



UvA-DARE (Digital Academic Repository)

Optimization and approximation on systems of geometric objects

van Leeuwen, E.J.

[Link to publication](#)

Citation for published version (APA):

van Leeuwen, E. J. (2009). Optimization and approximation on systems of geometric objects

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 4

Geometric Intersection Graphs and Their Representation

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

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

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

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

4.1 Scalable and ϵ -Separated Objects

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The same theorem holds for open scalable objects.

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

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

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

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

This fact is useful when proving the following corollary.

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

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

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

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

4.2 Finite Representation

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

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

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

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

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

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

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

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

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

The following theorem follows straightforwardly from Theorem 4.2.1.

Theorem 4.2.3 *Let \mathcal{A} be a family of closed scalable objects, each of rational size. Then any \mathcal{A} -intersection graph has a representation \mathcal{S} such that the*

distinguished point of each object in \mathcal{S} has rational coordinates and all objects in \mathcal{S} have rational size.

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

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

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

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

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

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

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

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

4.3 Polynomial Representation and Separation

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

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

Definition 4.3.1 *For a family \mathcal{A} of scalable objects, an \mathcal{A} -intersection graph G has a q -representation for some $q \geq 0$ if G has a representation \mathcal{S} such that*

the distinguished point of each object in \mathcal{S} has q -bit rational coordinates and each object has q -bit rational size.

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

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

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

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

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

4.3.1 From Representation to Separation

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

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

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

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

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

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

and thus u and v are δ -separated.

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

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

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

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

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

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

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

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

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

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

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

4.3.2 From Separation to Representation

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The theorem follows. \square

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

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

We can now prove the following result.

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

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

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

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

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

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

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