an odd cycle strictly shorter than C_{og} . Indeed, if P_1 is a shortest path between v_ℓ and u and P_2 is a shortest path between w and $v_{\ell'}$, then P_1, uw, P_2 is a walk of length $k + k' + 1$. However, a shortest path between v_ℓ and $v_{\ell'}$ within C_{og} has length $\min(\ell' - \ell, \ell + g - \ell') \geq \frac{4}{\beta\epsilon} + 1 > k + k' + 1$. Hence, the walk P_1, uw, P_2 can be extended into an odd closed walk of length strictly smaller than g by taking the path from $v_{\ell'}$ to v_ℓ in C_{og} with the same parity as $k + k'$; a contradiction.

Therefore, if there is a monochromatic edge uw in B , it must be between L_k^ℓ and $L_{k'}^{\ell'}$ with $|k - k'| \in \{0, 1\}$, $\min(\ell' - \ell, \ell + g - \ell') < \frac{4}{\beta\varepsilon} + 1$, and $\{\ell, \ell'\} \cap [\gamma z + 1, (\gamma + 1)z] = \emptyset$. We fix k, k', ℓ, ℓ' satisfying those conditions. We call *small interval of ℓ and ℓ'* , denoted by $\text{si}(\ell, \ell')$, the integer interval $[\ell, \ell']$ if $\min(\ell' - \ell, \ell + g - \ell') = \ell' - \ell$ and $[\ell', g] \cup [1, \ell]$ if $\min(\ell' - \ell, \ell + g - \ell') = \ell + g - \ell'$. What we showed in the previous paragraph implies that if $\{\ell, \ell'\} \cap [\gamma z + 1, (\gamma + 1)z] = \emptyset$, then $\text{si}(\ell, \ell') \cap [\gamma z + 1, (\gamma + 1)z] = \emptyset$. Indeed, the small interval of ℓ and ℓ' is a circular interval over $[1, g]$ of length less than $\frac{4}{\beta\varepsilon} + 1 < z$. In particular, the vertices of C_{og} indexed by the small interval of ℓ and ℓ' are all in B .

Assume first that $k = k'$. There is a path P_1 from v_ℓ to u , and a path P_2 from w to $v_{\ell'}$, both of length k . Since by assumption the color for L_k^ℓ is the same as the color for $L_{k'}^{\ell'}$, the vertices v_ℓ and $v_{\ell'}$ have the same color (by construction of the 2-coloring). Thus, we have, in $C_{\text{og}} - S^Y$, a path P indexed by $\text{si}(\ell, \ell')$ from $v_{\ell'}$ to v_ℓ of even length less than $\frac{4}{\beta\varepsilon} + 1$. We emphasize that it is crucial that $\text{si}(\ell, \ell') \cap [\gamma z + 1, (\gamma + 1)z] = \emptyset$ (meaning that all the vertices of C_{og} indexed by $\text{si}(\ell, \ell')$ are still in B) to deduce that there is a path of even length between v_ℓ and $v_{\ell'}$. It follows from the mere fact that v_ℓ and $v_{\ell'}$ have the same color. Finally, the concatenation P_1, uw, P_2, P yields an odd cycle of length less than $2k + 1 + \frac{4}{\beta\varepsilon} + 1 \leq \frac{8}{\beta\varepsilon} + 2 < g$.

Now let us assume that $|k - k'| = 1$; say, without loss of generality, $k' = k + 1$. In that case, by construction of the 2-coloring, the edge can only be monochromatic if v_ℓ and $v_{\ell'}$ receive distinct colors. Furthermore, there is a path P_1 from v_ℓ to u , and a path P_2 from w to $v_{\ell'}$ with length of distinct parities (k and $k + 1$, respectively). Moreover, since v_ℓ and $v_{\ell'}$ get distinct colors, there is in $C_{\text{og}} - S^Y$ a path P indexed by $\text{si}(\ell, \ell')$ from $v_{\ell'}$ to v_ℓ of odd length at most $\frac{4}{\beta\varepsilon} + 1$. Again, we crucially use that all the vertices of C_{og} indexed by $\text{si}(\ell, \ell')$ are in B to deduce that P is of odd length from the fact that v_ℓ and $v_{\ell'}$ get distinct colors. Finally, the concatenation P_1, uw, P_2, P is an odd cycle of length less than $k + 1 + k + 1 + \frac{4}{\beta\varepsilon} + 1 \leq \frac{8}{\beta\varepsilon} + 3 < g$; a contradiction.

We conclude that the 2-coloring is indeed proper. \square

Since the sets of $\{S^Y\}_{Y \in [0, \lfloor \frac{g}{\beta\varepsilon} \rfloor]}$ are pairwise disjoint, a smallest set of the collection satisfies $|S^Y| \leq \frac{\beta\varepsilon}{2}n \leq \frac{\varepsilon}{2}|I|$. By Claim 2, removing this S^Y from H'' makes the graph bipartite. We finally compute a maximum independent set in polynomial time in $H'' - S^Y$. We return the best solution found. The pseudo-code is detailed in Algorithm 1.

THEOREM 1.4. *There is a randomized EPTAS for MAXIMUM CLIQUE on disk graphs, even if a geometric representation is not given. Its running time is $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$ for a $(1 - \varepsilon)$ -approximation on a graph with n vertices.*

PROOF. With a geometric representation, we can invoke the following argument to get a linear maximum stable set: Recall that the piercing number of a family of geometric objects is the minimum number of points such that each object contains at least one of those points. The piercing number of a collection of pairwise intersecting disks in the plane is 4 [24, 30, 57]. The number of faces in an arrangement of n circles (disk boundaries) is $O(n^2)$, and all the points within one face hit the same disks. In time $O(n^8)$, one can therefore exhaustively guess four points piercing a maximum clique C . We can remove all the disks that are not hit by any of those four points, since they are not part of C . This new instance G can have its vertices partitioned into four cliques, hence, $\alpha(\bar{G}) \geq |V(G)|/4$.

Without a geometric representation, we suggest the following: Disk graphs always contain a vertex whose neighborhood has independence number at most 6 (think of a vertex that has the smallest radius in one representation). For each vertex v , we run the robust PTAS of Chan and Har-Peled [27] for MAXIMUM INDEPENDENT SET in $G[N(v)]$ with ratio strictly larger than $6/7$ (say,

ALGORITHM 1: EPTAS for MAXIMUM INDEPENDENT SET on $\mathcal{X}(d, \beta, 1)$ **Precondition:** H satisfies $d := \text{VCdim}(G) = O(1)$, $\alpha(G) \geq \beta|V(G)|$, and $\text{iocp}(G) \leq 1$

```

1: function STABLE( $H, \varepsilon$ ):
2:    $c \leftarrow 8(1/(\beta\varepsilon)^2 + 1/(\beta\varepsilon) + 1)$ 
3:    $\delta \leftarrow \frac{\varepsilon}{c}$ 
4:    $s \leftarrow \frac{10d}{\delta} \log \frac{1}{\delta}$ 
5:   if  $\beta|V(H)| > 2s$  then solve  $H$  optimally by brute-force ▷  $|V(H)| = \tilde{O}(1/\varepsilon^3)$ 
6:   end if
7:   for  $_ \leftarrow 1$  to  $t = 2^{\tilde{O}(1/\varepsilon^3)}$  do
8:      $S \leftarrow$  uniform sample of  $V(G)$  of size  $s$  ▷  $S \subseteq I$  with probability  $< (\frac{\beta}{2})^s$ 
9:     if  $G[S]$  contains an edge then break and go to the next iteration
10:    end if
11:     $H' \leftarrow H - N[S]$  ▷ remove  $S$  and its neighborhood
12:     $C_{\text{og}} \leftarrow$  shortest odd cycle in  $H'$  ▷ in polynomial time [6]
13:     $g \leftarrow |C_{\text{og}}|$  ▷  $C_{\text{og}} = v_1 v_2 \dots v_g$ 
14:    if  $g \leq c$  then ▷ short induced odd cycle
15:       $S \leftarrow S \cup$  max stable on the bipartite  $H' - N[C_{\text{og}}]$  ▷  $\text{iocp}(G) = 1$ 
16:    else ▷ long induced odd cycle
17:       $L_\ell \leftarrow$  vertices of  $H'$  at distance exactly  $\ell$  from  $C_{\text{og}}$ 
18:       $L_i \leftarrow$  smallest layer among  $\{L_\ell\}_{1 \leq \ell \leq \lceil 2/\beta\varepsilon \rceil}$  ▷  $|L_i| \leq \frac{\varepsilon}{2}\alpha(G)$ 
19:       $H'' \leftarrow H'[C_{\text{og}} \cup \bigcup_{1 \leq j < i} L_j]$ 
20:       $S_k \leftarrow$  vertices of  $H''$  whose closest vertex on  $C_{\text{og}}$  of minimum index is  $v_k$ 
21:       $z \leftarrow \lceil \frac{4}{\beta\varepsilon} \rceil + 2$ 
22:       $S^Y \leftarrow$  smallest set among  $\{ \bigcup_{k \in [yz+1, (y+1)z]} S_k \}_{Y \in [0, \lfloor \frac{2}{\beta\varepsilon} \rfloor]}$  ▷  $|S^Y| \leq \frac{\varepsilon}{2}\alpha(G)$ 
23:       $S \leftarrow S \cup$  max stable on the bipartite  $H'[\bigcup_{j > i} L_j]$  ▷  $\text{iocp}(G) = 1$ 
24:       $S \leftarrow S \cup$  max stable on the bipartite  $H'' - S^Y$  ▷ Claim 2
25:    end if
26:  end for
27:  return  $S$  at the iteration maximizing its cardinality
28: end function

```

Postcondition: output S is a stable set of size at least $(1 - \varepsilon)\alpha(G)$ with high probability

7/8). By *robust*, we mean that their local-search based algorithm does not require a geometric representation. By the previous remark, at least one run has to report a value of at most 6 (indeed, $7(7/8) > 6$). Let u be a vertex corresponding to such a run. For every disk graph G , $\chi(G) \leq 6\omega(G)$ (actually a better bound of $6\omega(G) - 6$ is known [48]). Then $\alpha(G)\omega(G) \geq \frac{\chi(G)\alpha(G)}{6} \geq \frac{|V(G)|}{6}$, hence, $\omega(G) \geq \frac{|V(G)|}{6\alpha(G)}$.

We branch on two outcomes. Either u is in a maximum clique: We run the approximation of Theorem 1.3 on $G_u := \overline{G[N[u]]}$, which satisfies $\omega(\overline{G_u}) = \alpha(G_u) \geq |V(G_u)|/36$ (recall that we chose u so $\alpha(\overline{G_u}) \leq 6$). Or this vertex is not in any maximum clique: We delete it from the graph. Our branching tree has size $2n + 1$, so it only costs an extra linear multiplicative factor.

The VC-dimension of the neighborhoods of disk graphs, and even pseudo-disk graphs [8], is at most 4. Since the VC-dimension of a graph is equal to the one of its complement, the VC-dimension of \overline{G} is also at most 4. Finally, by Theorem 1.1, $\text{iocp}(\overline{G}) \leq 1$. We only call the approximation algorithm (a polynomial number of times) with disk graphs G' such that $\overline{G'} \in \mathcal{X}(4, \frac{1}{36}, 1)$ (argument without the geometric representation) or $\overline{G'} \in \mathcal{X}(4, \frac{1}{4}, 1)$ (argument with the geometric representation). Hence, we conclude by Theorem 1.3. \square

It is folklore that unit ball graphs have geometric VC-dimension 4. One can observe that, in the case of *unit* ball graphs, the geometric VC-dimension coincides with the VC-dimension of the neighborhoods. At the price of a multiplicative factor n in the running time, one can guess a vertex v in a maximum clique of a unit ball graph G and look for a clique in its neighborhood $H := G[N(v)]$. As the kissing number for unit spheres is bounded (it is 12), one can also show that the neighborhood of this vertex (in fact, of any vertex) can be partitioned into a constant number of cliques. The exact number is irrelevant here: An easy volume-based argument can show that this number is no greater than 30. Thus, $\alpha(\overline{H}) \geq |V(H)|/30$. Therefore, from Theorem 1.3 and Theorem 1.2, we immediately obtain the following:

THEOREM 1.5. *There is a randomized EPTAS in time $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$ for MAXIMUM CLIQUE on unit ball graphs, even if a geometric representation is not given.*

We can extend the EPTAS to work for constant (not necessarily 1) induced odd cycle packing number.

THEOREM 6.1. *For any constants $d, i \in \mathbb{N}, 0 < \beta \leq 1$, for every $\varepsilon > 0$, there is a randomized $(1 - \varepsilon)$ -approximation algorithm running in time $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$ for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, \beta, i)$ with n vertices.*

PROOF. Let I be a maximum independent set. We show by induction on i that for any ε' , we can find in time $2^{\tilde{O}(1/\varepsilon'^3)} n^{O(1)}$ a stable set of size $(1 - i\varepsilon')|I|$. The base case is Theorem 1.3. We assume that there is such an algorithm when the induced odd cycle packing number is $i - 1$. We follow Algorithm 1 with a graph H such that $\text{iocp}(H) = i$. On lines 15 and 23, the graph is not necessarily bipartite anymore (on line 24, the resulting graph is still bipartite in this case). Although, the induced odd cycle packing number is decreased to $i - 1$. So, by the induction hypothesis, we get a stable set within a factor $(1 - (i - 1)\varepsilon')$ of the optimum. To get there, we removed a subset of vertices of size at most $\varepsilon'|I|$. Therefore, the solution S that we obtain satisfies $|S| \geq (1 - i\varepsilon')|I|$.

We obtain the theorem by setting $\varepsilon := i\varepsilon'$, since i is absorbed in the \tilde{O} in the running time. \square

To the detriment of the efficiency of the approximation scheme, when $\text{iocp}(G) \leq 1$ (or even $\text{iocp}(G) = O(1)$), we can discard one of the two other assumptions of Theorem 1.3. Namely, we do not need bounded VC-dimension or that the optimum solution is a positive fraction of the number of vertices.

THEOREM 6.2. *There is a randomized PTAS for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(\infty, \beta, 1)$.*

PROOF. We observe that $\text{VCdim}(G)$ is always at most $\log |V(G)| = \log n$. So, following Algorithm 1, we now sample a set S of size $\frac{10 \log n}{\delta} \log \frac{1}{\delta}$. The probability for S to be successfully contained in an optimum solution I is at least $\frac{1}{n^{O(1/\varepsilon^3)}}$. Thus, by repeating this experience $n^{\tilde{O}(1/\varepsilon^3)}$ times, $S \subseteq I$ holds in at least one branch, with high probability. The rest of the algorithm is unchanged. \square

THEOREM 6.3. *There is a deterministic PTAS for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, 0, 1)$.*

PROOF. Now, instead of sampling $s = \frac{10d}{\delta} \log \frac{1}{\delta}$ vertices, we try all the $n^s = n^{f(\varepsilon)}$ subsets of size s . One of them falls entirely in an optimum solution I and is a desired ε -net (note that if the largest independent set contains less than s vertices, then we can find it in time $O(n^s)$). We also change lines 18 and 22 of Algorithm 1: Instead of deleting the lightest layer among the first $O(\frac{1}{\varepsilon})$, and the lightest block of strata among $O(\frac{1}{\varepsilon})$ disjoint blocks, we try all possible such pairs. This only adds

a multiplicative factor in $O(1/\varepsilon^2)$ to the running time. One of the deleted pair contains less than $\varepsilon|I|$ vertices of I , and we conclude similarly. \square

Both results generalize to $\mathcal{X}(\infty, \beta, i)$ and $\mathcal{X}(d, 0, i)$, for every integer i , with the same arguments as in the proof of Theorem 6.1.

7 OTHER INTERSECTION GRAPHS

In this section, we discuss the impossibility of generalizing our algorithms to higher dimensions and related classes of intersection graphs. We first show some hardness of approximation for MAXIMUM INDEPENDENT SET on the class of all the 2-subdivisions, hence, the same lower bound for MAXIMUM CLIQUE on all the co-2-subdivisions.

It is folklore that from the PCP of Moshkovitz and Raz [52], which roughly implies that MAX 3-SAT cannot be $(7/8 + \varepsilon)$ -approximated in subexponential time under the ETH, one can derive such inapproximability in subexponential time for many hard graph and hypergraph problems; see for instance Reference [16]. The following inapproximability result for MAXIMUM INDEPENDENT SET on bounded-degree graphs was shown by Chlebík and Chlebíková [28]. As their reduction is almost linear, the PCP of Moshkovitz and Raz boosts this hardness result from ruling out polynomial-time up to ruling out subexponential time 2^{n^γ} for any $\gamma < 1$.

THEOREM 7.1 ([28, 52]). *For any $\Delta \geq 3$, and $\gamma < 1$, there is a constant $\eta < 1$ such that MAXIMUM INDEPENDENT SET on graphs with n vertices and maximum degree Δ cannot be η -approximated in time 2^{n^γ} , unless the ETH fails.*

We could actually state a slightly stronger statement for the running time but will settle for this for the sake of clarity.

THEOREM 7.2. *For any $\gamma < 1$, there is a constant $\zeta < 1$ such that MAXIMUM INDEPENDENT SET on the class of all the 2-subdivisions has no ζ -approximation algorithm running in subexponential time 2^{n^γ} , unless the ETH fails.*

PROOF. Let G be a graph with maximum degree a constant $\Delta \geq 3$, with n vertices v_1, \dots, v_n and m edges e_1, \dots, e_m , and let H be its 2-subdivision. Recall that to form H , we subdivided every edge of G exactly twice. These $2m$ vertices in $V(H) \setminus V(G)$, representing edges, are called *edge vertices* and are denoted by $v^+(e_1), v^-(e_1), \dots, v^+(e_m), v^-(e_m)$, as opposed to the other vertices of H , which we call *original vertices*. If $e_k = v_i v_j$ is an edge of G , then $v^+(e_k)$ (respectively, $v^-(e_k)$) has two neighbors: $v^-(e_k)$ and v_i (respectively, $v^+(e_k)$ and v_j).

Observe that there is a maximum independent set S that contains exactly one of $v^+(e_k), v^-(e_k)$ for every $k \in [m]$. Indeed, S cannot contain both $v^+(e_k)$ and $v^-(e_k)$, since they are adjacent. However, if S contains neither $v^+(e_k)$ nor $v^-(e_k)$, then adding $v^+(e_k)$ to S and potentially removing the other neighbor of $v^+(e_k)$ that is v_i (with $e_k = v_i v_j$) can only increase the size of the independent set. Hence, S contains m edge vertices and $s \leq n$ original vertices, and there is no larger independent set in H .

We observe that the s original vertices in S form an independent set in G . Indeed, if $v_i v_j = e_k \in E(G)$ and $v_i, v_j \in S$, then neither $v^+(e_k)$ nor $v^-(e_k)$ could be in S .

Now, assume there is an approximation with ratio $\zeta := (1 + \frac{2(1-\eta)}{\eta(\Delta+1)^2})^{-1}$ for MAXIMUM INDEPENDENT SET on 2-subdivisions running in subexponential time, where $\eta < 1$ is a ratio that is not attainable for MAXIMUM INDEPENDENT SET on graphs of maximum degree Δ according to Theorem 7.1. On instance H , this algorithm would output a solution with m' edge vertices and s' original vertices. As we already observed, this solution can be easily (in polynomial time) transformed into an at-least-as-good solution with m edge vertices and s'' original vertices forming

an independent set in G . Further, we may assume that $s'' \geq n/(\Delta + 1)$, since for any independent set of G , we can obtain an independent set of H consisting of the same set of original vertices and m edge vertices. Since $m \leq n\Delta/2$ and $s'' \geq n/(\Delta + 1)$, we obtain $m \leq s''\Delta(\Delta + 1)/2$ and $2m/(\Delta + 1)^2 \leq s''\Delta/(\Delta + 1)$. From $\frac{m+s''}{m+s} \geq \zeta$ and $\Delta \geq 3$, we have

$$s \leq m \cdot \frac{2(1-\eta)}{\eta(\Delta+1)^2} + s'' \cdot \left(1 + \frac{2(1-\eta)}{\eta(\Delta+1)^2}\right) \leq s'' \left(\frac{(1-\eta)\Delta}{\eta(\Delta+1)} + 1 + \frac{2(1-\eta)}{\eta(\Delta+1)^2}\right) \leq s'' \left(1 + \frac{1-\eta}{\eta}\right) = \frac{s''}{\eta}.$$

This contradicts the inapproximability of Theorem 7.1. Indeed, note that the number of vertices of H is only a constant times the number of vertices of G (recall that G has bounded maximum degree, hence, $m = O(n)$). \square

We get the following as a direct corollary:

COROLLARY 7.3. *For every $\gamma < 1$, there is a constant $\zeta < 1$, MAXIMUM CLIQUE on the class of all the co-2-subdivisions has no ζ -approximation algorithm running in subexponential time 2^{n^γ} , unless the ETH fails.*

For exact algorithms, the subexponential time that we rule out under the ETH is not only 2^{n^γ} (with $\gamma < 1$) but actually any $2^{o(n)}$.

7.1 Balls and Higher Dimensions

In sharp contrast to the algorithms for disk and unit ball graphs, we will prove that with an extra dimension or with different radii (even arbitrarily close to each other), a PTAS is highly unlikely.

By Corollary 7.3, we just need to show that all co-2-subdivisions can be realized by our geometric objects. This appears like a simple and powerful method to rule out a PTAS (QPTAS, and even SUBEXPAS) for a geometric clique problem.

The co-2-subdivision of a graph with n vertices and m edges can be thought of as follows: It is made of a clique on n vertices, representing the initial vertices, and a clique on $2m$ vertices minus a perfect matching, representing endpoints of the initial edges. Each anti-matched pair of vertices corresponds to an edge in the initial graph. Each vertex representing one endpoint of an initial edge is adjacent to all the vertices representing the initial vertices but this endpoint.

THEOREM 7.4. *The class of 4-dimensional unit ball graphs contains all the co-2-subdivisions.*

PROOF. Given any simple graph $G = (V, E)$ with n vertices and m edges, we want to build a set S of $n + 2m$ points in \mathbb{R}^4 where each vertex v is represented by a point $p(v)$ and each edge e , by two points $p^+(e)$ and $p^-(e)$. A pair of points in S should be at distance at most 2, except if it is the two points of the same edge e , or if it is a point $p(v)$ with either a $p^+(vw)$ or a $p^-(wv)$; those pairs should be at distance strictly more than 2. We denote by x, y, z, t the coordinates of \mathbb{R}^4 . Let \mathcal{P} be the plane defined by the intersection of the hyperplanes of equation $z = 0$ and $t = 0$. The projection π onto \mathcal{P} of the points $p^+(e)$ (respectively, $p^-(e)$) fall regularly on the “top” part (respectively, “bottom” part) of a circle C of \mathcal{P} centered at the origin $O = (0, 0, 0, 0)$ and of diameter 2, such that for each edge e , $\pi(p^+(e))$ and $\pi(p^-(e))$ are antipodal on C . We just defined the points $\pi(p^+(e))$ and $\pi(p^-(e))$. The actual points $p^+(e)$ and $p^-(e)$ will be fixed later by moving them very slightly away from their projection in a two-dimensional plane orthogonal to \mathcal{P} . Let η be the maximum distance between a pair $(\pi(p^+(e)), \pi(p^-(e')))$ with $e \neq e'$. By construction $\eta < 2$.

Let \mathcal{P}^\perp be the (2-dimensional) plane containing the center O of C and orthogonal to \mathcal{P} ; in other words, the intersection of the hyperplanes of equation $x = 0$ and $y = 0$. We observe that all the points of \mathcal{P}^\perp are equidistant to all the points $\pi(p^+(e))$ and $\pi(p^-(e))$. We place all the points $p(v)$ in \mathcal{P}^\perp regularly spaced on a arc of a circle lying on \mathcal{P}^\perp centered at O and of radius $\sqrt{3} - \varepsilon$. One can




Fig. 11. The overall construction for 4-dimensional unit balls. We only represent the centers.

notice that for any $(v, e, s) \in V \times E \times \{+, -\}$, $d(p(v), \pi(p^s(e))) = \sqrt{4 - 2\sqrt{3}\epsilon + \epsilon^2}$. We will choose $\epsilon \ll 2 - \eta \ll 1$ so this shared distance is just below 2, and the points $\{p^+(e), p^-(e)\}_{e \in E}$ realize the same adjacencies than their projection by π .

For every $e = uv \in E$, we place $p^+(e)$ such that $\overrightarrow{\pi(p^+(e))p^+(e)} = -\frac{(\epsilon + \epsilon')}{\|Op(u)\|} \overrightarrow{Op(u)}$ and $p^-(e)$ such that $\overrightarrow{\pi(p^-(e))p^-(e)} = -\frac{(\epsilon + \epsilon')}{\|Op(v)\|} \overrightarrow{Op(v)}$. In words, we push very slightly $p^+(e)$ (respectively, $p^-(e)$) away from $\pi(p^+(e))$ (respectively, $\pi(p^-(e))$) in the opposite direction of $\overrightarrow{Op(u)}$ (respectively, $\overrightarrow{Op(v)}$). This way, $p^+(e)$ is at distance more than 2 from $p(u)$. Indeed,

$$\begin{aligned} d(p^+(e), p(u)) &= \|\overrightarrow{p^+(e)\pi(p^+(e))} + \overrightarrow{\pi(p^+(e))O} + \overrightarrow{Op(u)}\| = \|\overrightarrow{\pi(p^+(e))O} + \overrightarrow{Op(u)} + \overrightarrow{p^+(e)\pi(p^+(e))}\| \\ &= \|\overrightarrow{\pi(p^+(e))O} + (\sqrt{3} - \epsilon + \epsilon + \epsilon') \frac{\overrightarrow{Op(u)}}{\|Op(u)\|}\| = \sqrt{1 + (\sqrt{3} + \epsilon')^2} > 2. \end{aligned}$$

Similarly, $d(p^-(e), p(v)) > 2$. See Figure 11 for an illustration of the construction. We choose ϵ' infinitesimal (in particular, $\epsilon' \ll \epsilon$) so, still, $d(p^+(e), p(w)) < 2$ for any $w \neq u$ and $d(p^-(e), p(w)) < 2$ for any $w \neq v$. \square

COROLLARY 7.5. *For every $\gamma < 1$, there is a constant $\zeta < 1$ such that MAXIMUM CLIQUE on 4-dimensional unit ball graphs is not ζ -approximable even in time 2^{n^γ} , unless the ETH fails. Moreover, MAXIMUM CLIQUE is NP-hard on 4-dimensional unit ball graphs.*

The proof of Theorem 7.4 can be tweaked for 3-dimensional balls of different radii. For a real $\epsilon > 0$, we say that the radii of a representation (or the representation itself) are ϵ -close if the radii are all contained in the interval $[1, 1 + \epsilon]$. We denote by $\mathcal{B}(1, 1 + \epsilon)$ the ball graphs having an ϵ -close representation.

THEOREM 7.6. *For any $\epsilon > 0$, the subclass of ball graphs $\mathcal{B}(1, 1 + \epsilon)$ contains all the co-2-subdivisions.*

PROOF. Let x, y, z be the coordinates of \mathbb{R}^3 . We start similarly and define the same $\pi(p^+(e))$ and $\pi(p^-(e))$ as in the previous construction for the points $p^+(e)$ and $p^-(e)$ on a circle C of diameter 2 centered at $O = (0, 0, 0)$ on a plane \mathcal{P} of equation $z = 0$. One difference is that $\pi(p^+(e))$ and $\pi(p^-(e))$ are no longer projections. We then place the points $p(v)$ regularly spaced along the z -axis. More precisely, if the vertices are numbered v_1, v_2, \dots, v_n , the position of $p(v_i)$ is $(0, 0, \sqrt{3} + i\epsilon')$. The radius of the disk representing v_i centered at $p(v_i)$ is set to $r_i := \sqrt{1 + (\sqrt{3} + i\epsilon')^2} - 1 + \epsilon''$.




Fig. 12. The overall construction for 3-dimensional balls. Blue centers indicate unit balls, while the red the center, the larger the radius of the ball. We only represented the two centers for the edge $e_6 = v_5v_8$, showing how $p^+(e_6)$ is pushed away from $p(v_5)$, and $p^-(e_6)$, from $p(v_8)$. One may observe that $p(v_1), \dots, p(v_n), O, \pi(p^+(e_6)), \pi(p^-(e_6)), p^+(e_6), p^-(e_6)$ are coplanar.

The radii associated to the centers $p^+(e)$ and $p^-(e)$ are all set to 1. We move $p^+(v_iv_j)$ (respectively, $p^-(v_iv_j)$) away from $\pi(p^+(v_iv_j))$ (respectively, $\pi(p^-(v_iv_j))$) in the opposite direction of $\overrightarrow{p^+(v_iv_j)p(v_i)}$ (respectively, $\overrightarrow{p^-(v_iv_j)p(v_j)}$) by an infinitesimal quantity to only suppress the overlap of the disks centered at $p^+(v_iv_j)$ and $p(v_i)$ (respectively, $p^-(v_iv_j)$ and $p(v_j)$). See Figure 12 for an illustration of the construction. We make ϵ' and ϵ'' small enough that all the values r_i are between 1 and $1 + \epsilon$. \square

We call *quasi unit ball graphs* those graphs in the intersection $\bigcap_{\epsilon > 0} \mathcal{B}(1, 1 + \epsilon)$. As a corollary, we get some strong inapproximability even for quasi unit ball graphs.

COROLLARY 7.7. *For every $\gamma < 1$, there is a constant $\zeta < 1$ such that MAXIMUM CLIQUE on quasi unit ball graphs is not ζ -approximable even in time 2^{n^γ} , unless the ETH fails. Moreover, MAXIMUM CLIQUE is NP-hard on quasi unit ball graphs.*

We observe that the lower bounds of Corollaries 7.5 and 7.7 also hold when the geometric representation is given in input. For that, we need to argue that we can compute the coordinates (and radii) of our constructions in polynomial time. Let us estimate η , the maximum distance between a pair $(\pi(p^+(e)), \pi(p^-(e')))$ with $e \neq e'$, for an m -edge graph with m sufficiently large. We place the $2m$ points $\pi(p^+(e_1)), \dots, \pi(p^+(e_m)), \pi(p^-(e_1)), \dots, \pi(p^-(e_m))$ regularly spaced every π/m radians. Thus,

$$\eta = 2\sqrt{1 - \sin^2\left(\frac{\pi}{2m}\right)} = 2\cos\left(\frac{\pi}{2m}\right) \leq 2\left(1 - \frac{\pi^2}{8m^2}\right).$$

Hence, $2 - \eta \leq \frac{\pi^2}{4m^2}$. This implies that in the proof of Theorem 7.4, we can choose $\epsilon := \frac{1}{100m^3}$ and $\epsilon' := \frac{1}{100m^4}$. Now all the centers can be snapped to grid points on a three-dimensional grid with precision $\frac{1}{100m^5}$ (i.e., two adjacent grid points are at distance $\frac{1}{100m^5}$). In the the proof of Theorem 7.6, we can choose $\epsilon' := \frac{\epsilon}{100m^3}$ and $\epsilon'' := \frac{\epsilon}{100m^4}$. And all the centers and radii can be defined up to precision $\frac{\epsilon}{100m^5}$.

Note that we realize both constructions with polynomial precision (in m and, for Theorem 7.6, in $1/\epsilon$), although even a single-exponential precision would be allowed in a polytime construction. Indeed, in binary, polynomially many bits yield exponentially precise coordinates.




Fig. 13. A graph (left), which has a representation with empty ellipses (right) but no representation with convex sets.

7.2 Filled Ellipses and Filled Triangles

A natural generalization of a disk is an *elliptical disk*, also called a *filled ellipse*, i.e., an ellipse plus its interior. Arguably the simplest convex set with non empty interior is a filled triangle (a triangle plus its interior).

APX-hardness was shown for MAXIMUM CLIQUE in the intersection graphs of (*non-filled*) ellipses and triangles by Ambühl and Wagner [7]. Their reduction also implies that there is no subexponential algorithm for this problem, unless the ETH fails. Moreover, they claim that their hardness result extends to filled ellipses, since “*intersection graphs of ellipses without interior are also intersection graphs of filled ellipses.*” Unfortunately, this claim is incorrect, as we prove below: We construct a graph, which can be represented by ellipses without their interior, but cannot be represented by any convex sets (and thus by filled ellipses).

THEOREM 7.8. *There is a graph G that has an intersection representation with ellipses without their interior, but has no intersection representation with convex sets.*

PROOF. The argument is similar to the one used by Brimkov et al. [20], which was in turn inspired by the construction by Kratochvíl and Matoušek [47]. Consider the graph G in Figure 13, containing what we will henceforth refer to as *black, gray, red, blue, and white* vertices. Gray, blue, and red vertices are called *connector vertices*. We aim to show that G cannot be represented by a family of convex sets.

For the sake of contradiction, suppose that G admits a representation by convex sets. We denote by R_v the convex set representing the vertex v . The high-level idea is as follows: First, we observe that the union of representatives of white vertices separates the rest of the plane into two disjoint regions. Furthermore, since the subgraph induced by black vertices is connected, their representatives must all lie in one of these two regions. Next, we look at the set $R_a \cup R_b \cup R_c$, for the black vertices a, b , and c , and observe that due to the connector vertices, certain parts of R_a, R_b, R_c need to appear in a prescribed order as we move along the boundary of $R_a \cup R_b \cup R_c$. This forces the representation of $G - \{d\}$ to have a somewhat rigid structure. We carefully analyze this structure and conclude that the only way to realize all the adjacencies of d is to place R_d where it partially overlaps R_c , which contradicts the non-edge between c and d .




Fig. 14. A hypothetical representation of G with intersecting convex sets.

We refer to Figure 14 for a picture of the construction. The union of the representatives of white vertices contains a (closed) Jordan curve that we will call the outer circle. Let us choose the outer circle in such a way that it intersects the representatives of all connector vertices and such that its intersection with each white or connector vertex representative is a connected set. The outer circle divides the plane into two faces—an interior and an exterior.

Since no black vertex is adjacent to a white vertex, the outer circle cannot be crossed by the representative of any black vertex. Moreover, as black vertices form a connected subgraph, they have to be represented in the same face F (with respect to the outer circle). Since the circle intersects vertex representatives only in connected sets, as we traverse it, it enters and leaves the representatives exactly once. We order the representatives of white vertices in such a way that when traversing the circle in a clockwise fashion the representatives of connector vertices are visited in the following order: $a_1, \bar{c}_1, c_1, b_1, c_2, \bar{c}_2, a_2, \bar{c}_3, c_3, b_2, c_4, \bar{c}_4$, where each z_i for $z \in \{a, b\}$ is a distinct gray neighbor of z , each c_i is a distinct blue neighbor of c , and each \bar{c}_i is a distinct red neighbor of c .

Clearly, each gray neighbor of a must intersect R_a outside $R_a \cap (R_b \cup R_c)$, each gray neighbor of b must intersect R_b outside $R_b \cap (R_a \cup R_c)$, and each blue or red neighbor of c must intersect R_c outside $R_c \cap (R_a \cup R_b)$. Thus, some parts of R_a, R_b , and R_c are exposed (i.e., outside the intersection with the union of the representatives of the remaining two vertices) in the order: a, c, b, c, a, c, b, c , as we move along the boundary of $R_a \cup R_b \cup R_c$.

Let z' be a connector vertex and let $z \in \{a, b, c\}$ be its neighbor. Notice that for each z' , the vertex z is uniquely defined. The set $R_{z'}$ contains a segment $s(z')$, whose one end is on the boundary of R_z and the other end is on the outer circle (recall that all representatives are convex).

Observe that the segments $s(a_1), s(a_2), s(b_1), s(b_2)$ are pairwise disjoint and they separate $F \setminus (R_a \cup R_b)$ into four regions Q_1, Q_2, Q_3, Q_4 , such that for $i \in [4]$ it holds that $R_{c_i} \cap F \subseteq Q_i$. Note that one of these regions may be unbounded if F is the unbounded face of the outer circle. For each $i \in [4]$, since c_i is non-adjacent to \bar{c}_i , we observe that $(R_{c_i} \cap F) \setminus R_c$ is contained in the subregion of Q_i bounded by $s(\bar{c}_i) \cup R_c$.

For $i = \{1, 2, 3, 4\}$, let p_i be a point in $R_d \cap R_{c_i}$. Such a point exists, since d is adjacent to c_i . By convexity of R_d , the segment p_1p_2 is contained in R_d . Since d is non-adjacent to c, \bar{c}_1 , and \bar{c}_2 , we observe that p_1p_2 must cross the $s(b_1) \cup R_b$. As d is non-adjacent to b_1 , the segment p_1p_2 intersects R_b . Let q_1 be an intersection point of p_1p_2 and R_b . In an analogous way, we define q_2 to be an intersection point of p_3p_4 and R_b .

Let us now consider the segment q_1q_2 . By convexity of R_d and R_b , we have $q_1q_2 \subseteq R_b \cap R_d$. The segment q_1q_2 must intersect $s(c_1) \cup R_c \cup s(c_2)$, and let q' be this intersection point. If $q' \in s(c_1) \cup s(c_2)$, then we get the contradiction from the fact that b is non-adjacent to c_1 and c_2 . However, if $q' \in R_c$, then we get the contradiction from the fact that d and c are non-adjacent.

Finally, it is easy to represent G with empty ellipses (see Figure 13, right). \square

This error and the confusion between filled ellipses and ellipses without their interior has propagated to other more recent papers [43]. Fortunately, we show that the hardness result does hold for filled ellipses (and filled triangles) with a different reduction. Our construction can be seen as streamlining the ideas of Ambühl and Wagner [7]. It is simpler and, in the case of (filled) ellipses, yields a somewhat stronger statement.

THEOREM 7.9. *For every $\gamma < 1$, there is a constant $\zeta < 1$ such that MAXIMUM CLIQUE on the intersection graphs of filled ellipses has no ζ -approximation algorithm running in subexponential time 2^{n^γ} , unless the ETH fails, even when the ellipses have arbitrarily small eccentricity and the different lengths of the major axis are arbitrarily close.*

This is in sharp contrast with our subexponential algorithm and PTAS when the eccentricity is 0 (case of disks). For any $\varepsilon > 0$, if the eccentricity is only allowed to be at most ε , then a SUBEXPAS is very unlikely. This result subsumes Reference [25] (where NP-hardness is shown for connected shapes contained in a disk of radius 1 and containing a concentric disk of radius $1 - \varepsilon$ for arbitrarily small $\varepsilon > 0$) and corrects Reference [7]. We show the same hardness for the intersection graphs of filled triangles.

THEOREM 7.10. *For every $\gamma < 1$, there is a constant $\zeta < 1$ such that MAXIMUM CLIQUE on the intersection graphs of filled triangles has no ζ -approximation algorithm running in subexponential time 2^{n^γ} , unless the ETH fails.*

Once again, to show Theorem 7.9 and Theorem 7.10, it is sufficient to show that intersection graphs of (filled) ellipses or of (filled) triangles contain all co-2-subdivisions. We start with (filled) triangles, since the construction is simpler.

LEMMA 7.11. *The class of intersection graphs of filled triangles contains all co-2-subdivisions.*

PROOF. Let G be any graph with n vertices v_1, \dots, v_n and m edges e_1, \dots, e_m , and H be its co-2-subdivision. We start with $n + 2$ points $p_0, p_1, p_2, \dots, p_n, p_{n+1}$ forming a convex monotone chain. Those points can be chosen as $p_i := (i, p(i))$ where p is the equation of a positive parabola taking its minimum at $(0, 0)$. For each $i \in [0, n + 1]$, let q_i be the reflection of p_i by the line of equation $y = 0$. Let $x := (n + 1, 0)$. For each vertex $v_i \in V(G)$ the filled triangle $\delta_i := p_iq_ix$ encodes v_i . Observe that the points $p_0 = q_0, p_{n+1}$, and q_{n+1} will only be used to define the filled triangles encoding edges.

To encode (the two new vertices of) a subdivided edge $e_k = v_iv_j$, we use two filled triangles Δ_k^+ and Δ_k^- . The triangle Δ_k^+ (respectively, Δ_k^-) has an edge that is supported by $\ell(p_{i-1}, p_{i+1})$




Fig. 15. A co-2-subdivision of a graph with 5 vertices (in red) represented with triangles. Only two edges are shown: one between vertices 1 and 4 (green) and one between vertices 2 and 3 (blue).

(respectively, $\ell(q_{j-1}, q_{j+1})$) and is prolonged so it crosses the boundary of each $\delta_{i'}$ but δ_i (respectively, but δ_j). A second edge of Δ_k^+ and Δ_k^- are parallel and make with the horizontal a small angle εk , where $\varepsilon > 0$ is chosen so εm is smaller than the angle formed by $\ell(p_0, p_1)$ with the horizontal line. Those almost horizontal edges intersect for each pair $\Delta_{k'}^+$ and $\Delta_{k''}^-$, with $k' \neq k''$ intersects close to the same point. Filled triangles Δ_k^+ and Δ_k^- do not intersect. See Figure 15 for the complete picture.

It is easy to check that the intersection graph of $\{\delta_i\}_{i \in [n]} \cup \{\Delta_k^+, \Delta_k^-\}_{k \in [m]}$ is H . The family $\{\delta_i\}_{i \in [n]}$ forms a clique, since they all contain for instance the point x . The filled triangle Δ_k^+ (respectively, Δ_k^-) intersects every other filled triangle except Δ_k^- (respectively, Δ_k^+) and δ_i (respectively, δ_j) with $e_k = v_i v_j$.

One may observe that no triangle is fully included in another triangle. So the construction works both as the intersection graph of filled triangles *and* triangles without their interior. The edges of Δ_k^+ and Δ_k^- crossing the boundary of all but one δ_i can be arbitrary prolonged to the right. The almost horizontal edges of these triangles can be arbitrary prolonged to the left. Thus, the triangles can all be made isosceles. \square

We use the same ideas for the construction with filled ellipses. Two tangents of the ellipse will play the role of the two important sides of the triangles encoding edges of the initial graph G .

LEMMA 7.12. *The class of intersection graphs of filled ellipses contains all co-2-subdivisions.*

PROOF. Let G be any graph with n vertices v_1, \dots, v_n and m edges e_1, \dots, e_m , and H be its co-2-subdivision. We start with the convex monotone chain $p_0, p_1, p_2, \dots, p_{n-1}, p_n, p_{n+1}$; only the gap between p_i and p_{i+1} is chosen very small compared to the positive y -coordinate of p_0 . The disks \mathcal{D}_i encoding the vertices $v_i \in G$ must form a clique. We also take p_0 with a large x -coordinate. For $i \in [0, n + 1]$, q_i is the symmetric of p_i with respect to the x -axis. For each $i \in [n]$, we define \mathcal{D}_i as the disk whose boundary is the unique circle that goes through p_i and q_i , and whose tangent at p_i has the direction of $\ell(p_{i-1}, p_{i+1})$. It can be observed that, by symmetry, the tangent of \mathcal{D}_i at q_i has the direction of $\ell(q_{i-1}, q_{i+1})$.

Let us call τ_i^+ (respectively, τ_i^-) the tangent of \mathcal{D}_i at p_i (respectively, at q_i) very slightly translated upward (respectively, downward). The tangent τ_i^+ (respectively, τ_i^-) intersects every disk $\mathcal{D}_{i'}$ but




Fig. 16. The blue line intersects every red disk but the third one.

\mathcal{D}_i (see Figure 16). Let denote by p'_i (respectively, q'_i) the projection of p_i (respectively, q_i) onto τ_i^+ (respectively, onto τ_i^-). For each $k \in [m]$, let ℓ_k be the line crossing the origin $O = (0, 0)$ and forming with the horizontal an angle εk , where εk is smaller than the angle formed by $\ell(p_0, p_1)$ with the horizontal. Let ℓ_k^+ (respectively, ℓ_k^-) be ℓ_k very slightly translated upward (respectively, downward). To encode an edge $e_k = v_i v_j$, we have two filled ellipses \mathcal{E}_k^+ and \mathcal{E}_k^- . The ellipse \mathcal{E}_k^+ (respectively, \mathcal{E}_k^-) is defined as being tangent with τ_i^+ at p'_i (respectively, with τ_j^- at q'_j) and tangent at ℓ_k^+ (respectively, ℓ_k^-) at the point of x -coordinate 0 (thus, very close to O), where $e_k = v_i v_j$. The proof that the intersection graph of $\{\mathcal{D}_i\}_{i \in [n]} \cup \{\mathcal{E}_k^+, \mathcal{E}_k^-\}_{k \in [m]}$ is H is similar to the case of filled triangles.

As no ellipse is fully contained in another ellipse, this construction works for both filled ellipses and ellipses without their interior.

We place p_0 at $P := (\sqrt{3}/2, 1/2)$ and make the distance between p_i and p_{i+1} very small compared to 1. All points p_i are very close to P and all points q_i are very close to $Q := (\sqrt{3}/2, -1/2)$. This makes the radius of all disks \mathcal{D}_i arbitrarily close to 1. We choose the convex monotone chain p_0, \dots, p_{n+1} so $\ell(p_0, p_1)$ forms a 30-degree angle with the horizontal. As the chain is strictly convex but very close to a straight-line, $\ell(p_0, p_1) \approx \ell(p_n, p_{n+1}) \approx \ell(p_i, p_{i+1}) \approx \ell(p_i, p_{i+2})$. Thus, all those lines almost cross P and form an angle of roughly 30 degrees with the horizontal. The same holds for points q_i . For the choice of an elliptical disk tangent to the x -axis at O and to a line with a 60-degree slope at P (respectively, at Q), we take a disk of radius 1 centered at $(0, 1)$ (respectively, at $(0, -1)$); see Figure 17.

The acute angle formed by ℓ_1 and ℓ_m (incident in O) is made arbitrarily small so, by continuity of the elliptical disk defined by two tangents at two points, the filled ellipses \mathcal{E}_k^+ and \mathcal{E}_k^- have eccentricity arbitrarily close to 0 and major axis arbitrarily close to 1. \square

In the construction, we made *both* the eccentricity of the (filled) ellipses arbitrarily close to 0 and the ratio between the largest major axis and the smallest major axis arbitrarily close to 1. We know that this construction is very unlikely to work for the extreme case of unit disks, since a polynomial algorithm is known for MAX CLIQUE. Note that even with disks of arbitrary radii, Theorem 1.1 unconditionally proves that the construction does fail. Indeed, the co-2-subdivision of $C_3 + C_3$ is the complement of $C_9 + C_9$, hence, not a disk graph.

As in the previous section, the constructions of Lemmas 7.11 and 7.12 require only polynomial precision (when single-exponential precision would be enough for a polytime algorithm). Hence,




Fig. 17. The layout of the disks \mathcal{D}_i , and the elliptical disks \mathcal{E}_k^+ and \mathcal{E}_k^- .

the lower bounds of Theorems 7.9 and 7.10 also hold when the geometric representation is part of the input.

7.3 Homothets of a Convex Polygon

Another natural direction of generalizing a result on disk intersection graphs is to consider *pseudodisk intersection graphs*, i.e., intersection graphs of collections of closed subsets of the plane (regions bounded by simple Jordan curves) that are pairwise in a *pseudodisk* relationship (see Kratochvíl [46]). Two regions A and B are in pseudodisk relation if both differences $A \setminus B$ and $B \setminus A$ are arc-connected. It is known that P_{hom} graphs, i.e., intersection graphs of homothetic copies of a fixed polygon P , are pseudodisk intersection graphs [4]. As shown by Brimkov et al., for every convex k -gon P , a P_{hom} graph with n vertices has at most n^k maximal cliques [20]. This clearly implies that MAXIMUM CLIQUE, but also CLIQUE p -PARTITION for fixed p is polynomially solvable in P_{hom} graphs. Actually, the bound on the maximum number of maximal cliques from Reference [20] holds for a more general class of graphs, called k_{DIR} -CONV, which admit a intersection representation by convex polygons, whose every side is parallel to one of k directions.

Moreover, we observe that Theorem 3.6 cannot be generalized to P_{hom} graphs or k_{DIR} -CONV graphs. Indeed, consider the complement \overline{P}_n of an n -vertex path P_n . The number of maximal cliques in \overline{P}_n , or, equivalently, maximal independent sets in P_n is $\Theta(c^n)$ for $c \approx 1.32$, i.e., exponential in n [35]. Therefore, for every fixed polygon P (or for every fixed k) there is n , such that \overline{P}_n is not a P_{hom} (k_{DIR} -CONV) graph.

8 REMARKS AND FURTHER DIRECTIONS

The algorithm of Theorem 1.3 also works for weighted graphs.

THEOREM 8.1. *MAXIMUM WEIGHTED INDEPENDENT SET admits a randomized EPTAS and a deterministic PTAS on disk and unit ball graphs.*

The slight modifications to approximate MAXIMUM WEIGHTED INDEPENDENT SET consist in sampling S proportionally to the weights, and to remove the *lightest* layer L_i (among the first $\lceil 2/\beta\epsilon \rceil$) and the *lightest* set S^V (rather than the ones of minimum size). We then use repeatedly that MAXIMUM WEIGHTED INDEPENDENT SET can be solved in polynomial time on bipartite graphs.

This implies a randomized EPTAS for MAXIMUM WEIGHTED CLIQUE on disk graphs and unit ball graphs, with the same arguments used to get $|I| \geq \beta|V(G)|$. Now, what we obtain is $w(I) \geq \beta w(V(G))$, where w is the weight function and $w(X) := \sum_{u \in X} w(u)$.

One might wonder what is the constant hidden in $O(1)$ in the time complexity of the randomized EPTAS $f(\varepsilon)n^{O(1)}$. Here is how to achieve near quadratic time $f(\varepsilon)n^2 \log n$ where n is the number of vertices of our *unit ball graph* G (the geometric representation is not required). The first observation is that instead of finding an optimum stable set in a bipartite graph (which we do several times as a subroutine), one only needs a $(1 - \varepsilon)$ -approximation of it. This can be done in time $f(\varepsilon)m$, where m is the number of edges [31]. We also need to overcome our first branching guess of a vertex that belongs to an optimal solution (and hence multiply by n our complexity). To achieve this, we start by packing greedily disjoint neighborhoods $N(v_1), N(v_2), \dots, N(v_k)$ of G , while it is possible. Then, we consider the set of subgraphs G_1, G_2, \dots, G_k induced by the vertices at distance at most 3 of v_1, v_2, \dots, v_k , respectively. Observe that by maximality of the packing, every vertex is at distance at most 2 of some v_i , and thus, every edge, and even every clique of G belongs to at least one G_i . By a volume argument, any vertex, and thus any edge, belongs to at most a constant number of graphs G_i . Thus, we only need to compute the maximum clique over the graphs G_i , each with number of edges m_i , with the property that $\sum_{i=1}^k m_i = O(m)$. Now, the difference is that the clique number of each G_i is at least a constant fraction of its number of vertices, so the sampling step succeeds with probability at least a positive constant. We consider the complement of each graph G_i and approximate the independence number. The main and only obstacle is that finding a shortest odd cycle in quadratic time seems hopeless. Indeed, the other elements of the algorithm: Removing $N[S]$, computing the sets L_i and S^Y , removing the lightest of them, and $(1 - \varepsilon)$ -approximating the maximum stable set in a bipartite graph can all be done in quadratic time.

We only compute one (not necessarily shorter) odd cycle of length h , via breadth-first search, for instance. We can assume that $h = \omega(1/\varepsilon^2)$, otherwise, we are done. We take potentially *thicker* slices for S^Y : Instead of $z = \Theta(1/\varepsilon)$ consecutive strata in each layer up to L_i , we take $10\beta\varepsilon h$ consecutive strata. Either the suggested 2-coloring of $H'' - S^Y$ is proper and we are done. Or there is a monochromatic edge, which, going through the different cases of Claim 2, yields a new odd cycle of length $O(1/\varepsilon)$ (which is an easy case to handle) or short-cutting the previous odd cycle by at least $\Omega(\varepsilon h)$ vertices. In the latter case, our new odd cycle is shorter by a constant multiplicative factor $1 - \varepsilon$. Hence, after at most $O(\log n)$ improvements, we find an odd cycle that is small enough to conclude. Therefore, the overall complexity of the algorithm is $f(\varepsilon)n^2 \log n$.

The obvious remaining question is the complexity of MAXIMUM CLIQUE in disk graphs and in unit ball graphs. We showed why the versatile approach of representing complements of even subdivisions of graphs forming a class on which MAXIMUM INDEPENDENT SET is NP-hard fails if the class is *general graphs*, *planar graphs*, or even any class containing the disjoint union of two odd cycles. This approach was used by Middendorf and Pfeiffer for some string graphs [51] (with the class of all graphs); Cabello et al. [23] to settle the then long-standing open question of the complexity of MAXIMUM CLIQUE for segments (with the class of planar graphs), in Section 7 of this article, for filled ellipses, filled triangles, quasi unit ball graphs, and 4-dimensional unit ball graphs (with the class of all graphs). Determining the complexity of MAXIMUM INDEPENDENT SET on graphs without two vertex-disjoint odd cycles as an induced subgraph is a valuable first step towards settling the complexity of MAXIMUM CLIQUE on disks.

An interesting direction would be to find a toy problem on which we could prove NP-hardness. A nice class, which appears to be a subclass of unit ball graphs, is that of the so-called *Borsuk graphs*: We are given some (small) real $\varepsilon > 0$ and a finite collection V of unit vectors in \mathbb{R}^3 . The

Borsuk graph $B(V, \varepsilon)$ has vertex set V and its edges are all pairs $\{v, v'\}$ whose dot product is at most $-1 + \varepsilon$ (i.e., near antipodal).

The difficulty of computing the (weighted) independence number on Borsuk graphs is also an open question. A notable subclass of Borsuk graphs where this problem is polynomial-time solvable is the class of the quadrangulations of the projective plane. These well-studied objects have the striking property to be either bipartite or 4-chromatic. Furthermore, the odd cycle packing number of these graphs is at most 1. Artmann et al. recently showed that so-called *bimodular integer programming*, that is, integer programming where the constraint matrix has full rank and all its subdeterminants are in $\{-2, -1, 0, 1, 2\}$, can be solved in strongly polynomial time [9]. They also observe that MAXIMUM WEIGHTED INDEPENDENT SET on graphs with $\text{ocp} \geq 1$ is a bimodular integer programming problem. This implies the tractability of computing the weighted independence number on quadrangulations of the projective plane.

A second question is to derandomize the EPTAS. The difficulty here is concentrated in the sampling. The VC-dimension argument seems easy to deal with, however, we need that our sampling falls in the maximum independent set (or at least in some independent set that is close to maximum).

Another natural question is to find a superclass of geometric intersection graphs that both contain unit ball graphs and disk graphs. More generally, is it possible to explain why we have the same forbidden induced subgraph (the complement of a disjoint union of two odd cycles) for disk graphs and unit ball graphs? The suggested proofs for Theorem 1.1 and Theorem 1.2 are quite different.

Let us call *quasi unit disk graphs* those disk graphs that can be realized for any $\varepsilon > 0$ with disks having all the radii in the interval $[1, 1 + \varepsilon]$. Recall that we showed that, for the clique problem, quasi unit ball graphs are unlikely to have a QPTAS, while unit ball graphs admit an EPTAS. In dimension 2, it can be easily shown that unit disk graphs form a proper subset of quasi unit disk graphs, which form themselves a proper subset of disk graphs. Can we find for this intermediate class an efficient exact algorithm solving MAXIMUM CLIQUE?

PROBLEM 1. *Is there a polynomial-time algorithm for MAXIMUM CLIQUE on quasi unit disk graphs?*

Our randomized EPTAS works for MAXIMUM INDEPENDENT SET under three assumptions. While it is clear that we crucially need that $\text{iocp} \leq 1$ (or at least that iocp is constant), as far as we can tell, the boundedness of the VC-dimension and the fact that the solution is of linear size might not be required. Recall that we did obtain a randomized PTAS in the class $\mathcal{X}(\infty, \beta, O(1))$, and a deterministic PTAS in $\mathcal{X}(d, 0, O(1))$.

PROBLEM 2. *Is there an EPTAS for MAXIMUM INDEPENDENT SET on graphs without the union of two odd cycles as an induced subgraph, or even with $\text{iocp} = O(1)$, and either one of the following conditions:*

- *there is a solution of size at least βn for some constant β ,*
- *the VC-dimension of the graph is bounded by a constant d ?*

As this article was under review, Dvořák and Pekárek [32] answered the first item positively. More precisely, they obtained a randomized EPTAS for the class $\mathcal{X}(\infty, \beta, i)$ with running time $f(\varepsilon)\tilde{O}(n^{i+4})$. Matching the $f(\varepsilon)\tilde{O}(n^2)$ -time for unit balls (hence, with $i = 1$), as well as the second item remain open.

It might also be that no additional condition is needed.

PROBLEM 3. *Is there a(n E)PTAS for MAXIMUM INDEPENDENT SET on graphs without the union of two odd cycles as an induced subgraph?*

Atminas and Zamaraev [10] showed that the complement of $K_2 + C_s$ is not a unit disk graph when s is odd (where K_2 is an edge and C_s is a cycle on s vertices). Is this obstruction enough to obtain an alternative polynomial-time algorithm for MAXIMUM CLIQUE on unit disk graphs?

PROBLEM 4. *Is MAXIMUM INDEPENDENT SET solvable in polynomial-time on graphs excluding the union of an edge and an odd cycle as an induced subgraph?*

REFERENCES

- [1] Peyman Afshani and Timothy M. Chan. 2005. Approximation algorithms for maximum cliques in 3D unit-disk graphs. In *Proceedings of the 17th Canadian Conference on Computational Geometry (CCCG'05)*. 19–22. Retrieved from <http://www.cccg.ca/proceedings/2005/69.pdf>.
- [2] Peyman Afshani and Hamed Hatami. 2008. Approximation and inapproximability results for maximum clique of disc graphs in high dimensions. *Inf. Proc. Lett.* 105, 3 (2008), 83–87. DOI: <https://doi.org/10.1016/j.ipl.2007.08.005>
- [3] Pankaj K. Agarwal and Nabil H. Mustafa. 2006. Independent set of intersection graphs of convex objects in 2D. *Comput. Geom.* 34, 2 (2006), 83–95. DOI: <https://doi.org/10.1016/j.comgeo.2005.12.001>
- [4] Pankaj K. Agarwal, János Pach, and Micha Sharir. 2008. State of the union (of geometric objects). In *Surveys in Discrete and Computational Geometry: Twenty Years Later. Contemporary Mathematics*, Vol. 453. American Mathematical Society, 9–48.
- [5] Jochen Alber and Jiri Fiala. 2004. Geometric separation and exact solutions for the parameterized independent set problem on disk graphs. *J. Algor.* 52, 2 (2004), 134–151. DOI: <https://doi.org/10.1016/j.jalgor.2003.10.001>
- [6] Noga Alon, Raphael Yuster, and Uri Zwick. 1997. Finding and counting given length cycles. *Algorithmica* 17, 3 (1997), 209–223. DOI: <https://doi.org/10.1007/BF02523189>
- [7] Christoph Ambühl and Uli Wagner. 2005. The clique problem in intersection graphs of ellipses and triangles. *Theor. Comput. Syst.* 38, 3 (2005), 279–292. DOI: <https://doi.org/10.1007/s00224-005-1141-6>
- [8] Boris Aronov, Anirudh Donakonda, Esther Ezra, and Rom Pinchasi. 2018. On pseudo-disk hypergraphs. *CoRR* abs/1802.08799 (2018).
- [9] Stephan Artmann, Robert Weismantel, and Rico Zenklusen. 2017. A strongly polynomial algorithm for bimodular integer linear programming. In *Proceedings of the 49th ACM SIGACT Symposium on Theory of Computing (STOC'17)*. 1206–1219. DOI: <https://doi.org/10.1145/3055399.3055473>
- [10] Aistis Atminas and Viktor Zamaraev. 2018. On forbidden induced subgraphs for unit disk graphs. *Disc. Comput. Geom.* 60, 1 (2018), 57–97. DOI: <https://doi.org/10.1007/s00454-018-9968-1>
- [11] Jørgen Bang-Jensen, Bruce Reed, Mathias Schacht, Robert Šámal, Bjarne Toft, and Uli Wagner. 2006. On six problems posed by Jarik Nešetřil. *Top. Disc. Math.* 26 (2006), 613–627.
- [12] Csaba Biró, Édouard Bonnet, Dániel Marx, Tillmann Miltzow, and Paweł Rzażewski. 2018. Fine-grained complexity of coloring unit disks and balls. *J. Comput. Geom.* 9, 2 (2018), 47–80. Retrieved from <http://jocg.org/index.php/jocg/article/view/414>.
- [13] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. 1989. Learnability and the Vapnik-Chervonenkis dimension. *J. ACM* 36, 4 (1989), 929–965. DOI: <https://doi.org/10.1145/76359.76371>
- [14] Adrian Bock, Yuri Faenza, Carsten Moldenhauer, and Andres J. Ruiz-Vargas. 2014. Solving the stable set problem in terms of the odd cycle packing number. In *Proceedings of the 34th International Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS'14)*. 187–198. DOI: <https://doi.org/10.4230/LIPIcs.FSTTCS.2014.187>
- [15] Marthe Bonamy, Édouard Bonnet, Nicolas Bousquet, Pierre Charbit, and Stéphan Thomassé. 2018. EPTAS for max clique on disks and unit balls. In *Proceedings of the 59th IEEE Symposium on Foundations of Computer Science (FOCS'18)*. 568–579. DOI: <https://doi.org/10.1109/FOCS.2018.00060>
- [16] Édouard Bonnet, Bruno Escoffier, Eun Jung Kim, and Vangelis Th. Paschos. 2015. On subexponential and FPT-time inapproximability. *Algorithmica* 71, 3 (2015), 541–565. DOI: <https://doi.org/10.1007/s00453-014-9889-1>
- [17] Édouard Bonnet, Panos Giannopoulos, Eun Jung Kim, Paweł Rzażewski, and Florian Sikora. 2018. QPTAS and subexponential algorithm for maximum clique on disk graphs. In *Proceedings of the 34th International Symposium on Computational Geometry (SoCG'18)*. 12:1–12:15. DOI: <https://doi.org/10.4230/LIPIcs.SoCG.2018.12>
- [18] Andreas Brandstädt, Van Bang Le, and Jeremy P. Spinrad. 1999. *Graph Classes: A Survey*. SIAM.
- [19] Heinz Breu and David G. Kirkpatrick. 1998. Unit disk graph recognition is NP-hard. *Comput. Geom.* 9, 1–2 (1998), 3–24. DOI: [https://doi.org/10.1016/S0925-7721\(97\)00014-X](https://doi.org/10.1016/S0925-7721(97)00014-X)
- [20] Valentin E. Brimkov, Konstanty Junosza-Szaniawski, Sean Kafer, Jan Kratochvíl, Martin Pergel, Paweł Rzażewski, Matthew Szczepankiewicz, and Joshua Terhaar. 2018. Homothetic polygons and beyond: Maximal cliques in intersection graphs. *Disc. Appl. Math.* 247 (2018), 263–277. DOI: <https://doi.org/10.1016/j.dam.2018.03.046>

- [21] Sergio Cabello. 2015. Maximum clique for disks of two sizes. *Open problems from Geometric Intersection Graphs: Problems and Directions CG Week Workshop*. Retrieved from <http://cgweek15.tcs.uj.edu.pl/problems.pdf>.
- [22] Sergio Cabello. 2015. *Open problems presented at the Algorithmic Graph Theory on the Adriatic Coast Workshop*. Retrieved from <https://conferences.matheo.si/event/6/picture/35.pdf>.
- [23] Sergio Cabello, Jean Cardinal, and Stefan Langerman. 2013. The clique problem in ray intersection graphs. *Disc. Comput. Geom.* 50, 3 (2013), 771–783. DOI: <https://doi.org/10.1007/s00454-013-9538-5>
- [24] Paz Carmi, Matthew J. Katz, and Pat Morin. 2018. Stabbing pairwise intersecting disks by four points. *CoRR abs/1812.06907* (2018).
- [25] Stephan Ceroi. 2002. *The Clique Number of Unit Quasi-disk Graphs*. Technical Report RR-4419. INRIA. Retrieved from <https://hal.inria.fr/inria-00072169>.
- [26] Timothy M. Chan. 2003. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algor.* 46, 2 (2003), 178–189. DOI: [https://doi.org/10.1016/S0196-6774\(02\)00294-8](https://doi.org/10.1016/S0196-6774(02)00294-8)
- [27] Timothy M. Chan and Sariel Har-Peled. 2012. Approximation algorithms for maximum independent set of pseudo-disks. *Disc. Comput. Geom.* 48, 2 (2012), 373–392. DOI: <https://doi.org/10.1007/s00454-012-9417-5>
- [28] Miroslav Chlebik and Janka Chlebiková. 2006. Complexity of approximating bounded variants of optimization problems. *Theor. Comput. Sci.* 354, 3 (2006), 320–338. DOI: <https://doi.org/10.1016/j.tcs.2005.11.029>
- [29] Brent N. Clark, Charles J. Colbourn, and David S. Johnson. 1990. Unit disk graphs. *Disc. Math.* 86, 1–3 (1990), 165–177. DOI: [https://doi.org/10.1016/0012-365X\(90\)90358-O](https://doi.org/10.1016/0012-365X(90)90358-O)
- [30] L. Danzer. 1986. Zur Lösung des Gallaischen Problems über Kreisscheiben in der Euklidischen Ebene. *Studia Sci. Math. Hungar* 21, 1–2 (1986), 111–134.
- [31] Ran Duan and Seth Pettie. 2014. Linear-time approximation for maximum weight matching. *J. ACM* 61, 1 (2014), 1:1–1:23. DOI: <https://doi.org/10.1145/2529989>
- [32] Zdenek Dvořák and Jakub Pekárek. 2020. Induced odd cycle packing number, independent sets, and chromatic number. *CoRR abs/2001.02411* (2020).
- [33] Thomas Erlebach, Klaus Jansen, and Eike Seidel. 2005. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.* 34, 6 (2005), 1302–1323. DOI: <https://doi.org/10.1137/S0097539702402676>
- [34] Aleksei V. Fishkin. 2003. Disk graphs: A short survey. In *Approximation and Online Algorithms, First International Workshop, WAOA 2003, Budapest, Hungary, September 16–18, 2003, Revised Papers (Lecture Notes in Computer Science)*, Klaus Jansen and Roberto Solis-Oba (Eds.), Vol. 2909. Springer, 260–264. DOI: https://doi.org/10.1007/978-3-540-24592-6_23
- [35] Zoltán Füredi. 1987. The number of maximal independent sets in connected graphs. *J. Graph Theor.* 11, 4 (1987), 463–470. DOI: <https://doi.org/10.1002/jgt.3190110403>
- [36] Matt Gibson and Imran A. Pirwani. 2010. Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier—(Extended abstract). In *Proceedings of the 18th European Symposium on Algorithms (ESA'10)*. 243–254. DOI: https://doi.org/10.1007/978-3-642-15775-2_21
- [37] Ervin Györi, Alexandr V. Kostochka, and Tomasz Łuczak. 1997. Graphs without short odd cycles are nearly bipartite. *Disc. Math.* 163, 1 (1997), 279–284. DOI: [https://doi.org/10.1016/0012-365X\(95\)00321-M](https://doi.org/10.1016/0012-365X(95)00321-M)
- [38] Sariel Har-Peled, Haim Kaplan, Wolfgang Mulzer, Liam Roditty, Paul Seiferth, Micha Sharir, and Max Willert. 2018. Stabbing pairwise intersecting disks by five points. In *Proceedings of the 29th International Symposium on Algorithms and Computation (ISAAC'18)*. 50:1–50:12. DOI: <https://doi.org/10.4230/LIPIcs.ISAAC.2018.50>
- [39] D. Haussler and E. Welzl. 1987. Epsilon-nets and simplex range queries. *Disc. Comput. Geom.* 2 (1987), 127–151.
- [40] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.* 63, 4 (Dec. 2001), 512–530.
- [41] Ross J. Kang and Tobias Müller. 2012. Sphere and dot product representations of graphs. *Disc. Comput. Geom.* 47, 3 (2012), 548–568. DOI: <https://doi.org/10.1007/s00454-012-9394-8>
- [42] Ken-ichi Kawarabayashi and Bruce A. Reed. 2010. Odd cycle packing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*. 695–704. DOI: <https://doi.org/10.1145/1806689.1806785>
- [43] Chaya Keller, Shakhar Smorodinsky, and Gábor Tardos. 2017. On Max-Clique for intersection graphs of sets and the Hadwiger-Debrunner numbers. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA'17)*. 2254–2263. DOI: <https://doi.org/10.1137/1.9781611974782.148>
- [44] Paul Koebe. 1936. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physikalische Klasse* 88 (1936), 141–164.
- [45] Jan Kratochvíl. 1993. Precoloring extension with fixed color bound. *Acta Math. Univ. Comen* 62 (1993), 139–153.
- [46] Jan Kratochvíl. 1996. Intersection graphs of noncrossing arc-connected sets in the plane. In *Proceedings of the Symposium on Graph Drawing (GD'96)*. 257–270. DOI: https://doi.org/10.1007/3-540-62495-3_53
- [47] Jan Kratochvíl and Jiří Matoušek. 1994. Intersection graphs of segments. *J. Comb. Theor., Ser. B* 62, 2 (1994), 289–315. DOI: <https://doi.org/10.1006/jctb.1994.1071>

- [48] Ewa Malesinska, Steffen Piskorz, and Gerhard Weïßenfels. 1998. On the chromatic number of disk graphs. *Networks* 32, 1 (1998), 13–22. DOI : [https://doi.org/10.1002/\(SICI\)1097-0037\(199808\)32:1\(13::AID-NET2\)3.0.CO;2-M](https://doi.org/10.1002/(SICI)1097-0037(199808)32:1(13::AID-NET2)3.0.CO;2-M)
- [49] Dániel Marx and Michal Pilipczuk. 2015. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. In *Proceedings of the 23rd European Symposium on Algorithms (ESA'15)*. 865–877. DOI : https://doi.org/10.1007/978-3-662-48350-3_72
- [50] Terry A. McKee and Fred R. McMorris. 1999. *Topics in Intersection Graph Theory*. SIAM.
- [51] Matthias Middendorf and Frank Pfeiffer. 1992. The max clique problem in classes of string-graphs. *Disc. Math.* 108, 1–3 (1992), 365–372. DOI : [https://doi.org/10.1016/0012-365X\(92\)90688-C](https://doi.org/10.1016/0012-365X(92)90688-C)
- [52] Dana Moshkovitz and Ran Raz. 2010. Two-query PCP with subconstant error. *J. ACM* 57, 5 (2010), 29:1–29:29. DOI : <https://doi.org/10.1145/1754399.1754402>
- [53] Tim Nieberg and Johann Hurink. 2005. A PTAS for the minimum dominating set problem in unit disk graphs. In *Proceedings of the 3rd International Workshop on Approximation and Online Algorithms (WAOA'15)*. 296–306. DOI : https://doi.org/10.1007/11671411_23
- [54] Tim Nieberg, Johann Hurink, and Walter Kern. 2004. A robust PTAS for maximum weight independent sets in unit disk graphs. In *Proceedings of the 30th International Workshop on Graph-theoretic Concepts in Computer Science (WG'04)*. 214–221. DOI : https://doi.org/10.1007/978-3-540-30559-0_18
- [55] Vijay Raghavan and Jeremy P. Spinrad. 2003. Robust algorithms for restricted domains. *J. Algor.* 48, 1 (2003), 160–172. DOI : [https://doi.org/10.1016/S0196-6774\(03\)00048-8](https://doi.org/10.1016/S0196-6774(03)00048-8)
- [56] Warren D. Smith and Nicholas C. Wormald. 1998. Geometric separator theorems & applications. In *Proceedings of the 39th Symposium on Foundations of Computer Science (FOCS'98)*. 232–243. DOI : <https://doi.org/10.1109/SFCS.1998.743449>
- [57] Lajos Stachó. 1981. A solution of Gallai's problem on pinning down circles. *Mat. Lapok* 32, 1–3 (1981), 19–47.
- [58] Peter Guthrie Tait. 1877. Some elementary properties of closed plane curves. *Mess. Math., New Series* 69 (1877), 270–272.
- [59] Erik Jan van Leeuwen. 2006. Better approximation schemes for disk graphs. In *Proceedings of the 10th Scandinavian Workshop on Algorithm Theory (SWAT'06)*. 316–327. DOI : https://doi.org/10.1007/11785293_30
- [60] Erik Jan van Leeuwen. 2009. *Optimization and Approximation on Systems of Geometric Objects*. Ph.D. Dissertation. Utrecht University.
- [61] Vladimir N. Vapnik and A. Ya Chervonenkis. 2015. On the uniform convergence of relative frequencies of events to their probabilities. In *Meas. Complex*. Springer, 11–30.

Received March 2019; revised March 2020; accepted November 2020