

EPTAS for Max Clique on Disks and Unit Balls

Marthe Bonamy¹, Édouard Bonnet², Nicolas Bousquet³, Pierre Charbit^{2,4}, and Stéphan Thomassé^{2,5}

- 1 CNRS, LaBRI, Université de Bordeaux
marthe.bonamy@u-bordeaux.fr
- 2 ENS Lyon, LIP, Lyon, France
edouard.bonnet@dauphine.fr, stephan.thomasse@ens-lyon.fr
- 3 CNRS, G-SCOP laboratory, Grenoble-INP, France
nicolas.bousquet@grenoble-inp.fr
- 4 Université Paris Diderot - IRIF, Paris, France
charbit@irif.fr
- 5 Institut Universitaire de France

Abstract

We propose a polynomial-time algorithm which takes as input a finite set of points of \mathbb{R}^3 and compute, up to arbitrary precision, a maximum subset with diameter at most 1. More precisely, we give the first randomized EPTAS and deterministic PTAS for MAXIMUM CLIQUE in unit ball graphs. Our approximation algorithm also works on disk graphs with arbitrary radii.

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. Recently, it was shown that the disjoint union of two odd cycles is never the complement of a disk graph [Bonnet, Giannopoulos, Kim, Rzażewski, Sikora; SoCG '18]. This enabled the authors to derive a QPTAS and a subexponential algorithm for MAX CLIQUE on disk graphs.

In this paper, we improve the approximability to a randomized EPTAS (and a deterministic PTAS). More precisely, we 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 large independence number. We then address the question of computing MAX CLIQUE for disks in higher dimensions. We show that intersection graphs of unit balls, like disk graphs, do not admit the complement of two odd cycles as an induced subgraph. This, in combination with the first result, straightforwardly yields a randomized EPTAS for MAX CLIQUE on unit ball graphs. In stark contrast, we show that on ball and unit 4-dimensional disk graphs, MAX CLIQUE is NP-hard and does not admit an approximation scheme even in subexponential-time, unless the Exponential Time Hypothesis fails.

Digital Object Identifier 10.4230/LIPIcs...

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 [12, 27]. 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 [20], and many standard optimization problems have been studied on disk graphs, see for example [34] 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 [25] on which several



licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of those problems are untractable. Although shifting techniques and separator theorems may often lead to subexponential classic or parameterized algorithms [3, 7, 26, 31]. Many approximation algorithms have been designed specifically on (unit) disk graphs, or more generally on geometric intersection graphs, see for instance [16, 19, 21, 28, 29, 33] 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 techniques that we mentioned are only amenable to packing and covering problems. The MAXIMUM CLIQUE problem is arguably the most prominent problem which 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. This is why, in this paper, we focus on solving MAXIMUM CLIQUE on (unit) disk graphs in dimension 2 or higher.

Previous results.

In 1990, Clark *et al.* [17] 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, remove all the centers that would contradict this maximality, observe that the resulting graph is co-bipartite, implying that an optimum solution can be found in polynomial time. However, recognizing unit disk graphs is NP-hard [13], and even $\exists\mathbb{R}$ -complete [24]. In particular, if the input is the mere unit disk graph, one cannot expect to efficiently compute a geometric representation in order to run the previous algorithm. Raghavan and Spinrad showed how to overcome this issue and suggested a polynomial-time algorithm which does not require the geometric representation [30]. 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 fifties (although the first published records of that result came later in the eighties) that the piercing number of pairwise intersecting disks is 4 [18, 32]. An account of this story can be found in a recent note by Har-Peled *et al.* [22]. Ambühl and Wagner observed that this yields a 2-approximation for MAXIMUM CLIQUE [5]. 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 both instance formed by one co-bipartite graph and return the larger solution. 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 [6], but was also explicitly mentioned as an open problem by Fishkin [20], Ambühl and Wagner [5]. Cabello even asked if there is a 1.99-approximation for disk graphs with

two sizes of radii [14, 15]. Recently, Bonnet *et al.* [10] showed that the disjoint union of two odd cycles is not the complement of a disk graph. From this result, they obtained a subexponential algorithm running in time $2^{\tilde{O}(n^{2/3})}$ for MAXIMUM CLIQUE on disk graphs, based on a win-win approach. They also got a QPTAS by calling a PTAS for MAXIMUM INDEPENDENT SET on graphs with sublinear odd cycle packing number due to Bock *et al.* [9], or branching on a low-degree vertex.

Our results.

Our main contributions are twofold. The first is an EPTAS (Efficient Polynomial-Time Approximation Scheme, that is, a PTAS in time $f(\varepsilon)n^{O(1)}$) for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, \beta, 1)$ with constant VC-dimension d , independence number at least βn , and no disjoint union of two odd cycles as an induced subgraph.

► **Theorem 1.** *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 $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$ for MAXIMUM INDEPENDENT SET on graphs of $\mathcal{X}(d, \beta, 1)$ with n vertices.*

Using the forbidden induced subgraph result of Bonnet *et al.* [10], it is then easy to reduce MAXIMUM CLIQUE on disk graphs to MAXIMUM INDEPENDENT SET on $\mathcal{X}(3, \beta, 1)$ for some constant β . We therefore obtain an EPTAS for MAXIMUM CLIQUE on disk graphs, settling almost¹ completely the approximability of this problem.

► **Theorem 2.** *There is a randomized EPTAS for MAXIMUM CLIQUE on disk graphs, even without geometric representation. Its running time is $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$ for a $(1 - \varepsilon)$ -approximation on a graph with n vertices.*

The second contribution is to show the same forbidden induced subgraph for unit ball graphs as the one obtained for disk graphs. The proofs are radically different and the classes are incomparable. So the fact that the same obstruction applies for disk graphs and unit ball graphs might simply be a coincidence.

► **Theorem 3.** *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 $\text{iocp}(\overline{G}) \leq 1$.*

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. Again, Theorem 1 and Theorem 3 naturally lead to:

► **Theorem 4.** *There is a randomized EPTAS in time $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$ for MAXIMUM CLIQUE on unit ball graphs, even without the geometric representation.*

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, and unit 4-dimensional disk graphs. Our lowerbounds also imply NP-hardness. To the best of our knowledge, the NP-hardness of MAXIMUM CLIQUE on unit d -dimensional disk graphs was only known when d is superconstant ($d = \Omega(\log n)$) [2].

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

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

EPTAS for Maximum Independent Set on $\mathcal{X}(d, \beta, 1)$

The first main result of this paper asserts that if a graph G satisfies that every two odd cycles are joined by an edge, the Vapnik-Chervonenkis dimension 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 computed in polytime 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 hopefully in a fixed optimum solution I . We include the selected vertices in our solution and remove their neighborhood from the graph. Due to the classic result of Haussler and Welzl [23] 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 the shortest among the ones 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 midpoint is attached to the origin, one can apply a continuous motion in order to turn it around (a motion à la Kakeya, henceforth *Kakeya motion*). If we now consider two needles, each with a Kakeya motion, then the two needles will go through the same position. This simply follows from the fact that two antipodal curves on the 2-sphere intersect. The second main result of this paper is a translation of this Jordan-type theorem in terms of intersection graphs: The complement of unit balls does not contain the disjoint union of two odd cycles. The proof can really be seen as two Kakeya motions, each one along the two odd cycles, leading to a contradiction when the needles achieve parallel directions.

Together with the first result, it implies a randomized EPTAS for MAXIMUM CLIQUE on disk graphs, and 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.

Organization

The rest of the paper 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 of graphs parameterized by three constants: the VC dimension, the ratio *independence number divided by number of vertices*, and the maximum number of odd cycles that can be found as a disjoint union in an induced subgraph. In Section 3, we design a randomized EPTAS for MAXIMUM INDEPENDENT SET on this class. In Section 4, we show that complements of unit ball graphs do not have a disjoint union of two odd cycles as an induced subgraph. This yields a randomized EPTAS for MAXIMUM CLIQUE on unit ball graphs (as well as on disk graphs). This is tight in two directions: having different values of radii, and the dimension of the ambient space. Indeed, we complement this positive result by showing that MAXIMUM CLIQUE is unlikely to even have a QPTAS on ball graphs where all the radii are arbitrarily close to 1, and on 4-dimensional unit ball graphs.

2 Preliminaries

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 $\text{iocp}(G)$, is the maximum number of vertex-disjoint odd cycles with no edge between any two of them.

VC-dimension

VC-dimension has been introduced by Vapnik and Chervonenkis in the seminal paper [35]. Let $H = (V, E)$ be a hypergraph. A set X 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 hypergraph $(V(G), \{N_G(v) \mid v \in V(G)\})$.

Geometric notations

For a positive integer d , we denote by \mathbb{R}^d the d -dimensional euclidean space. If x and y are two points of \mathbb{R}^d , xy is the straight-line segment whose endpoints are x and y . We denote by $d(x, y)$ the euclidean distance between x and y . A *d -dimensional closed disk* 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 the quantity $\sup_{x, y \in S} d(x, y)$.

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 disk graph* is the intersection graph of d -dimensional closed disks of \mathbb{R}^d . We shorten 2-dimensional disk graph in *disk graph*, and 3-dimensional disk graph in *ball graph*. A *d -dimensional unit disk graph* is the intersection graph of unit d -dimensional closed disks of \mathbb{R}^d , that is, disks with radius 1. Unit d -dimensional disk graphs can be thought of only with points: vertices are points (at the center of the disks) 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 subcollection of points whose diameter is at most a fixed value.

Bonnet *et al.* established the following obstruction for disk graphs.

► **Theorem 5** ([10]). *The complement of a disk graph cannot have the disjoint union of two odd cycles as an induced subgraph.*

This can be equivalently rephrased as: if G is a disk graph, then $\text{iocp}(\overline{G}) \leq 1$.

Approximation schemes

A PTAS (Polynomial-Time Approximation Scheme) for a minimization (resp. maximization) problem is an approximation algorithm which takes an additional parameter $\varepsilon > 0$ and outputs in time $n^{f(\varepsilon)}$ a solution of value at most $(1 + \varepsilon)\text{OPT}$ (resp. 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) is in time $\frac{1}{\varepsilon}n^{O(1)}$, a QPTAS (Quasi PTAS) is in time $n^{\text{polylog } n}$ for every ε . Finally, and this is quite informal and not standard, we call SUBEXPAS (subexponential AS) an approximation scheme with running time $2^{n^{0.99}}$ for every ε . All those approximation schemes can come deterministic or randomized.

The class $\mathcal{X}(d, \beta, i)$

In the next section, 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 induced odd cycle packing number equals 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 $\mathcal{X}(d, \beta, i)$. It will turn out that the complement of the interesting instances are in $\mathcal{X}(3, \frac{1}{6}, 1)$ (or $\mathcal{X}(3, \frac{1}{4}, 1)$) for disk graphs, and in $\mathcal{X}(4, \frac{1}{12}, 1)$ for unit ball graphs.

3 EPTAS for Maximum Independent Set on $\mathcal{X}(d, \beta, i)$

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

► **Theorem 1.** *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 $2^{\tilde{O}(1/\varepsilon^3)}n^{O(1)}$ 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 vertex-set 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 $\text{iocp}(H) \leq 1$. For instance, this implies that for any odd cycle C , the set $N[C]$ is an odd-cycle transversal.

Let $c := 8(\frac{1}{(\beta\varepsilon)^2} + \frac{1}{\beta\varepsilon} + 1) = O(1/\varepsilon^2)$, $\delta := \frac{\varepsilon}{c} = O(\varepsilon^3)$, and $s := \frac{10d}{\delta} \log \frac{1}{\delta}$. We call *short odd cycle* an odd cycle of length at most c , and *long 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, $|I| > 2s$.

► **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 d . By definition of K , all its edges have size at least $\delta|V(K)|$. A celebrated result in VC-dimension theory by Haussler and Welzl [23], later improved by Blumer *et al.* [8], ensures that every such hypergraph K admits a hitting set (a set of vertices that intersects every edge) of size at most $\frac{10d}{\delta} \log \frac{1}{\delta}$. ◀

Algorithmically, we have two ways of selecting the set S , leading to a deterministic PTAS 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 $n^{f(\varepsilon)}$ possible sets), or we use another result proven in [23]: 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 $t := \lceil \frac{\log(10^{-10})}{\log(1 - (\beta/2)^s)} \rceil = 2^{\tilde{O}(1/\varepsilon^3)}$ 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). 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 [4]) that determines whether the graph is bipartite and, if not, outputs a shortest odd cycle C_{og} 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 odd cycle, then $|N_{H'}(C_{\text{og}}) \cap I| \leq c\delta|I| = \varepsilon|I|$, and therefore $T := N_H(S) \cup N_{H'}[C_{\text{og}}]$ is an odd-cycle transversal of H with $|T \cap I| \leq \varepsilon|I|$.

We can now safely assume that $g > c$, i.e., C_{og} is a long 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 $L_i := \{v \mid \exists u \in L_{i-1} \text{ with } uv \in E(H') \text{ and } v \notin L_j \text{ for } j < i\}$. 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, there is one with index at most $\frac{2}{\beta\varepsilon}$ that contains at most $\frac{\varepsilon\beta}{2}n \leq \frac{\varepsilon}{2}|I|$ 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.

- Secondly, we show that this component admits an odd-cycle transversal of size at most $\frac{\varepsilon}{2}|I|$ (informally, the neighborhood at distance up to $O(\frac{1}{\varepsilon})$ of $O(\frac{1}{\varepsilon})$ consecutive vertices on the cycle C_{og}).

In other words, we can find $\varepsilon|I|$ vertices whose deletion yields a bipartite graph (see Figure 1), which together with $N(S)$ form the desired odd-cycle transversal.




Figure 1 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 $\lceil \frac{2}{\beta\varepsilon} \rceil$ layers. To the right of this cutset, we know that the graph is bipartite. This brings us back to the case with fewer than $\frac{2}{\beta\varepsilon}$ layers, where we can find a small odd cycle transversal (in red).

If $\lambda > \frac{2}{\beta\varepsilon}$, then there is some index $i \leq \lceil \frac{2}{\beta\varepsilon} \rceil$ such that L_i is of size at most $\frac{\beta\varepsilon}{2}n \leq \frac{\varepsilon}{2}|I|$. 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 $H'' := H'[C_{\text{og}} \cup \bigcup_{1 \leq j < i} L_j]$.

The graph H'' has less than $\frac{2}{\beta\varepsilon}$ layers emanating from C_{og} . We will find an odd-cycle transversal of size at most $\frac{\varepsilon}{2}|I|$. 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} which 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 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 $z := \lceil \frac{4}{\beta\varepsilon} \rceil + 2$ and for any non-negative integer γ at most $\frac{g}{z} - 1$, $S^\gamma := \bigcup_{\gamma z + 1 \leq j \leq (\gamma + 1)z} S_j$.

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

Proof. Observe that $\frac{g}{z} - 1 \geq \frac{2}{\beta\varepsilon}$, so each S^γ of the claim is well-defined. It holds that $C_{\text{og}} \cap S^\gamma = \bigcup_{\gamma z + 1 \leq j \leq (\gamma + 1)z} \{v_j\}$. We exhibit a proper 2-coloring of B by coloring its vertices as follows. We start by coloring each vertex of the path $C_{\text{og}} \setminus S^\gamma$ in an alternated manner, 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 $L_0 := C_{\text{og}}$). 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, $|P_1| = |P_2| = k < i \leq \frac{2}{\beta\varepsilon}$. 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 $2k+1 < \frac{4}{\beta\varepsilon} + 1 < g$, 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 $\min(\ell' - \ell, \ell + g - \ell') \geq \frac{4}{\beta\varepsilon} + 1$ since otherwise it would be possible to construct a shorter odd cycle 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 an odd walk of length $k + k' + 1$. However, a shortest path between v_ℓ and $v_{\ell'}$ within C_{og} has length $\min(\ell' - \ell, \ell + g - \ell') \geq \frac{4}{\beta\varepsilon} + 1 > k + k' + 1$. 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 $C_{\text{og}} - S^\gamma$, 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 P is a path of even length 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 odd 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 $C_{\text{og}} - S^\gamma$, 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. \blacktriangleleft

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

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  $\triangleright |V(H)| = \tilde{O}(1/\varepsilon^3)$ 
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  $s$   $\triangleright S \subseteq I$  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:     $H' \leftarrow H - N[S]$   $\triangleright$  remove  $S$  and its neighborhood
12:     $C_{\text{og}} \leftarrow$  shortest odd cycle in  $H'$   $\triangleright$  in polynomial time [4]
13:     $g \leftarrow |C_{\text{og}}|$   $\triangleright C_{\text{og}} = v_1v_2 \dots v_g$ 
14:    if  $g \leq c$  then  $\triangleright$  short odd cycle
15:       $S \leftarrow S \cup$  max stable on the bipartite  $H' - N[C_{\text{og}}]$   $\triangleright \text{iocp}(G) = 1$ 
16:    else  $\triangleright$  long odd cycle
17:       $L_\ell \leftarrow$  vertices of  $H'$  at distance exactly  $\ell$  from  $C_{\text{og}}$ 
18:       $L_i \leftarrow$  smallest layer among  $\{L_\ell\}_{1 \leq \ell \leq \lceil 2/\beta\varepsilon \rceil}$   $\triangleright |L_i| \leq \frac{\varepsilon}{2}\alpha(G)$ 
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:       $S^\gamma \leftarrow$  smallest set among  $\{\bigcup_{k \in [\gamma z + 1, (\gamma + 1)z]} S_k\}_{\gamma \in [0, \lfloor \frac{2}{\beta\varepsilon} \rfloor]}$   $\triangleright |S^\gamma| \leq \frac{\varepsilon}{2}\alpha(G)$ 
23:       $S \leftarrow S \cup$  max stable on the bipartite  $H'[\bigcup_{j > i} L_j]$   $\triangleright \text{iocp}(G) = 1$ 
24:       $S \leftarrow S \cup$  max stable on the bipartite  $H'' - S^\gamma$   $\triangleright$  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

► **Theorem 2.** *There is a randomized EPTAS for MAXIMUM CLIQUE on disk graphs, even without geometric representation. 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 could 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 [18, 32]. The number of faces in an arrangement of n circles (disk boundaries) is $O(n^2)$, and all the points within one face hits the same disks. In time $O(n^8)$, one can therefore exhaustively guess four points piercing a maximum clique \mathcal{C} . We can remove all the disks which are not hit by any of those four points, since they are not part of \mathcal{C} . This new instance G can have its vertices partitioned into four cliques, hence $\alpha(\overline{G}) \geq |V(G)|/4$.

Without a geometric representation, we suggest the following. Disk graphs are 6ω -degenerate (and closed by induced subgraphs), i.e., there is a vertex of degree at most 6ω , where ω is the clique number of the graph. Furthermore, the neighborhood of this vertex is easily partitionable into 6 cliques. We can find such a vertex in polynomial time. We branch on two outcomes. Either this vertex is in a maximum clique: we run the approximation of Theorem 1 on its closed neighborhood G which satisfies $\alpha(\overline{G}) \geq |V(G)|/6$. 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.

It is folklore that disk graphs have VC-dimension at most 3 (see for instance [11]) and, 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 3. Finally, by Theorem 5 [10], $\text{iocp}(\overline{G}) \leq 1$. We only call the approximation algorithm (a polynomial number of times) with disk graphs G such that $\overline{G} \in \mathcal{X}(3, \frac{1}{6}, 1)$ (argument without the geometric representation) or $\overline{G} \in \mathcal{X}(3, \frac{1}{4}, 1)$ (argument with the geometric representation). Hence, we conclude by Theorem 1. ◀

It is known that the VC-dimension of unit balls and the piercing number of pairwise intersecting unit balls are both constant. In Section 4, we will show that the induced odd cycle packing number of unit ball graphs is at most 1. This is the missing element for the EPTAS to also work for MAXIMUM CLIQUE on unit balls.

We conclude this section by extending the EPTAS to work for constant (not necessarily 1) induced odd cycle packing number.

► **Theorem 6.** *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 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. 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 line 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. ◀

4 Balls and higher dimensions

We start by showing the main result of this section: for every unit ball graph G , $iocp(\overline{G}) \leq 1$. A *closed polygonal chain* C in \mathbb{R}^d is defined by a set 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 7.** *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 which is common to $\text{Needle}(C_1)$ and $\text{Needle}(C_2)$. For that, we yield a continuous map from \mathbb{R}^+ to $\text{Needle}(C_1)$ that we naturally interpret as a curve on the 2-sphere. Indeed the set of directions in \mathbb{R}^3 is isomorphic to the set of points on the 2-sphere. As we do not need a specific parameterization of the curve, we just describe how we continuously move a vector \overrightarrow{ab} whose direction runs through the entire set $\text{Needle}(C_1)$. Let x_1, x_2, \dots, x_p be the vertices of 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 . For the next step, a is fixed at x_3 and b continuously moves from x_2 to 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 in x_1 and a is in x_2 . We do the same one more time, so that a is back to x_1 and b is back to x_2 , and we stop. One may observe that this process spans $\text{Needle}(C_1)$.

This defines a closed curve \mathcal{C}_1 on the 2-sphere since we are finally back to $\text{dir}(\overrightarrow{x_1x_2})$ from where we started. Furthermore, \mathcal{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 attain the opposite direction in $[x_{i-1} \rightarrow x_i \leftarrow x_{i+1}]$. Similarly $\text{Needle}(C_2)$ draws a closed antipodal curve on the 2-sphere \mathcal{C}_2 . The curves \mathcal{C}_1 and \mathcal{C}_2 intersect since they are closed and antipodal. An intersection point corresponds to a direction shared by $\text{Needle}(C_1)$ and $\text{Needle}(C_2)$. ◀

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 3.** *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 $iocp(\overline{G}) \leq 1$.*

Proof. Let $x_1, x_2, \dots, x_p \in \mathbb{R}^3$ ($y_1, y_2, \dots, y_q \in \mathbb{R}^3$) be the centers of unit balls representing the complement of an odd cycle of length p (resp. q), such that x_i and x_{i+1} (resp. y_i and y_{i+1}) encode the non-adjacent pairs. Let C_1 (resp. C_2) be the closed polygonal chain obtained from the centers x_1, x_2, \dots, x_p (resp. y_1, y_2, \dots, y_q). By Lemma 7, 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 $y_{j-1}y_{j+1}$.



■ **Figure 2** The two cases for the collinear vectors $\overrightarrow{x_i x}$ and $\overrightarrow{y_j y}$.

Let suppose that $\overrightarrow{x_i x}$ and $\overrightarrow{y_j y}$ have the same direction (left of Figure 2). 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, 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 is 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 now assume that $\overrightarrow{x_i x}$ and $\overrightarrow{y_j y}$ have opposite direction (right of Figure 2). 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, 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 $d(x, y) < d(x_i, x)$. Besides, $d(x_i, y_j) \leq 2 < d(y, y_j)$. So we arrive at the contradiction $d(x, y) + d(x_i, y_j) < d(x, x_i) + d(y, y_j)$. ◀

It is folklore that unit ball graphs have VC-dimension 4. 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 $H := G[N(v)]$. It can be easily shown that the neighborhood of this vertex (in fact, of any vertex) can be partitioned into, say 50 cliques. But actually 12 cliques are enough following the proof of the kissing number for spheres. Thus $\alpha(\overline{H}) \geq |V(H)|/12$. Therefore, from Theorem 1 and Theorem 3, we immediately obtain the following.

► **Theorem 4.** *There is a randomized EPTAS in time $2^{\tilde{O}(1/\varepsilon^3)} n^{O(1)}$ for MAXIMUM CLIQUE on unit ball graphs, even without the geometric representation.*

From Theorem 3, we also get an exact subexponential algorithm running in time $2^{\tilde{O}(n^{2/3})}$ for MAXIMUM CLIQUE on unit ball graphs. Indeed, such an algorithm was obtained for the class of graphs G such that $\text{iocp}(\overline{G}) \leq 1$ [10].

² or a plane if it is not unique

In sharp contrast, we will prove that with an extra dimension or with different radii (even arbitrarily close to each other), a PTAS is highly unlikely. In both cases, we show that there is a constant ratio which is not attainable even in subexponential time, unless the ETH fails. A *2-subdivision* of a simple graph is the graph obtained by subdividing each edge exactly twice. A *co-2-subdivision* is the complement graph of a 2-subdivision. The class of the 2-subdivisions (resp. co-2-subdivisions) is the set of all graphs that are the 2-subdivision (co-2-subdivision) of a simple graph.

The following result was shown by Bonnet *et al.*

► **Theorem 8** ([10]). *There is a constant $\alpha > 1$ such that MAXIMUM CLIQUE on co-2-subdivisions is not α -approximable even in time $2^{n^{0.99}}$, unless the ETH fails. Moreover, MAXIMUM CLIQUE is NP-hard on co-2-subdivisions.*

Therefore, 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 by a clique of n vertices representing the initial vertices, and a clique of $2m$ vertices minus a perfect matching; each anti-matched pair of vertices corresponding to an edge in the initial graph. Each vertex representing half an initial edge is adjacent to all the vertices representing the initial vertices but one.

► **Theorem 9.** *The class of 4-dimensional unit disks 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 call 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 $p^+(e)$ (resp. $p^-(e)$) fall regularly on the "top" part (resp. "bottom" part) of a circle \mathcal{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 \mathcal{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 \mathcal{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 notice that for any $(v, e, s) \in V \times E \times \{+, -\}$, $d(p(v), \pi(p^s(e))) = \sqrt{4 - 2\sqrt{3}\varepsilon + \varepsilon^2}$. We will choose $\varepsilon \ll 2 - \eta \ll 1$ so that 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 $\overrightarrow{\pi(p^+(e))p^+(e)} = -\frac{(\varepsilon + \varepsilon')}{\|Op(u)\|} \overrightarrow{Op(u)}$ and $p^-(e)$ such that $\overrightarrow{\pi(p^-(e))p^-(e)} = -\frac{(\varepsilon + \varepsilon')}{\|Op(v)\|} \overrightarrow{Op(v)}$. In words, we push very slightly $p^+(e)$ (resp. $p^-(e)$) away from $\pi(p^+(e))$ (resp. $\pi(p^-(e))$) in the opposite direction of $\overrightarrow{Op(u)}$ (resp. $\overrightarrow{Op(v)}$). This way, $p^+(e)$ is at distance more than 2 from $p(u)$. Indeed, $d(p^+(e), p(u)) =$



■ **Figure 3** The overall construction for 4-dimensional unit disks. We only represent the centers.

$\|\overrightarrow{p^+(e)\pi(p^+(e))} + \overrightarrow{\pi(p^+(e))O} + \overrightarrow{Op(u)}\| = \|\overrightarrow{\pi(p^+(e))O} + \overrightarrow{Op(u)} + \overrightarrow{p^+(e)\pi(p^+(e))}\| = \|\overrightarrow{\pi(p^+(e))O} + (\sqrt{3} - \varepsilon + \varepsilon + \varepsilon') \frac{\overrightarrow{Op(u)}}{\|Op(u)\|}\| = \sqrt{1 + (\sqrt{3} + \varepsilon')^2} > 2$. Similarly, $d(p^-(e), p(v)) > 2$. We choose ε' infinitesimal (in particular, $\varepsilon' \ll \varepsilon$) so that, still, $d(p^+(e), p(w)) < 2$ for any $w \neq u$ and $d(p^-(e), p(w)) < 2$ for any $w \neq v$. ◀

► **Corollary 10.** *There is a constant $\alpha > 1$ such that MAXIMUM CLIQUE on 4-dimensional unit disks is not α -approximable even in time $2^{n^{0.99}}$, unless the ETH fails. Moreover, MAXIMUM CLIQUE is NP-hard on 4-dimensional unit disks.*

The proof of Theorem 9 can be tweaked for balls of different radii. For a real $\varepsilon > 0$, we say that the radii of a representation (or the representation itself) are ε -close if the radii are all contained in the interval $[1, 1 + \varepsilon]$. We denote by $\mathcal{B}(1, 1 + \varepsilon)$ the ball graphs having an ε -close representation.

► **Theorem 11.** *For any $\varepsilon > 0$, the subclass of ball graphs $\mathcal{B}(1, 1 + \varepsilon)$ 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 \mathcal{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 $(0, 0, \sqrt{3} + i\varepsilon')$. The radius of the disk representing v_i centered at $p(v_i)$ is set to $r_i := \sqrt{1 + (\sqrt{3} + i\varepsilon')^2} - 1 + \varepsilon''$. The radii associated to the centers $p^+(e)$ and $p^-(e)$ are all set to 1. We move $p^+(v_i v_j)$ (resp. $p^-(v_i v_j)$) away from $\pi(p^+(v_i v_j))$ (resp. $\pi(p^-(v_i v_j))$) in the opposite direction of $\overrightarrow{p^+(v_i v_j)p(v_i)}$ (resp. $\overrightarrow{p^-(v_i v_j)p(v_j)}$) by an infinitesimal quantity, in order to only suppress the overlap of the disks centered at $p^+(v_i v_j)$ and $p(v_i)$ (resp. $p^-(v_i v_j)$ and $p(v_j)$). We make ε' and ε'' small enough that all the values r_i are between 1 and $1 + \varepsilon$. ◀

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

► **Corollary 12.** *There is a constant $\alpha > 1$ such that MAXIMUM CLIQUE on quasi unit ball graphs is not α -approximable even in time $2^{n^{0.99}}$, unless the ETH fails. Moreover, MAXIMUM CLIQUE is NP-hard on quasi unit ball graphs.*

References

- 1 P. Afshani and T. M. Chan. Approximation algorithms for maximum cliques in 3d unit-disk graphs. In *Proceedings of the 17th Canadian Conference on Computational Geometry, CCCG'05, University of Windsor, Ontario, Canada, August 10-12, 2005*, pages 19–22, 2005.
- 2 P. Afshani and H. Hatami. Approximation and inapproximability results for maximum clique of disc graphs in high dimensions. *Inf. Process. Lett.*, 105(3):83–87, 2008.
- 3 J. Alber and J. Fiala. Geometric separation and exact solutions for the parameterized independent set problem on disk graphs. *J. Algorithms*, 52(2):134–151, 2004.
- 4 N. Alon, R. Yuster, and U. Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.
- 5 C. Ambühl and U. Wagner. The Clique Problem in Intersection Graphs of Ellipses and Triangles. *Theory Comput. Syst.*, 38(3):279–292, 2005.
- 6 J. Bang-Jensen, B. Reed, M. Schacht, R. Šámal, B. Toft, and U. Wagner. On six problems posed by Jarik Nešetřil. *Topics in Discrete Mathematics*, pages 613–627, 2006.
- 7 C. Biró, É. Bonnet, D. Marx, T. Miltzow, and P. Rzażewski. Fine-grained complexity of coloring unit disks and balls. In *33rd International Symposium on Computational Geometry, SoCG 2017, July 4-7, 2017, Brisbane, Australia*, pages 18:1–18:16, 2017.
- 8 A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *J. ACM*, 36(4):929–965, 1989.
- 9 A. Bock, Y. Faenza, C. Moldenhauer, and A. J. Ruiz-Vargas. Solving the Stable Set Problem in Terms of the Odd Cycle Packing Number. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, pages 187–198, 2014.
- 10 É. Bonnet, P. Giannopoulos, E. J. Kim, P. Rzażewski, and F. Sikora. QPTAS and subexponential algorithm for maximum clique on disk. SoCG 2018 (to appear).
- 11 N. Bousquet, A. Lagoutte, Z. Li, A. Parreau, and S. Thomassé. Identifying codes in hereditary classes of graphs and vc-dimension. *SIAM J. Discrete Math.*, 29(4):2047–2064, 2015.
- 12 A. Brandstädt, V. B. Le, and J. P. Spinrad. *Graph classes: a survey*. SIAM, 1999.
- 13 H. Breu and D. G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Comput. Geom.*, 9(1-2):3–24, 1998.
- 14 S. Cabello. Maximum clique for disks of two sizes. Open problems from Geometric Intersection Graphs: Problems and Directions CG Week Workshop, Eindhoven, June 25, 2015 (<http://cgweek15.tcs.uj.edu.pl/problems.pdf>), 2015. [Online; accessed 07-December-2017].
- 15 S. Cabello. Open problems presented at the Algorithmic Graph Theory on the Adriatic Coast workshop, Koper, Slovenia (<https://conferences.matheo.si/event/6/picture/35.pdf>). June 16-19 2015.
- 16 T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003.
- 17 B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1-3):165–177, 1990.
- 18 L. Danzer. Zur lösung des gallaischen problems über kreisscheiben in der euklidischen ebene. *Studia Sci. Math. Hungar*, 21(1-2):111–134, 1986.
- 19 T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005.

- 20 A. V. Fishkin. Disk graphs: A short survey. In K. Jansen and R. Solis-Oba, editors, *Approximation and Online Algorithms, First International Workshop, WAOA 2003, Budapest, Hungary, September 16-18, 2003, Revised Papers*, volume 2909 of *Lecture Notes in Computer Science*, pages 260–264. Springer, 2003.
- 21 M. Gibson and I. A. Pirwani. Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier - (extended abstract). In *Algorithms - ESA 2010, 18th Annual European Symposium, Liverpool, UK, September 6-8, 2010. Proceedings, Part I*, pages 243–254, 2010.
- 22 S. Har-Peled, H. Kaplan, W. Mulzer, L. Roditty, P. Seiferth, M. Sharir, and M. Willert. Stabbing pairwise intersecting disks by five points. *CoRR*, abs/1801.03158, 2018.
- 23 D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. pages 61–71, 1986.
- 24 R. J. Kang and T. Müller. Sphere and Dot Product Representations of Graphs. *Discrete & Computational Geometry*, 47(3):548–568, 2012.
- 25 P. Koebe. *Kontaktprobleme der konformen Abbildung*. Hirzel, 1936.
- 26 D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 865–877, 2015.
- 27 T. A. McKee and F. R. McMorris. *Topics in intersection graph theory*. SIAM, 1999.
- 28 T. Nieberg and J. Hurink. A PTAS for the minimum dominating set problem in unit disk graphs. In *Approximation and Online Algorithms, Third International Workshop, WAOA 2005, Palma de Mallorca, Spain, October 6-7, 2005, Revised Papers*, pages 296–306, 2005.
- 29 T. Nieberg, J. Hurink, and W. Kern. A robust PTAS for maximum weight independent sets in unit disk graphs. In *Graph-Theoretic Concepts in Computer Science, 30th International Workshop, WG 2004, Bad Honnef, Germany, June 21-23, 2004, Revised Papers*, pages 214–221, 2004.
- 30 V. Raghavan and J. P. Spinrad. Robust algorithms for restricted domains. *J. Algorithms*, 48(1):160–172, 2003.
- 31 W. D. Smith and N. C. Wormald. Geometric separator theorems & applications. In *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, pages 232–243, 1998.
- 32 L. Stachó. A solution of gallai’s problem on pinning down circles. *Mat. Lapok*, 32(1-3):19–47, 1981.
- 33 E. J. van Leeuwen. Better approximation schemes for disk graphs. In *Algorithm Theory - SWAT 2006, 10th Scandinavian Workshop on Algorithm Theory, Riga, Latvia, July 6-8, 2006, Proceedings*, pages 316–327, 2006.
- 34 E. J. van Leeuwen. *Optimization and Approximation on Systems of Geometric Objects*. PhD thesis, Utrecht University, 2009.
- 35 V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pages 11–30. Springer, 2015.