

International Journal of Computational Geometry & Applications
© World Scientific Publishing Company

A POLYNOMIAL-TIME APPROXIMATION ALGORITHM FOR A GEOMETRIC DISPERSION PROBLEM*

MARC BENKERT[†]

XP Software, 8-10 Purdue Str, Belconnen, Australia.
marc.benkert@cpsoftware.com

JOACHIM GUDMUNDSSON[‡]

National ICT Australia Ltd, Sydney, Australia.
joachim.gudmundsson@nicta.com.au

CHRISTIAN KNAUER

Institute of Computer Science, Freie Universität Berlin, Germany.
christian.knauer@inf.fu-berlin.de

RENÉ VAN OOSTRUM

ORTEC b.v., Gouda, The Netherlands.
Rene.vanOostrum@ortec.com

ALEXANDER WOLFF[†]

Faculteit Wiskunde en Informatica, Technische Universiteit Eindhoven, The Netherlands.
http://www.win.tue.nl/~awolff

Received December 24, 2006

Revised October 3, 2007

Communicated by D. T. Lee and Danny Chen, Guest Editor

We consider the following packing problem. Let α be a fixed real in $(0, 1]$. We are given a bounding rectangle ρ and a set \mathcal{R} of n possibly intersecting unit disks whose centers lie in ρ . The task is to pack a set \mathcal{B} of m disjoint disks of radius α into ρ such that no disk in \mathcal{B} intersects a disk in \mathcal{R} , where m is the maximum number of *unit* disks that can be packed. In this paper we present a polynomial-time algorithm for $\alpha = 2/3$. So far only the case of packing squares has been considered. For that case, Baur and Fekete have given a polynomial-time algorithm for $\alpha = 2/3$ and have shown that the problem cannot be solved in polynomial time for any $\alpha > 13/14$ unless $\mathcal{P} = \mathcal{NP}$.

Keywords: Geometric packing, approximation algorithm, APX-hard.

*Part of this work was done while M. Benkert and A. Wolff were affiliated with Karlsruhe University, while M. Benkert was visiting NICTA on a DAAD grant, and while R. van Oostrum was affiliated with Utrecht University.

[†]Supported by grant WO 758/4-2 of the German Research Foundation (DFG).

[‡]Funded by the Australian Government's Backing Australia's Ability initiative, in part through the Australian Research Council.



Fig. 1. The packing problem we consider in this paper occurs in non-photorealistic rendering, e.g., when using an oil painting style.

1. Introduction

The geometric dispersion problem we consider in this paper is about placing points in a restricted area of the plane such that the points are located as far away as possible from a set of points that is already present, and simultaneously as far away from each other as possible as well. This problem has applications in surveying, in non-photorealistic rendering, and in obnoxious facility location.

In surveying one may be interested in certain soil or water parameters. For a given domain, measurements of some parameter are known for a given set of locations, and one has the means to do extra measurements or set up a number of new measuring stations. It is desired that these new locations are nicely distributed over the area that has not yet been covered by the given fixed locations.

In non-photorealistic systems 3D models are to be rendered, e.g., in an oil painting style. For an example, see Figure 1. When a model is rendered “painterly”, instead of computing the color of each pixel by ray casting, the color is solely based on lighting and model properties. The idea is to generate a number of brush strokes. Each of these starts at a selected pixel that defines the color for all pixels that are covered by this brush stroke. The aim is to distribute the locations of the brush strokes more or less evenly.

In the general facility location problem a set of customers is given that is to be served from a set of facilities. The goal is to place a number of facilities in such a way that the distance to the closest facility is minimized over all customers. In other words, facilities are desirable, and customers like to be close to them. In obnoxious, or undesirable, facility location problems, the opposite is true: customers now consider it undesirable to be in the proximity of these facilities, and the goal




Fig. 2. (a) The input: a rectangle ρ and a set \mathcal{R} of (possibly overlapping) unit disks. (b) An optimal solution where $m = 6$ unit disks are packed in $\rho \setminus \bigcup \mathcal{R}$. (c) A solution with m disks of radius $3/4$.

is to maximize the distance to the closest facility over all customers. Examples of such undesirable facilities are, for instance, nuclear power plants or garbage dumps.

There are several models for and variants of the problem, see the surveys by Cappanera⁵ and Tóth,¹⁴ and the monographs by Rogers¹² and Zong and Talbot.¹⁶ The problem we will study in this paper is a special case of the following problem.

Problem 1 (PACKINGSCALEDTRANSLATES). Let $\alpha \in (0, 1]$ be a fixed real, let $S \subset \mathbb{R}^2$ be a shape of positive area, and let αS be the result of scaling S at the origin with a factor of α . Given an area $A \subset \mathbb{R}^2$, pack at least m disjoint translates of αS into A , where m is the maximum number of disjoint translates of S that can be packed into A .

Note that we do not know the value of m a priori. The special case that we study is referred to as the PACKINGSCALEDDISKS problem. In this case S is the unit disk and $A = \rho \setminus \bigcup \mathcal{R}$ is a given rectangle ρ minus a set \mathcal{R} of n given obstacle disks (see Figure 2). It is easy to see (and we will give the proof in Section 2) that greedily placing radius- α disks in A solves PACKINGSCALEDDISKS for $\alpha \leq 1/2$. The main contribution of this paper is a polynomial-time algorithm for $\alpha = 2/3$.

In their pioneering work, Hochbaum and Maass⁸ described a polynomial-time approximation scheme (PTAS) for the problem of packing the maximum number of translates of a fixed shape into a region given by the union of cells of the unit grid. The problem is known to be NP-hard.⁷

Even though the corresponding geometric dispersion problem looks very similar, inapproximability results have been shown. Baur and Fekete³ investigated the complexity of different variants of the geometric dispersion problem, one of which is a special case of PACKINGSCALEDTRANSLATES where S is the axis-parallel unit square and A is a rectilinear polygon with n vertices. We refer to this problem as PACKINGSCALEDBOXES. Baur and Fekete showed that it cannot be solved in polynomial time for $\alpha > 13/14$ unless $\mathcal{P} = \mathcal{NP}$.

They also gave a $2/3$ -approximation for PACKINGSCALEDBOXES. Their algorithm can be sketched as follows. First they compute a set S of at least $2m/3$ disjoint unit squares in the given polygon A , using the PTAS of Hochbaum and

Maass. Then they use the space occupied by the squares in S to pack at least m squares of side length $2/3$ in A . Their algorithm runs in $O(n^{40} \log m)$ time. If the user does not insist on an explicit list of all squares in the output, the running time becomes $O(n^{38})$, i.e., strongly polynomial.

The two main steps of our $2/3$ -approximation for PACKINGSCALEDISKS are conceptually the same as those of Baur and Fekete’s algorithm. To make it work for disks, however, we need a more refined analysis. First of all we have to adjust the PTAS of Hochbaum and Maass to packing disks. From this PTAS we need $8m/9$ unit disks. We also need a more involved geometric argument that guarantees sufficient space for the radius- $2/3$ disks. It is based on a matching in a nearest-neighbor graph of the unit disks. As the algorithm of Baur and Fekete, our algorithm is exclusively of theoretical interest—we roughly estimate its running time to be $O(n^{1620})$, see Section 4.

For completeness we mention that Ravi et al.¹¹ considered two *abstract* dispersion problems to which they refer as MAX-MIN and MAX-AVG. In those problems a set of n objects with their pairwise distances and a number $k < n$ is given, and the task is to pick k of the given objects such that the minimum (resp., average) distance among them is maximized. It is known that both versions are NP-hard. Ravi et al. showed that unless $\mathcal{P} = \mathcal{NP}$ there is no constant-factor approximation for MAX-MIN in the case of arbitrary distances. For the case that the distances fulfill the triangle inequality they gave factor-2 and factor-4 approximation algorithms for MAX-MIN and MAX-AVG, respectively. The former is optimal. Wang and Kuo¹⁵ investigated a variant of MAX-MIN where the objects are points in Euclidean space. They gave an $O(\max\{kn, n \log n\})$ -time algorithm for the one-dimensional case and showed that it is NP-hard to solve the two-dimensional case. For MAX-AVG, Ravi et al.¹¹ solved the one-dimensional case within the same time bound while the two-dimensional problem is still open.

Our paper is structured as follows. We show that greedily packing works for $\alpha \leq 1/2$ in the next section. We outline our algorithm for $\alpha = 2/3$ in Section 3. We give the details in Sections 4–8, and conclude in Section 9.

2. A Simple Greedy Strategy

Here we present a simple greedy approach that guarantees a $1/2$ -approximation. We use the term r -disk as shorthand for a disk of radius r . Consider an initially empty set \mathcal{B} . The greedy algorithm will iteratively add $1/2$ -disks to \mathcal{B} until no more disks can be added. More exactly, in each round a $1/2$ -disk D is added to \mathcal{B} if D lies entirely within ρ and does not intersect any of the disks in \mathcal{B} or \mathcal{R} .

Theorem 1. *For $\alpha \leq 1/2$ greedily packing region A solves PACKINGSCALEDISKS in $O((m+n) \log(m+n))$ time.*

Proof. The following simple charging argument shows that the greedy algorithm places at least m $1/2$ -disks. Consider an arbitrary (optimal) placement Π of m

disjoint unit disks in A . We charge a disk D in Π whenever the center of an α -disk falls into D . We claim that all disks in Π get charged at least once during the execution of the greedy algorithm. Assume this was not the case. Then we could place a α -disk concentrically into an uncharged unit disk. This would contradict the termination condition of the greedy algorithm. Thus each of the m disks in Π has in fact been charged. Since the placement of an α -disk causes the charging of at most one disk in Π , the greedy algorithm places at least m α -disks as stated.

The greedy algorithm can be implemented as follows. First, we compute the union of the disks in \mathcal{R} scaled by a factor of $(1 + \alpha)$. Second, we intersect this union with a copy of ρ shrunk by α in each direction. The resulting region—call it A' —contains the centers of all α -disks that lie in ρ and do not intersect any disk in \mathcal{R} . Then we pick any vertex v on the boundary of A' , add the α -disk centered at v to \mathcal{B} , subtract the (2α) -disk centered at v from A' , and repeat this process until A' is empty.

As Aurenhammer¹ has observed, (the boundary of) a union of disks can be represented via the *power diagram* of the disks. A power diagram is a generalization of a Voronoi diagram where the distance used is the *power distance* instead of the Euclidean distance. The power distance of a point p from a disk of radius r and center c is defined by $|pc|^2 - r^2$. Given s disks of arbitrary radii, their power diagram can be computed incrementally in $O(s^2)$ worst-case time.² Thus the greedy algorithm can be implemented to run in $O((\tilde{m} + n)^2)$ time, where $\tilde{m} \geq m$ is the number of disks in the output.

A faster approach is as follows: we first compute the boundary β of the scaled copies of the disks in \mathcal{R} in batch in $O(n \log n)$ time via the Voronoi diagram \mathcal{V} of their centers. (Note that all vertices of β lie on Voronoi edges and that each Voronoi edge contains at most one vertex of β .) Then we use the fact that the disks in \mathcal{B} are disjoint. We overlay ρ with a unit-square grid and maintain for each grid cell the constant number of 2α -arcs that lie on the current boundary of the union of all disks. Whenever we place a new 2α -disk, we use \mathcal{V} and the grid to search for all pieces of the old boundary that will be (partially) covered by the new disk. In this search we can afford traversing all pieces of the boundary that disappear since they will never appear again. The number of grid cells is $O(\tilde{m} + n)$. Thus this implementation runs in $O((\tilde{m} + n) \log(\tilde{m} + n))$ time.

The lemma follows by observing that $\tilde{m} = O(m + n)$. This holds again due to a simple charging argument: if we charge each disk placed by the greedy algorithm to the closest disk in the set $\Pi \cup \mathcal{R}$, then each of the $m + n$ disks in that set gets charged at most a constant number of times. \square

Clearly the greedy algorithm for $\alpha = 1/2$ works also for other shapes that are convex and point-symmetric. The running time for such shapes depends on how fast their unions can be computed incrementally.

3. Algorithm Outline

We now give a rough outline of our 2/3-approximation algorithm for the problem PACKINGSCALEDISKS, see Algorithm 1. We refer to our algorithm as DISKPACKER. For $r > 0$ and a set $R \subseteq \mathbb{R}^2$ let the r -freespace of R , denoted $\mathcal{F}_r(R)$, be the set of the centers of all r -disks that are completely contained in R . Let the extended r -freespace $\mathcal{F}_r^\otimes(R)$ be the Minkowski sum of $\mathcal{F}_r(R)$ and an r -disk.

Algorithm 1: DISKPACKER(ρ, \mathcal{R})

- 1 Compute the 1-freespace $\mathcal{F}_1 = \mathcal{F}_1(\rho \setminus \bigcup \mathcal{R})$ and the extended 1-freespace \mathcal{F}_1^\otimes .
 - 2 Use the PTAS of Hochbaum and Maass⁸ to compute a set of at least $8m/9$ disjoint unit disks in \mathcal{F}_1^\otimes .
 - 3 Greedily fill possible gaps in the remaining freespace with further unit disks.
 - 4 Let \mathcal{B} be the set of unit disks placed by PTAS and post-processing.
 - 5 Define a metric $\text{dist}(\cdot, \cdot)$ on \mathcal{B} .
 - 6 Compute the nearest-neighbor graph $G = (\mathcal{B}, \mathcal{E})$ with respect to dist .
 - 7 Find a matching in G .
 - 8 **for** each pair $\{D, D'\}$ of unit disks in the matching **do**
 - 9 | place three 2/3-disks in a region defined by D and D'
 - 10 **end**
 - 11 **for** each unmatched unit disk $D \in \mathcal{B}$ **do**
 - 12 | place one 2/3-disk in D
 - 13 **end**
 - 14 **return** $\mathcal{B}_{2/3}$, the set of all placed 2/3-disks (of cardinality at least m).
-

Recall that $A = \rho \setminus \bigcup \mathcal{R}$. We first compute the sets $\mathcal{F}_1 = \mathcal{F}_1(A)$ and $\mathcal{F}_1^\otimes = \mathcal{F}_1^\otimes(A)$, see Figure 3. Then, we apply the PTAS of Hochbaum and Maass⁸ to \mathcal{F}_1^\otimes . For any positive integer t , the PTAS places at least $(1 - 1/t)^2 \cdot m$ unit disks into \mathcal{F}_1^\otimes , where m is the maximum number of unit disks that can be packed into \mathcal{F}_1^\otimes . For more details, see Section 4. As we aim for a set of at least $8m/9$ unit disks we set $t = 18$ since this is the smallest integer for which $(1 - 1/t)^2$ exceeds $8/9$. After running the PTAS we greedily fill possible gaps in the remaining freespace with further unit disks. Let \mathcal{B} be the set of unit disks placed by PTAS and post-processing, and let $m' \geq 8m/9$ be the cardinality of \mathcal{B} .

Then we compute a set $\mathcal{B}_{2/3}$ of at least $9m'/8 \geq m$ disjoint 2/3-disks in A . We obtain $\mathcal{B}_{2/3}$ in two steps. First, we compute a matching in the nearest-neighbor graph $G = (\mathcal{B}, \mathcal{E})$ of \mathcal{B} (see Section 6) with respect to a metric $\text{dist}(\cdot, \cdot)$ that we will specify in Section 5. Second, we define a region for each pair of unit disks in the matching such that (a) we can place three 2/3-disks in each region (see Figure 4) and (b) the regions are pairwise disjoint. The main part of the paper (see Section 7) is proving that property (b) actually holds. For each unmatched unit disk we place a 2/3-disk such that their centers coincide, see Section 8. Next we describe each




Fig. 3. The 1-freespace \mathcal{F}_1^\otimes (light shaded) and a shortcut vv' (dashed) between the connected components C and C' of \mathcal{F}_1^\otimes .




Fig. 4. Packing three $2/3$ -disks (bold) in the region (shaded) spanned by a pair of unit disks (dotted).

step of Algorithm DISKPACKER in more detail.

4. Adjusting the PTAS of Hochbaum and Maass

The PTAS of Hochbaum and Maass⁸ packs a near-optimal number of squares into a grid-aligned rectilinear region. In this section we detail how to adjust their PTAS to our situation: we want to pack disks instead of squares, our region A is not rectilinear but a rectangle with holes that are unions of disks, and we do not have an underlying grid.

For packing (square) objects of size $k \times k$ into some given region R , the shifting technique of Hochbaum and Maass⁸ works as follows. Let t be a positive integer. Take a square grid of cell size $kt \times kt$. For each grid cell that intersects R , solve the corresponding local packing problem in $f(t)$ time optimally, where f is some function. Let the set of all packed objects be $L_{0,0}$. Then shift the grid by the vector (ki, kj) for $i = 0, \dots, t-1$ and $j = 0, \dots, t-1$. Each shifted grid $\Gamma_{i,j}$ yields a locally optimal solution $L_{i,j}$. Let OPT be an globally optimal solution and let $\text{OPT}_{i,j}$ be the set of all objects in OPT that do not intersect the boundaries of the grid cells in $\Gamma_{i,j}$. Note that a fixed object in OPT intersects a horizontal grid line for one out

of the t different vertical positions of the grid. The same object intersects a vertical grid line for one out of the t different horizontal positions of the grid. Thus the object does *not* intersect any grid line for $(t - 1)^2$ of the t^2 positions of the grid. In other words $\sum_{i,j} |\text{OPT}_{i,j}| = (t - 1)^2 \cdot |\text{OPT}|$. Now by the pigeon-hole principle there is a pair (i^*, j^*) such that OPT_{i^*,j^*} contains at least $(t - 1)^2/t^2 = (1 - 1/t)^2$ times as many objects as OPT. Clearly each locally optimal solution $L_{i,j}$ contains at least as many objects as $\text{OPT}_{i,j}$. Thus L_{i^*,j^*} is a $(1 - 1/t)^2$ -approximation of OPT. Hochbaum and Maass state that each local packing problem can be solved in $f(t) = \tilde{N}^{t^2}$ time, where \tilde{N} is the number of objects that can be packed in the corresponding $(kt \times kt)$ -square. This yields a total running time of $O(k^2 t^2 N^{t^2})$, where N is the number of unit squares needed to cover region R .

Now we detail how to apply the shifting technique to the problem of packing unit disks in the region $A = \rho \setminus \bigcup \mathcal{R}$. Note that a unit disk fits into a (2×2) -square. Thus if we choose $k = 2$, the above analysis concerning the approximation factor carries over. It remains to show how to solve the local packing problems optimally. So let S be a grid cell of size $(2t \times 2t)$ that intersects A and that may contain or intersect obstacle disks, i.e., disks in \mathcal{R} . Let OPT_S be an optimal solution for packing unit disks in S .

We first make OPT_S canonical as follows. We go through the disks in OPT_S from left to right. For each disk D we move D as far left as possible – either horizontally or following the boundary of obstacle disks or previously moved disks. If D happens to hit another disk D' such that their centers lie on the same horizontal line, we move D along the lower half of the boundary of D' until D hits another object or until D reaches the lowest point of D' . In the former case D has reached its destination, in the latter case we continue to move D horizontally until the next event occurs, which we treat as before. If D hits the left edge of S , we move D downwards along the edge until D hits either another disk or the bottom edge of S .

This process yields the following. Every disk D in the canonical version OPT'_S of OPT_S touches at least two other objects (disks or square edges) with its left side (including the south pole). These are the objects that stop D from moving further to the left, and they completely determine the position of D . Using this fact we now compute OPT'_S by generating *all* canonical solutions, each of them by incrementally adding unit disks from left to right.

Let n_S be the number of obstacle disks that intersect S . In each step of our incremental procedure we have at most $c = \max\{n_S + 1, 2t\}$ choices to place a unit disk such that its left side touches two previously placed objects (again including obstacles and edges of S). Figure 5 shows a set \mathcal{R} of red disks and a set $\mathcal{B}_{\text{curr}}$ of blue disks that have already been placed. The set of possible choices to place the next disk has been marked by dotted circles. Their centers correspond to locally leftmost points in $\mathcal{F}_1^{\text{curr}}(S)$, the 1-freespace in S with respect to the obstacle disks and the previously placed disks (see the shaded region in Figure 5). Let \tilde{N} be the number of disks in OPT_S . In 1910 Thue¹³ proved that the density of any arrangement of non-overlapping unit disks in the plane is at most $\pi/\sqrt{12}$, the ratio of the area of

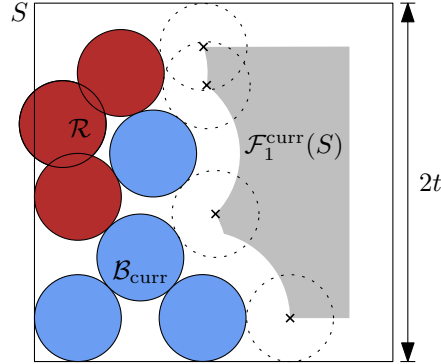


Fig. 5. Example situation for one step in the incremental procedure of the disk-placing PTAS.

the unit disk to the area of the circumscribed regular hexagon. Using Thue's result, we know that $\tilde{N} \leq (\pi/\sqrt{12}) \cdot (2t+4)^2 < 5t^2$ for $t \geq 12$. Thus there are at most c^{5t^2-1} choices, which can be enumerated within $O(c^{5t^2})$ time.

The total running time $T_{i,j}(n, t)$ for all local packing problems for a fixed position $\Gamma_{i,j}$ of the grid depends on the distribution of the obstacle disks. Our bound for $T_{i,j}(n, t)$ is maximized if $n_S = n$ for one subproblem S and $n_{S'} = 0$ for all other subproblems S' . This yields $T_{i,j}(n, t) = O(\max\{n+1, 2t\}^{5t^2})$. There are t^2 grid positions. Thus the total running time is $O(t^2 \cdot \max\{n+1, 2t\}^{5t^2})$. Recall from the previous section that we have to apply this algorithm for $t = 18$. Since we can assume $n > 2t$, the running time becomes $O(n^{1620})$.

If after applying the PTAS the remaining 1-freespace $\mathcal{F}_1^{\text{curr}}(\rho)$ of the whole instance is not empty, we greedily place further unit disks centered on points in $\mathcal{F}_1^{\text{curr}}(\rho)$ until $\mathcal{F}_1^{\text{curr}}(\rho)$ is empty. This is to make sure that small components (especially those which offer space for only one unit disk) are actually packed optimally. We will need this for our final counting argument, see Section 8. Let \mathcal{B} be the set of all disjoint unit disks that are placed by the PTAS and in the greedy post-processing phase. Let m' be the cardinality of \mathcal{B} .

5. The Freespace and a Metric on Unit Disks

We briefly recall the setting. We are given a set \mathcal{R} of n unit disks whose centers lie in a rectangle ρ , see Figure 2(a). The disks in \mathcal{R} are allowed to intersect. Let $A = \rho \setminus \bigcup \mathcal{R}$. We first compute the freespace $\mathcal{F}_1 = \mathcal{F}_1(A)$. According to Aurenhammer¹ the union of n disks can be computed in $O(n \log n)$ time and the complexity of its boundary is linear in n . We apply Aurenhammer's algorithm to the disks in \mathcal{R} scaled by a factor of 2. Then we intersect the resulting union with ρ shrunk by 1 unit. This yields \mathcal{F}_1 and \mathcal{F}_1^\otimes in $O(n \log n)$ time.

Next, we want to introduce a metric $\text{dist}(\cdot, \cdot)$ on unit disks in \mathcal{F}_1^\otimes . The idea of our algorithm is to use the connected components of \mathcal{F}_1 to identify all maximal regions

where we can place 2/3-disks. To guarantee that all such regions are discovered we need to join components of \mathcal{F}_1 that are not connected but still can hold 2/3-disks in the space between them. This is the idea behind the following definition.

Definition 1. Let C and C' be two connected components of \mathcal{F}_1 , and let v and v' be vertices on the boundaries of C and C' , respectively. We say that the straight-line segment vv' is a *shortcut* if $|vv'| \leq 2/3 \cdot \sqrt{11} \approx 2.21$, where $|pq|$ denotes the Euclidean distance of points p and q . Let $\mathcal{S}(C, C')$ be the set of all shortcuts induced by C and C' . We set

$$\mathcal{F}_1^+ = \mathcal{F}_1 \cup \bigcup_{\substack{C, C' \in \mathcal{F}_1 \\ s \in \mathcal{S}(C, C')}} s.$$

Figure 3 depicts \mathcal{F}_1 , \mathcal{F}_1^\otimes , and a shortcut vv' . Throughout the paper we will use upper-case letters to denote disks and the corresponding lower-case letters to denote their centers. Now, we are ready to define our metric for a connected component of \mathcal{F}_1^+ , see Figure 6.

Definition 2. Let D and D' be unit disks that lie in \mathcal{F}_1^\otimes . Let d and d' be their respective centers. The distance $\text{dist}(D, D')$ of D and D' is the length of the geodesic $g(d, d')$ of d and d' with respect to \mathcal{F}_1^+ . The *tunnel* $T(D, D')$ of D and D' is the union of all points in A within distance 1 of a point on $g(d, d')$.

From the definition of \mathcal{F}_1^\otimes it is easy to see that any 2/3-disk $D_{2/3}$ centered at a point of $g(d, d')$ does not intersect any disk in \mathcal{R} . (This will also follow from Lemma 2.) Thus $D_{2/3}$ is contained in the tunnel $T(D, D')$. Since \mathcal{R} is the union of a set of unit disks the geodesic between two points in \mathcal{F}_1^+ can only consist of line segments and arcs of radius 2, see Figure 6(b).

Recall that our algorithm computes a matching in the nearest-neighbor graph $G = (\mathcal{B}, \mathcal{E})$ induced by the metric $\text{dist}(\cdot, \cdot)$ on the set \mathcal{B} of unit disks that we get from the PTAS by Hochbaum and Maass. For each pair $\{D, D'\}$ of disks in the matching we now define a region $\mathcal{A}(D, D')$ into which we later place three 2/3-disks as in Figure 4. An obvious way to define this region would be to take the union of all 2/3-disks centered at points of the geodesic between d and d' in $\mathcal{F}_{2/3}$. Our definition is slightly more involved (for illustration see Figure 6), but it will simplify the proof of the main theorem in Section 6. The theorem states that two such regions are disjoint if the corresponding unit disks are pairwise disjoint. This makes sure that the 2/3-disks which we place into the regions are disjoint.

Definition 3. Let D and D' be unit disks in \mathcal{F}_1^\otimes . Let $g_{2/3}(d, d')$ be a geodesic from d to d' in $\mathcal{F}_{2/3}(T(D, D'))$. Then the *2/3-tunnel* $T_{2/3}(D, D')$ of D and D' consists of all points in A within distance 2/3 of a point on $g(d, d')$. Finally let the *placement region* $\mathcal{A}(D, D')$ of D and D' be $D \cup D' \cup T_{2/3}(D, D')$.

According to Chang et al.⁶ the geodesics $g(d, d')$ and $g_{2/3}(d, d')$ from d to d' can be computed in $O(n^2 \log n)$ time.




Fig. 6. The geodesic $g(d, d')$ (a) in the unrestricted case, (b) in the presence of obstacles, and (c) in case of a shortcut.

6. The Nearest-Neighbor Graph

Recall that m is the maximum number of disjoint unit disks that fit in \mathcal{F}_1^\otimes . From (our variant of) the PTAS of Hochbaum and Maass⁸ we get a set \mathcal{B} of $m' \geq 8m/9$ disjoint unit disks in \mathcal{F}_1^\otimes . Then we compute the nearest-neighbor graph $G = (\mathcal{B}, \mathcal{E})$ induced by the metric $\text{dist}(\cdot, \cdot)$. We consider G a directed graph, where an edge (C, D) is in G if D is the nearest neighbor of C , for $C, D \in \mathcal{B}$. In case of a tie, we pick any of the nearest neighbors of C , so every vertex in G has exactly one outgoing edge.

As for the algorithm we next find a matching in G and place three $2/3$ -disks in the placement region $\mathcal{A}(C, D)$ for each pair $\{C, D\}$ in the matching. Finally we place a $2/3$ -disk for each unmatched disk in \mathcal{B} . In Section 8 we show that in this way we place at least $9m'/8 \geq m$ disks of radius $2/3$ in total.

For nearest neighbor graphs of points in Euclidean space it is well-known that the maximum degree of the graph is bounded by 6. However, our setting is quite different so for completeness we include the following lemma.

Lemma 1. *The nearest-neighbor graph $G = (\mathcal{B}, \mathcal{E})$ with respect to dist is plane and has maximum degree 7.*

Proof. The following standard argument shows that G is plane. Assume that (C, D) and (E, F) are in G and that $g(c, d)$ and $g(e, f)$ intersect. Let p be one of the intersection points of $g(c, d)$ and $g(e, f)$. This implies that the lengths of $g(c, p)$, $g(d, p)$, $g(e, p)$ and $g(f, p)$ must all be equal. Then, however, either $g(c, d)$ or $g(e, f)$ cannot be a geodesic—a contradiction to (C, D) or (E, F) being in G .




Fig. 7. Illustration for the proof of Lemma 1.

It remains to show the degree bound. Let C be an arbitrary unit disk in \mathcal{B} . Since by construction of G the outdegree of each vertex is 1, it suffices to bound the indegree of C by 6.

Let $D, D' \in \mathcal{B}$ be two disks for which (D, C) and (D', C) are in \mathcal{E} . Consider the geodesics $g(c, d)$ and $g(c, d')$. By construction they end in C with a straight-line segment, see Figure 7. Now, it is enough to show that the angle α between $g(c, d)$ and $g(c, d')$ in C is larger than $360^\circ/7$. We analyze the setting for which α is minimized. Note that if D and D' touched C , then we would immediately have $\alpha \geq 60^\circ$ since D and D' are disjoint. To minimize α , the disks D and D' must lie at some distance from C , but such that the inequalities $\text{dist}(C, D) \leq \text{dist}(D, D')$ and $\text{dist}(C, D') \leq \text{dist}(D, D')$ are still obeyed. Hence there must be some disk in \mathcal{R} that lies in the wedge-like region between the geodesics $g(c, d)$ and $g(c, d')$. Among these disks let R be the one closest to C .

Due to the way the disks in \mathcal{B} have been placed, R cannot be very far from C , otherwise there would be a disk in \mathcal{B} closer to D or D' than C . For the same reason, there must be other disks in \mathcal{R} on the opposite sides of $g(c, d)$ and $g(c, d')$. The disks in \mathcal{R} must leave a corridor of width at least $2 \cdot 2/3$, otherwise there would not be a shortcut between them that the geodesics can use. This yields $|rc| < 11/3$. If the distance were larger, the presence of an additional disk in \mathcal{B} would immediately contradict (D, C) and (D', C) being in \mathcal{E} , see the indicated dotted disks in Figure 7. By construction the distance of any point on the geodesics $g(c, d)$ and $g(c, d')$ to r is at least $5/3$. Now some simple trigonometry yields $\alpha > 52.2^\circ$, which is greater than $360^\circ/7 \approx 51.4^\circ$. If we repeat the above construction we can place at most six disks D, D', D'', \dots such that C is the nearest neighbor of all of them. \square

From now on we call $\{C, D\} \subseteq \mathcal{B}$ a *nearest pair* if (C, D) or (D, C) is an edge in G , i.e., if D is closest to C or C is closest to D (among the disks in \mathcal{B}). Recall that the placement region $\mathcal{A}(C, D)$ of C and D was defined as $C \cup D \cup T_{2/3}(C, D)$. As

a nearest pair $\{C, D\}$ can be in the matching, we have to prove the following two statements:

- (i) three $2/3$ -disks fit into $\mathcal{A}(C, D)$ and
- (ii) for any nearest pair $\{E, F\}$, where $C, D, E,$ and F are pairwise disjoint, it holds that $\mathcal{A}(C, D) \cap \mathcal{A}(E, F) = \emptyset$.

Note that we do not have to care whether $\mathcal{A}(C, D)$ intersects $\mathcal{A}(C, E)$ because a matching in G contains at most one pair out of $\{C, D\}$ and $\{C, E\}$. Three $2/3$ -disks obviously fit into $\mathcal{A}(C, D)$ since C and D do not intersect, see Figure 4. Thus, statement (i) is true. The remaining part of the paper will focus on proving statement (ii).

7. Placement Regions of Nearest Pairs are Disjoint

By the definition of shortcuts, unit disks whose centers lie in different components of \mathcal{F}_1^+ do not intersect. This immediately yields that $\mathcal{A}(C, D) \cap \mathcal{A}(E, F) = \emptyset$ holds for $\{C, D\}$ and $\{E, F\}$ being nearest pairs from different components. Thus it remains to show that $\mathcal{A}(C, D) \cap \mathcal{A}(E, F) = \emptyset$ for nearest pairs $\{C, D\}$ and $\{E, F\}$ that lie in the same component.

We split the proof into two parts. The first part (Lemma 4) shows that $T_{2/3}(C, D)$ does not intersect any disk other than C and D . The second part (Theorem 2) shows that no two $2/3$ -tunnels $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ intersect. We start with two technical lemmas that we need to prove the first part.

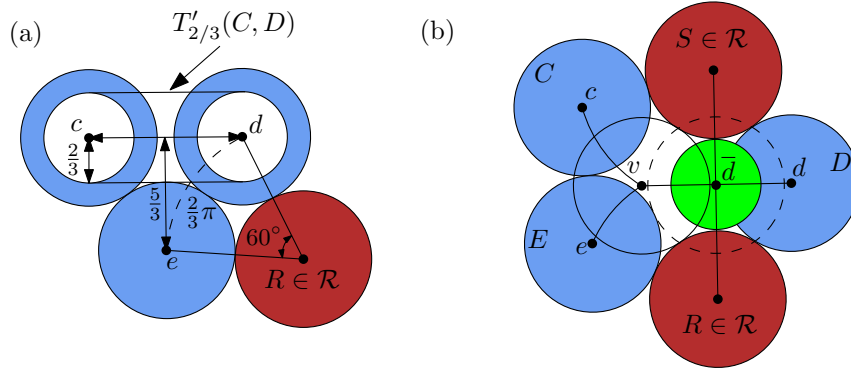


Fig. 8. Illustrations for (a) the proofs of Lemmas 2-3 and (b) case 2 in the proof of Lemma 4.

Lemma 2. *Let C and D be two unit disks in \mathcal{F}_1^\otimes . If $|cd| \leq \frac{2}{3}\sqrt{11}$ then $g_{2/3}(c, d)$ is a straight-line segment.*

Proof. Let $T'_{2/3}(C, D)$ be the Minkowski sum of a $2/3$ -disk and the line segment cd , see Figure 8(a). If $g_{2/3}(c, d)$ is not a line segment, then a disk E in $\mathcal{B} \cup \mathcal{R}$ intersects

$T'_{2/3}(C, D)$. We establish a lower bound on $|cd|$ for this to happen. Note that C , D and E are pairwise disjoint as C and D are disks in \mathcal{B} .

Clearly, the minimum distance between c and d is attained if E and $T'_{2/3}(C, D)$ only intersect in a single point and furthermore, both E and C as well as E and D intersect in a single point. This means that $|ce| = |de| = 2$. Moreover, the Euclidean distance between e and the straight-line segment cd is $1 + \frac{2}{3} = \frac{5}{3}$. By Pythagoras' theorem we calculate $|cd|$ to be at least $\frac{2}{3}\sqrt{11}$. This means that $T'_{2/3}(C, D)$ is contained in $A = \rho \setminus \bigcup \mathcal{R}$.

If C and D belong to different components of \mathcal{F}_1 , they must be connected via a shortcut according to Definition 1. Thus, $g_{2/3}(c, d)$ is a line segment. \square

Lemma 3. *Let D and E be two unit disks in \mathcal{F}_1^\otimes that touch each other. Then $\text{dist}(D, E) \leq \frac{2}{3}\pi$.*

Proof. Let D and E be two unit disks in \mathcal{F}_1^\otimes that touch each other, as illustrated in Figure 8(a). The length of the curve $g(D, E)$ is maximized if there is an obstacle disk R that touches D and E and no shortcut could be taken. In this case $g(D, E)$ describes a circular arc of radius 2 spanning 60° , thus its length is $\frac{1}{6} \cdot 2 \cdot 2\pi = \frac{2}{3}\pi$. \square

Now, we are ready to prove the first part:

Lemma 4. *Let $\{C, D\} \subseteq \mathcal{B}$ be a nearest pair. No disk of $\mathcal{B} \cup \mathcal{R} \setminus \{C, D\}$ intersects $T_{2/3}(C, D)$.*

Proof. From the definition of freespace and Definitions 2 and 3 it immediately follows that neither $T(C, D)$ nor $T_{2/3}(C, D)$ are intersected by a disk in \mathcal{R} . Thus, it remains to prove that apart from C and D no disk in \mathcal{B} intersects $T_{2/3}(C, D)$.

Without loss of generality, assume that C is the disk in \mathcal{R} closest to D . The proof is by contradiction, i.e., we assume that there is a disk $E \in \mathcal{B}$ that intersects $T_{2/3}(C, D)$.

First, we move a unit disk on $g(C, D)$ from the position of D to the first position in which it hits E . Denote the disk in this position by \overline{D} . We claim that the length of $g(\overline{d}, e)$, within \mathcal{F}_1^+ , is shorter than $g(c, \overline{d})$. This obviously contradicts C being the nearest neighbor of D , and would thereby complete the proof of the lemma.

We have to consider two cases for the upper bound on the length of $g(\overline{d}, e)$.

Case 1: If \overline{d} is in \mathcal{F}_1 and there is an obstacle disk $R \in \mathcal{R}$ that touches \overline{D} and E , then the length of $g(\overline{d}, e)$ is maximized and $g(\overline{d}, e)$ is an arc of radius 2 and spanning 60° . Lemma 3 yields $g(\overline{d}, e) \leq \frac{2}{3}$. It might be that $g(\overline{d}, e)$ does not lie entirely within \mathcal{F}_1 . However, as $\overline{D}, E \subseteq \mathcal{F}_1^\otimes$, there must then be a shortcut that would shorten the length of $g(\overline{d}, e)$ even further.

Next we give a lower bound on the length of $g(c, \overline{d})$. Since E touches $T_{2/3}(C, D)$ and \overline{D} it follows that C and \overline{D} are disjoint, otherwise E could not intersect $T_{2/3}(C, D)$. Consequently C, \overline{D} and E are pairwise disjoint and, according to

Lemma 2, the Euclidean distance between c and \bar{d} is greater than $\frac{2}{3}\sqrt{11}$. Putting the two bounds together we get:

$$g(\bar{d}, e) \leq \frac{2}{3} < \frac{2}{3}\sqrt{11} < g(c, \bar{d}).$$

Case 2: If \bar{d} is not in \mathcal{F}_1 then \bar{d} must lie on a shortcut vv' and the unit disk \bar{D} intersects at least one disk in \mathcal{R} . Let v be the endpoint of the shortcut (v, v') closest to E and let $\bar{D}_{2/3}$ be the $2/3$ -disk centered at \bar{d} . Note that v must lie in the same component as e , thus $g(\bar{d}, e)$ consists of a straight-line segment from \bar{d} to v followed by the geodesic $g(v, e)$. The length of $g(\bar{d}, e)$ is maximized if the angle $\angle v\bar{d}e$ is maximized. This is the case if a unit disk $R \in \mathcal{R}$ touches E and $\bar{D}_{2/3}$ and a unit disk $S \in \mathcal{R}$ touches $\bar{D}_{2/3}$ such that \bar{d}, r and s are collinear, as shown in Figure 8(b). By parametrization it follows that the length of $g(\bar{d}, e)$ is maximized if the length of $\bar{d}v$ is maximized, which is bounded by $1/3\sqrt{11}$, according to Lemma 2. By Pythagoras' theorem we can now compute the coordinates of e , they are $\approx (-1.8182, -0.8333)$, where the coordinate system is fixed by $\bar{d} = (0, 0)$ and $r = (0, -5/3)$. The geodesic $g(v, e)$ consists of an arc of radius 2, applying the cosine theorem then yields that the length of $g(v, e)$ is approximately 1.1105, which gives that $g(\bar{d}, e) < \frac{1}{3}\sqrt{11} + 1.105$.

Next we need a lower bound on the length of $g(c, \bar{d})$. Using the same ideas as above, the position of c can be computed by Pythagoras' theorem to be $\approx (-1.9356, 1.1632)$ and the length of the geodesic $g(\bar{c}, v)$ is at least 1.4607 using the cosine theorem. Thus,

$$g(\bar{d}, e) \leq \frac{1}{3}\sqrt{11} + 1.105 < \frac{1}{3}\sqrt{11} + 1.4607 < g(c, \bar{d}).$$

Since $g(\bar{d}, e)$ has been shown to be shorter than $g(c, \bar{d})$, in all cases, \bar{d} must be closer to c than to e , which is a contradiction to the initial assumption. This completes the proof of the lemma.

Note that the disks involved in this construction can be moved such that the lengths of $g(\bar{d}, e)$ and $g(c, \bar{d})$ changes. However, the above construction minimizes their difference. \square

Lemma 4 proves that no other disks apart from C and D intersect $T_{2/3}(C, D)$. It remains to prove that no two $\frac{2}{3}$ -tunnels $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ intersect.

Theorem 2. *Let $\{C, D\}, \{E, F\} \subseteq \mathcal{B}$ be two nearest pairs such that C, D, E and F are pairwise disjoint, it holds that $T_{2/3}(C, D) \cap T_{2/3}(E, F) = \emptyset$.*

Proof. The proof is by contradiction again. We assume that $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ intersect and show that this would either contradict $\{C, D\}$ or $\{E, F\}$ being nearest neighbors. Obviously it is enough to exclude the case that $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ intersect in a single point. We first characterize such an instance which helps us to conduct the contradiction proof.

For the sake of completeness we first have to exclude the case that $g(c, d)$ and $g(e, f)$ use the same shortcut: if they did, the geodesics $g_{2/3}(c, d)$ and $g_{2/3}(e, f)$

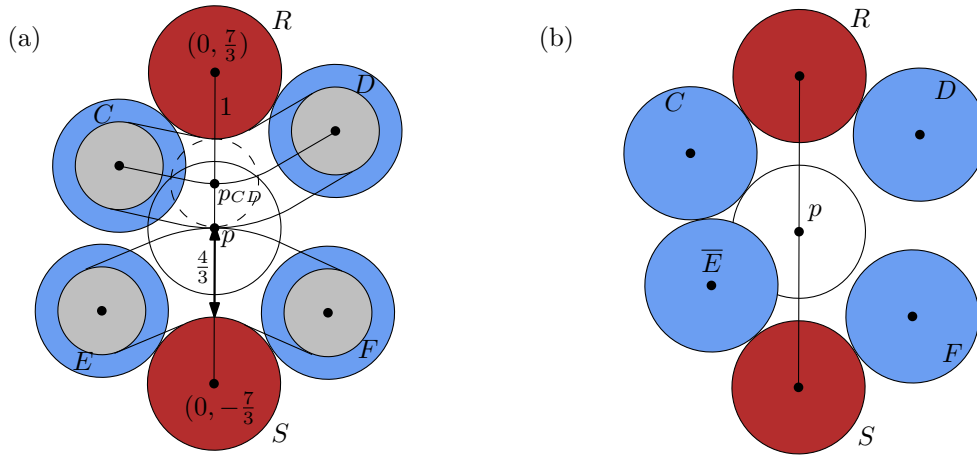


Fig. 9. Illustrating the proof of (a) Theorem 2 and (b) case (ii) in the proof of Theorem 2.

would intersect which immediately yields that $g(c, d)$ and $g(e, f)$ would intersect—a contradiction to Lemma 1.

Thus, we can assume that the intersection point p of $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ lies in \mathcal{F}_1 and neither $g(c, d)$ nor $g(e, f)$ takes a shortcut containing p . This in turn means that we can assume that no shortcut is taken at all. We observe that at least one of the disks $\{C, D, E, F\}$ intersects the unit disk τ with center p ; otherwise there would be another disk in \mathcal{B} located in the space between C, D, E and F which would immediately contradict $\{C, D\}$ as well as $\{E, F\}$ being nearest pairs. Without loss of generality, let C be a disk that intersects τ .

Let p_{CD} be the point on $g_{2/3}(C, D)$ such that $|pp_{CD}| = 2/3$, see Figure 9(a). Define p_{EF} similarly. We will assume that there is a vicinity of p_{CD} and p_{EF} in which $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ are arcs. The case when one vicinity of p_{CD} and p_{EF} is a straight line is easier and can be handled using similar arguments.

The curvature of $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ in a vicinity of p_{CD} and p_{EF} induces the existence of two disks $R, S \in \mathcal{R}$ as illustrated in Figure 9(a). Since R and S forces the curvature of $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ we may introduce the following coordinate system. The origin is p and the coordinates of r and s are $(0, \frac{7}{3})$ and $(0, -\frac{7}{3})$, respectively.

As a consequence of Lemma 2 we get that $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ start with a straight-line segment of length at least $\frac{1}{3}\sqrt{11}$, again see Figure 9(a). Thus, the curvature of $g_{2/3}(C, D)$ in p_{CD} infers that $|cp_{CD}| \geq \frac{1}{3}\sqrt{11}$ holds, which means that C either lies completely to the left of the y -axis or to the right. This holds analogously for the other disks. Without loss of generality, we assume that C and E lie to the left of the y -axis and D and F lie to the right, see Figure 9(a).

Note that we have to take into account the exact relationship behind the pairs

$\{C, D\}$ and $\{E, F\}$ being nearest pairs, e.g., C could be the nearest neighbor of D , or D could be the nearest neighbor of C . We will prove the following:

- (i) $\text{dist}(C, E) < \text{dist}(E, F)$
- (ii) $\text{dist}(C, E) < \text{dist}(C, D)$
- (iii) $\text{dist}(D, F) < \text{dist}(C, D)$

Item (i) says that C is closer to E than F is. Thus, in order for $\{E, F\}$ to be a nearest pair, E must be the nearest neighbor of F . We use this fact to show that (ii) and (iii) hold. Together, (ii) and (iii) comprise the contradiction: (ii) says that D is not the nearest neighbor of C , while (iii) says that C is not the nearest neighbor of D . Hence, $\{C, D\}$ cannot be a nearest pair.

(i): To prove that $\text{dist}(C, E) < \text{dist}(E, F)$ we will argue that $T(E, F)$ intersects C , i.e., there is a unit disk \overline{E} whose center lies on $g(E, F)$ that intersects C and not F . Let \overline{E} be defined by the left and bottommost point \overline{e} on $g(E, F)$ such that \overline{E} intersects C . This is illustrated in Figure 9(b). The proof of (i) can then be completed by showing that $\text{dist}(C, \overline{E}) < \text{dist}(\overline{E}, F)$.

First we prove that there exists a position of \overline{e} such that \overline{E} intersects C , i.e., a unit disk cannot pass between C and S without intersecting C . This could only be achieved by maximizing $|cs|$. Recall that C intersects τ , thus, $|cs|$ is maximized if C touches R and τ , i.e., C takes its left and topmost position, as shown in Figure 9(b). Using Pythagoras' theorem we can compute the coordinates of c for this setting to be $\approx (-1.62, 1.17)$. From now on we will omit the sign \approx when stating results of the calculations. Hence, it holds that $|cs| \leq 3.86$ which in turn yields that no unit disk can pass between C and S since this would require $|cs| \geq 4$.

Next, we minimize the distance $\text{dist}(\overline{E}, F)$ in order to get F to be closer to \overline{E} than to C . For this, \overline{E} should take its rightmost position touching C . This position is attained if C is as far as possible from S , i.e., \overline{E} takes position $(-1.62, 1.17)$ again. Using Pythagoras' theorem, the coordinates of \overline{e} is $(-1.29, -0.80)$. This means that $|\overline{e}p_{EF}| \geq 1.29$ and thus $\text{dist}(\overline{E}, F) \geq 1.29 + \frac{1}{3}\sqrt{11}$ as $|p_{EFF}| \geq \frac{1}{3}\sqrt{11}$ holds. According to Lemma 3 $\text{dist}(C, \overline{E}) \leq \frac{2}{3}\pi$, and we have:

$$\text{dist}(C, \overline{E}) \leq \frac{2}{3}\pi < 1.29 + \frac{1}{3}\sqrt{11} \leq \text{dist}(\overline{E}, F),$$

which concludes (i).

(ii): To prove that $\text{dist}(C, E) < \text{dist}(C, D)$ always holds we first establish a lower bound on $\text{dist}(C, D)$ and then an upper bound on $\text{dist}(C, E)$. For the lower bound on we try to push C and D as close as possible together under the restriction that E can still be the nearest neighbor of F . To minimize $\text{dist}(C, D)$, C should take its right and bottommost position and D should take its left and bottommost position.

For the bound on D we only use the fact that D is not allowed to intersect the tunnel $T(E, F)$. If it does, we would get $\text{dist}(D, E) < \text{dist}(E, F)$ by a similar argument as in (i). (Here, the corresponding point \overline{f} can even lie further to the right than $(1.29, -0.80)$ as D does not have to intersect τ .) However, $\text{dist}(D, E) <$

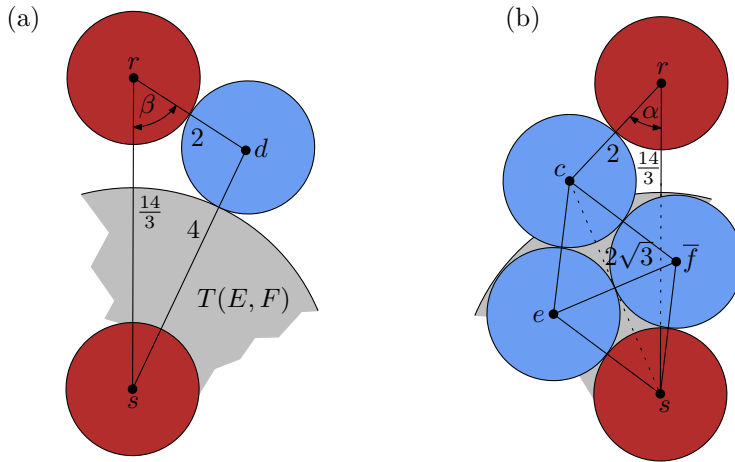


Fig. 10. Illustration of the proof of the lower bound on $|cd|$ in case (ii).

$\text{dist}(E, F)$ together with (i) would immediately contradict $\{E, F\}$ to be a nearest pair. Disk D takes its left and bottommost position without intersecting $T(E, F)$ if D touches R and is infinitesimal close to $T(E, F)$. For simplicity we assume that D touches $T(E, F)$, see Figure 10(a). Standard trigonometric calculations give the left and bottommost coordinates of d to be $(1.70, 1.29)$.

For the right and bottommost position of C we use the following arguments. Let \bar{f} be the rightmost point on $g(E, F)$ such that \bar{F} touches either C or E . We use that E has to be touched by C otherwise C is closer to F than E . We compute the right and bottommost position of C if \bar{F} touches C and E at the same time, see Figure 10(b). Note, that this actually yields a position in which C is closer to F than E (with respect to our metric dist). Again, standard trigonometry gives that the right and bottommost coordinates of c is $(-1.35, 0.86)$.

Now a lower bound on $|cd|$ is the Euclidean distance between $(1.70, 1.29)$ and $(-1.35, 0.86)$ which is 3.08. We obtain the final lower bound on $\text{dist}(C, D)$ by noting that both C and D touch R , if not the Euclidean distance between c and d could be shortened, hence the geodesic has to follow the circular arc around R , thus we get $\text{dist}(C, D) > 3.49$.

By Pythagoras' theorem we can also compute the coordinates of \bar{f} to be $(0.25, -0.35)$ – we will need them in the proof of (iii).

To prove (ii) it remains to show an upper bound on $\text{dist}(C, E)$ which is less than 3.49. We try to push C and E as far away from each other as possible, under the restriction that E is still the nearest neighbor of F . It is clear that C has to take its left and topmost position, which is already known as $(-1.62, 1.17)$, from (i), while E should take its left and bottommost position. Again, we only use that on the rightmost point \bar{f} on $g(E, F)$ such that either C or E is touched by \bar{F} , it must be E that is touched. First, we compute the rightmost point \bar{f} on $g(E, F)$ where \bar{F} touches

C taking position $(-1.62, 1.17)$ and touches S . As we know the coordinates of c and s , we can compute the coordinates of \bar{f} by Pythagoras' theorem. It holds that \bar{f} is $(-0.33, -0.36)$, see Figure 11(a). Now, E takes its left and bottommost position if it touches S and \bar{F} . Using Pythagoras' theorem again, we get that $e = (-1.87, -1.64)$. This yields the upper bound on $|ce|$ of 2.82. However, $g(C, E)$ may have to curve around S so we get $\text{dist}(C, E) < \pi$. Putting it together we get:

$$\text{dist}(C, E) < \pi < 3.49 < \text{dist}(C, D),$$

and we are done with (ii).

(iii): We use the lower bound on $\text{dist}(C, D)$ that was derived in (ii). Thus, we only have to show an upper bound on $\text{dist}(D, F)$ which is less than 3.49. For the upper bound we try to push D and F as far away from each other as possible, under the restriction that E can still be the nearest neighbor of F . For this, D has to take its right and topmost position while F has to take its right and bottommost position. We can assume that D takes position $(1.70, 1.29)$, the position of D which was responsible for the lower bound on $\text{dist}(C, D)$. This assumption is justified since, if D does not take position $(1.70, 1.29)$, we move D on $g(C, D)$ to this position, say \bar{D} , and show that $\text{dist}(\bar{D}, F) < 3.49$. Then, $\text{dist}(D, F) < \text{dist}(C, D)$ also holds since $\text{dist}(C, \bar{D}) > 3.49$.

To maximize $\text{dist}(D, F)$ we need F to take its bottommost position. Consider the disk P' that touches C and S and lies to the right of cs , see Figure 11(b). For the right and bottommost position of C , which is $(-1.35, 0.86)$, we already computed the position of P' to be $(0.25, -0.35)$. This implies that D does not intersect P' as $|dp'| > 2$ (recall that the position of D was decided in the previous section). We have shown that C , D and E do not intersect P' , however, then F has to intersect P' otherwise there would be another disk in \mathcal{B} located in the space between C , D , E and F which would immediately contradict $\{C, D\}$ as well as $\{E, F\}$ being nearest pairs. The disk F now takes its right and bottommost position if it touches P' and S . Using Pythagoras' theorem the position of F is shown to be $(1.84, -1.55)$.

As a result we get $|df| < 2.84$ but since $g(D, F)$ may have to curve around some other disk we get $\text{dist}(D, F) < 3.18$. Putting it together we get:

$$\text{dist}(D, F) < 3.18 < 3.49 < \text{dist}(C, \bar{D}) \leq \text{dist}(C, D),$$

and we are done with (iii) and, hence also the theorem. \square

8. Placing the 2/3-Disks

In the previous sections we have detailed how we compute a set \mathcal{B} of $m' \geq 8m/9$ unit disks in A and how we determine the nearest-neighbor graph $G = (\mathcal{B}, \mathcal{E})$ with respect to the metric $\text{dist}(\cdot, \cdot)$. In order to finally place the 2/3-disks we compute a matching in G . Observe that G can consist of more than one connected component. We consider each connected component separately. We start with components of

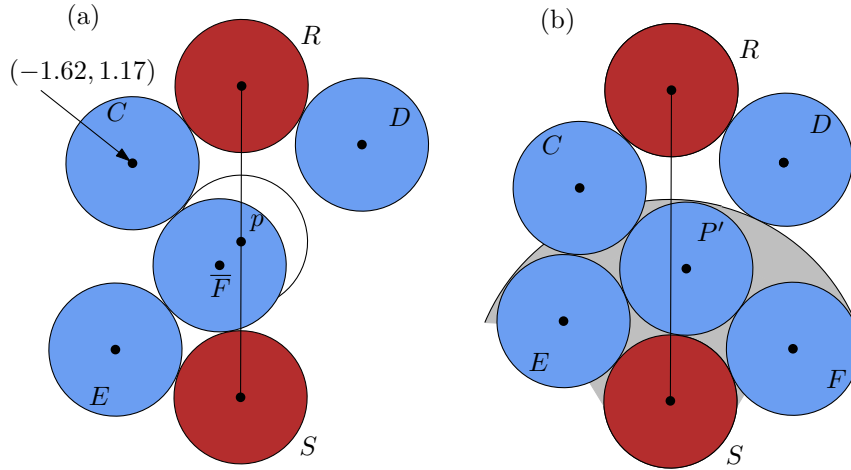


Fig. 11. (a) The upper bound on $|ce|$ in case (ii). (b) Illustrating the proof of case (iii).

size 1. Let \mathcal{B}_1 be the set of all unit disks in \mathcal{B} that are singletons in G . Observe that each disk in \mathcal{B}_1 corresponds to a connected component of \mathcal{F}_1 that has been packed optimally since we applied a greedy post-processing step to the PTAS, see Section 4. Thus we do not lose anything if we place only one $2/3$ -disk for each of the $m'_1 = |\mathcal{B}_1|$ singleton unit disks. Let $\mathcal{B}_2 = \mathcal{B} \setminus \mathcal{B}_1$ and let $m'_2 = m' - m'_1$ be the number of disks in \mathcal{B}_2 . Note that $m'_2 \geq 8(m - m'_1)/9$ since m'_2 is the number of disks in an $8/9$ -approximation of packing unit disks into those components of \mathcal{F}_1 that contain at least two unit disks. Finally let G_2 be the subgraph of G restricted to \mathcal{B}_2 . We show that G_2 contains a matching with at least $m'_2/8$ edges.

By Lemma 1 each connected component \mathcal{C} contains a spanning tree of degree at most 7. In that tree we repeatedly match a leaf with its unique incident vertex. When two vertices are matched they are removed from the tree. This may partition the tree into a forest. For each tree in the forest the process continues iteratively. This yields a matching in \mathcal{C} that contains at least $\lceil |\mathcal{C}|/8 \rceil$ edges. Repeating this argument for each connected component yields a matching in G_2 that contains at least $\lceil m'_2/8 \rceil$ edges.

According to Theorem 2 and Lemma 4 we can pack three $2/3$ -disks in $\mathcal{A}(C, D)$ for every matched pair $\{C, D\}$ such that these $2/3$ -disks are pairwise disjoint. For each of the remaining unmatched disks in G we place one $2/3$ -disk in each unit disk. Let $\mathcal{B}_{2/3}$ be the set of the $2/3$ -disks that we placed, including those for the m'_1 unit disks in \mathcal{B}_1 . By construction no two disks in $\mathcal{B}_{2/3}$ intersect. The cardinality of $\mathcal{B}_{2/3}$ is at least $3\lceil m'_2/8 \rceil + 3\lceil m'_2/4 \rceil + m'_1 \geq 9m'_2/8 + m'_1$. Since $m'_2 \geq 8(m - m'_1)/9$, the set $\mathcal{B}_{2/3}$ contains at least m disks, and we can conclude with the following theorem.

Theorem 3. *Algorithm DISKPACKER is a polynomial-time $2/3$ -approximation for the problem PACKINGSCALEDISKS.*

9. Conclusion

Naturally our main result is purely of theoretical interest. In terms of running time the bottleneck is our disk-packing variant of the PTAS of Hochbaum and Maass, which we apply with an approximation factor of $8/9$. To obtain an algorithm with better running time one might try to consider larger components, or improve the number of matched disks. However, to get a considerable improvement it seems unavoidable to use a completely different approach. Are there fast algorithms for values of α between $1/2$ and $2/3$?

From a theoretical point of view it would be desirable to reconsider the square case in order to narrow the gap between the lower bound $2/3$ and the upper bound $13/14$ of Baur and Fekete.³ Their inapproximability result can be adapted to the disk case, but would yield an upper bound very close to 1, leaving an even larger gap than in the square case.

Acknowledgments

We thank Esther Moet for contributing to an earlier version of this paper. Our thanks also go to Marc van Kreveld who found a mistake in our runtime analysis of the greedy algorithm in the proof of Theorem 1, and to Chris Gray who helped us fix that mistake. We are also very indebted to the anonymous referees whose comments have helped us to improve the paper.

References

1. Franz Aurenhammer. Improved algorithms for discs and balls using power diagrams. *J. Algorithms*, 9(2):151–161, 1988.
2. Franz Aurenhammer and Herbert Edelsbrunner. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, 1984.
3. Christoph Baur and Sándor P. Fekete. Approximation of geometric dispersion problems. *Algorithmica*, 30(3):451–470, 2001.
4. Peter Brass, William O. Moser, and János Pach. *Research problems in discrete geometry*. Springer-Verlag, Berlin, 2005.
5. Paola Cappanera. A survey on obnoxious facility location problems. Technical Report TR-99-11, University of Pisa, 1999.
6. Ee-Chien Chang, Sung Woo Choi, DoYong Kwon, Hyungju Park, and Chee-Keng Yap. Shortest path amidst disc obstacles is computable. In *Proc. 21st ACM Symp. Comput. Geom. (SoCG'05)*, pages 116–125, 2005.
7. Robert J. Fowler, Mike Paterson, and Steven L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Inf. Process. Lett.*, 12(3):133–137, 1981.
8. Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985.
9. Rolf Klein, Kurt Mehlhorn, and Stefan Meiser. Randomized incremental construction of abstract Voronoi diagrams. *Comput. Geom. Theory Appl.*, 3(3):157–184, 1993.
10. Trygve Nagell, editor. *Selected mathematical papers of Axel Thue*. Universitetsforlaget Oslo, 1977.

11. S. S. Ravi, Daniel J. Rosenkrantz, and Giri Kumar Tayi. Heuristic and special case algorithms for dispersion problems. *Operations Research*, 42(2):299–310, 1994.
12. C. Ambrose Rogers. *Packing and Covering*. Cambridge University Press, 1964.
13. Axel Thue. Über die dichteste Zusammenstellung von kongruenten Kreisen in der Ebene. *Christiania Vid. Selsk. Skr.*, 1:3–9, 1910. Reprinted in Ref. 10.
14. Gábor Fejes Tóth. Packing and covering. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 19–41. CRC Press, 2004.
15. D. W. Wang and Yue-Sun Kuo. A study on two geometric location problems. *Inform. Process. Letters*, 28(6):281–286, 1988.
16. Chuanming Zong and John Talbot. *Sphere Packings*. Springer-Verlag, 1999.