A polynomial-time approximation algorithm for a geometric dispersion problem

Marc Benkert[1⋆], Joachim Gudmundsson[2], Christian Knauer[3], Esther Moet[4], René van Oostrum[4], and Alexander Wolff[1⋆]

[1] Faculty of Computer Science, Karlsruhe University, P.O. Box 6980, D-76128 Karlsruhe, Germany. `http://i11www.iti.uka.de/algo/group`
[2] National ICT Australia Ltd[‡], Sydney, Australia, `joachim.gudmundsson@nicta.com.au`
[3] Institute of Computer Science, Freie Universität Berlin, Germany. `christian.knauer@inf.fu-berlin.de`
[4] Department of Computing and Information Sciences, Utrecht University, The Netherlands. {`esther,rene`}`@cs.uu.nl`

Abstract. We consider the problem of placing a maximal number of disks in a polygonal region containing obstacles such that no two disks intersect. Let α be a fixed real in $(0, 1]$. We are given a bounding polygon P and a set $\mathcal{R}$ of possibly intersecting unit disks whose centers lie in P. The task is to pack a set $\mathcal{B}$ of m disjoint disks of radius α into P 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. Baur and Fekete showed that the problem cannot be solved in polynomial time for $\alpha \geq 13/14$, unless $\mathcal{P} = \mathcal{NP}$. In this paper we present an algorithm for $\alpha = 2/3$.

1 Introduction

Obnoxious facility location problems consider the placement of facilities of which clients consider it undesirable to be in the proximity, for instance, nuclear power plants or garbage dumps. There are several models for and variations to the problem; see the survey by Cappanera [2]. We consider the following instance: we are given a bounding rectangle P, a set $\mathcal{R}$ of n points in P (the red points), and an integer k, and we should construct a set $\mathcal{B}$ of k (blue) points such that the minimum distance from a blue point to another point (either red or blue) is maximized over all points in $\mathcal{B}$. If the optimal distance is denoted by r_{opt}, then we can reformulate the problem as follows: we are given a set of n centers of possibly overlapping red disks with unknown radius r_{opt}, and we are to determine r_{opt} and to find a set of k blue disks with radius r_{opt} such that no blue disk overlaps any other disk, whether red or blue.

⋆ Supported by grant WO 758/4-1 of the German Science Foundation (DFG). Part of this work was done while M. Benkert visited NICTA on a DAAD grant.
‡ Funded by the Australian Government's Backing Australia's Ability initiative, in part through the Australian Research Council.

The problem of packing objects into a bounded region is one of the classic problems in mathematics and theoretical computer science, see for example the monographs [10, 12] which are solely devoted to this problem, and the survey by Tóth [11]. In this paper we consider problems related to packing disks into a polygonal region. As pointed out by Baur and Fekete in [1], even when the structure of the region and the objects are simple, only very little is known, see for example [5, 8].

Consider the following decision problem corresponding to our optimization problem: we are given $\mathcal{R}$ and k and a radius r, and we must decide whether $r > r_{\text{opt}}$ or $r \leq r_{\text{opt}}$. In the latter case, we must also give a set $\mathcal{B}$ of k blue disk with radius r such that no blue disk overlaps any other disk. If we had an algorithm at our disposal that solves the decision problem in polynomial time, then we could solve the original optimization problem in polynomial time by applying Megiddo's parametric search [9]. Unfortunately, the decision problem is known to be NP-complete [4]. Therefore we are looking for an algorithm that approximates the decision problem in the following sense: if m disks of radius r can be placed, then our algorithm places m disks of radius αr, for some fixed $\alpha \in (0, 1]$. If $m < k$, then we know that $r > r_{\text{opt}}$, and if $m \geq k$, then either $r \leq r_{\text{opt}}$, or $r > r_{\text{opt}}$ and $\alpha r \leq r_{\text{opt}}$. In other words, placement of at least k disks of radius less than αr_{opt} is guaranteed, and we may or may not be fortunate for radii in between αr_{opt} and r_{opt}.

Obviously, we would like to maximize α while keeping a polynomial running time. Given such an algorithm, we can use it to compute an α-approximation to the original optimization problem, again by using parametric search, albeit in a somewhat non-standard way.

By rescaling r to 1, we can regard the decision problem as that of packing $m \geq k$ unit disks into a rectangle that is already partially covered by n unit disks. In this paper, we consider the following problem:

Problem 1 (ApproxSize) *Let $\alpha \in (0, 1]$ be a fixed real. Given a bounding polygon P and a set $\mathcal{R}$ of possibly intersecting unit disks whose centers lie in P, pack m non-intersecting disks of radius α into P, where m is the maximal number of unit disks that can be packed in P.*

Note that we do not know the value of m a priori. For $\alpha = 1/2$, the problem can be solved by placing disks with radius $1/2$ greedily, i.e., as long as there is space to place a disk, we place one at an arbitrary feasible position. The following simple charging argument shows that we will place at least m disk of radius $1/2$ in this way. Consider an arbitrary placement of m unit disks, and charge a disk C with radius $1/2$ to a unit disk D if the center of C lies inside D. After the greedy algorithm has finished, all of the m unit disks have a charge of at least one. Otherwise, we can place a disk with radius $1/2$ in an uncharged unit disk such that their centers coincide, and this contradicts the termination condition of the greedy algorithm.

In their pioneering work [6] Hochbaum and Maas gave a polynomial-time approximation scheme (PTAS) for the problem of packing a maximal number of unit disks into a region. The problem is known to be NP-complete [4]. Even though the corresponding geometric dispersion problem looks very similar, inapproximability results have been shown. Baur and Fekete [1] proved hardness results for a variety of geometric dispersion problems, and their results can be modified to our setting with a bit of effort. Specifically, they showed that ApproxSize cannot be solved in polynomial time for any radii that exceed 13/14, unless $P = NP$. Furthermore, for the case when the objects are squares, Baur and Fekete gave an $O(\log k \cdot n^{40})$-time 2/3-approximation algorithm, where k is the number of squares. However, since a square is a simpler shape and easier to pack than a disk their approach cannot be generalized to disks. The main contribution of this paper is a polynomial-time 2/3-approximation algorithm. Actually, we conjecture that 2/3 is indeed the largest value for which the problem can be solved in polynomial time.

ApproxSize has applications in non-photorealistic rendering system, where 3D models are to be rendered in an oil painting style, as well as in random examinations of, e.g., soil or water.

2 Algorithm outline

We now give a rough outline of our algorithm DiskPacking. We use the term r-disk as shorthand for a disk of radius r. For $r > 0$ and a set $R \subseteq \mathbb{R}^2$ let the *r-freespace* $\mathcal{F}_r(R)$ of R be the set of the centers of all r-disks in R. By $\mathcal{F}_r^{\otimes}(R)$ we denote the Minkowski sum of $\mathcal{F}_r(R)$ and an r-disk.

We first compute the sets $\mathcal{F}_1 = \mathcal{F}_1(P \setminus \bigcup \mathcal{R})$ and $\mathcal{F}_1^{\otimes} = \mathcal{F}_1^{\otimes}(P \setminus \bigcup \mathcal{R})$, see Fig. 1. Then, we apply the PTAS of Hochbaum and Maas [6] to $\mathcal{F}_1^{\otimes}$. For any positive integer t, the PTAS packs in $O(n^{t^2})$ time 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 $\mathcal{F}_1^{\otimes}$ and n is the minimum number of unit squares whose union covers $\mathcal{F}_1^{\otimes}$. Setting $t = 25$ we obtain in $O(n^{625})$ time a set $\mathcal{B}$ of $m' \geq 12/13 \cdot m$ unit disks.

Note that the approximation scheme by Hochbaum and Maas can be modified such that the algorithm is strongly polynomial with respect to the size of our input. If the number of disks that can be packed is not polynomial in the size of P and $\mathcal{R}$ then there must exist a huge empty square region within P. This can be "cut out" and packed almost optimally by using a naîve approach. The added error obtained is bounded by $O(1/\tilde{n}^2)$ where $\tilde{n}$ is the optimal number of disks that can be packed in the square. This step can be repeated until there are no more huge empty squares.

Given $\mathcal{B}$ we compute a set $\mathcal{B}_{2/3}$ of disks of radius 2/3 that has cardinality at least $13/12 \cdot m' \geq m$ and is contained in $P \setminus \bigcup R$. We obtain $\mathcal{B}_{2/3}$ in two steps. First, we compute a sufficiently large matching in the nearest-neighbor graph $G = (\mathcal{B}, E)$ of $\mathcal{B}$ with respect to a metric $\mathrm{dist}(\cdot, \cdot)$ that we will specify later. Second, we define a region for each pair of matching unit disks such that we can place three 2/3-disks in each region (see Fig. 2) and all regions are pairwise disjoint. For each unmatched unit disk D we place a 2/3-disk D' such that the centers of D and D' coincide.

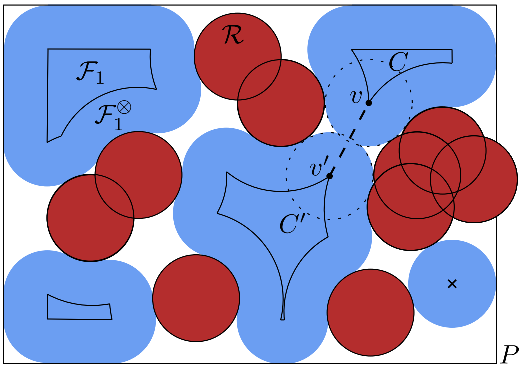

Fig. 1. The 1-freespace $\mathcal{F}_1$ (light shaded) and a shortcut vv' (dashed) between the connected components C and C' of $\mathcal{F}_1$.

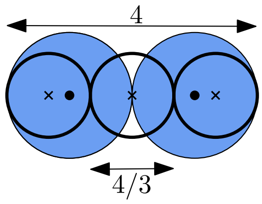

Fig. 2. Packing three 2/3-disks in the region spanned by a pair of unit disks.

In the next sections we describe each step of Algorithm DiskPacking in more detail.

3 The freespace and a metric on unit disks

We briefly recall the setting. We are given a set $\mathcal{R}$ of unit disks whose centers lie in a polygon P. The disks in $\mathcal{R}$ are allowed to intersect. We first compute the freespace $\mathcal{F}_1$ of $P \setminus \bigcup \mathcal{R}$. According to Kedem et al. [7] the union of s disks can be computed in $O(s \log^2 s)$ time and its complexity is linear in s. Applying their algorithm to the disks in $\mathcal{R}$ scaled by a factor of 2 and intersecting the resulting union with P, we can compute $\mathcal{F}_1$ in $O(|\mathcal{R}| \log^2 |\mathcal{R}|)$ time, where $|\mathcal{R}|$ is the cardinality of $\mathcal{R}$.

Next, we want to introduce a metric $\mathrm{dist}(\cdot,\cdot)$ on unit disks in $\mathcal{F}_1^{\otimes}$. With the current definition of $\mathcal{F}_1$ we have the problem that two unit disks centered on points in different connected components of $\mathcal{F}_1$ can intersect. We solve this problem by considering a superset $\mathcal{F}_1^+$ of $\mathcal{F}_1$ that connects close components of $\mathcal{F}_1$. By $|p,q|$ we denote the Euclidean distance between two points p and q in the plane.

Definition 1. *Let $\mathcal{C}_1$ be the set of connected components of $\mathcal{F}_1$, and let $C, C' \in \mathcal{C}_1$. Let v and v' be vertices on the boundaries of C and C', respectively. We say that the line segment vv' is a* shortcut *if $|v,v'| < 2$. Let $\mathcal{S}(C,C')$ be the set of all shortcuts induced by C and C'. We set $\mathcal{F}_1^+ = \mathcal{F}_1 \cup \bigcup_{C,C' \in \mathcal{C}_1;\, s \in \mathcal{S}(C,C')} s$.*

Figure 1 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 Fig. 3.

Definition 2. *Let D and D' be unit disks whose centers d and d' lie in $\mathcal{F}_1$. The distance* $\mathrm{dist}(D, D')$ *of D and D' is the length of the geodesic $g(d, d')$ of d and d' in $\mathcal{F}_1^+$. The* tunnel *$T(D, D')$ of D and D' is the union of all 1-disks and 2/3-disks in $P \setminus \bigcup \mathcal{R}$ centered at points of $g(d, d')$.*

It is easy to see that a 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 $T(D, D')$. The geodesic between two points in $\mathcal{F}_1^+$ can only consist of line segments and arcs of radius 2, see Fig. 3.

Recall that our algorithm computes a matching in the nearest-neighbor graph $G(\mathcal{B}, E)$ induced by the metric $\mathrm{dist}(\cdot, \cdot)$ on the set $\mathcal{B}$ of unit disks that we get from the PTAS by Hochbaum and Maas. For each pair $\{D, D'\}$ in the matching we define a region $T_{2/3}(D, D')$ into which we will then place three 2/3-disks as in Fig. 2. An obvious way to define $T_{2/3}(D, D')$ would be to take the union of all 2/3-disks centered at points of the geodesic between c and d in $\mathcal{F}_{2/3}$. Our definition is not as straight-forward, but will simplify the proof that $T_{2/3}(D, D')$ and $T_{2/3}(F, F')$ are disjoint if D, D', F, and F' are pairwise disjoint.

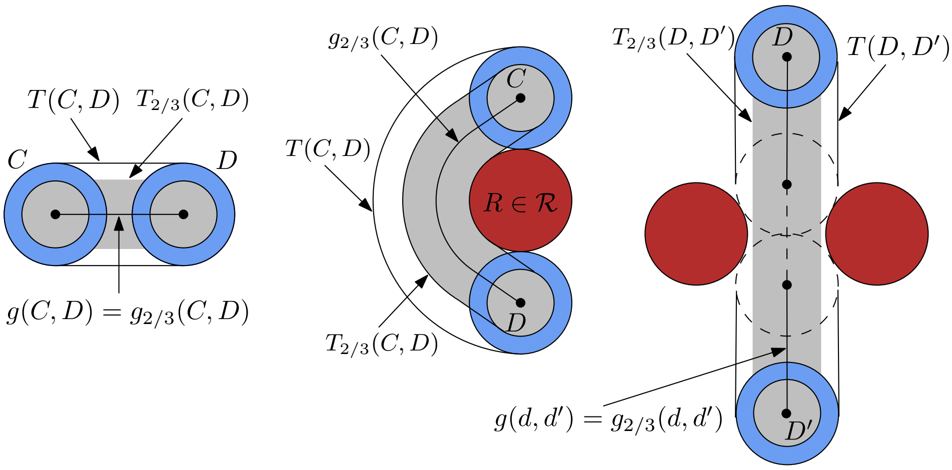

Fig. 3. The geodesic $g(c, d)$. Left: unrestricted. Center: obstacle $R \in \mathcal{R}$. Right: shortcut.

Definition 3. *Let D and D' be unit disks whose centers d and d' lie in $\mathcal{F}_1$. 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' is the union of all 2/3-disks centered at points of $g_{2/3}(d, d')$.*

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

4 The nearest-neighbor graph

Recall that m is the maximum number of disjoint unit disks that fit into $\mathcal{F}_1^{\otimes}$. For $t = 25$ the $(1-1/t)^2$-approximation of Hochbaum and Maas [6] yields a set $\mathcal{B}$ of $m' \geq 12m/13$ unit disks in $\mathcal{F}_1^{\otimes}$. Our plan is to compute the nearest-neighbor graph $G(\mathcal{B}, E)$ induced by the metric $\mathrm{dist}(\cdot,\cdot)$, find a matching of sufficient size in G, and finally place three 2/3-disks in the 2/3-tunnel $T_{2/3}(C,D)$ for each pair $\{C,D\}$ in the matching. If we place another 2/3-disk for each unmatched disk in $\mathcal{B}$, we show that we place at least $13m'/12 \geq m$ disks of radius 2/3 as desired.

By construction, two unit disks D_1 and D_2 whose centers lie in different components of $\mathcal{F}_1^{+}$ have an empty intersection, so we can consider each connected component of $\mathcal{F}_1^{+}$ separately.

After running the algorithm of Hochbaum and Maas we greedily add to $\mathcal{B}$ disjoint unit disks in $\mathcal{F}_1^{\otimes} \setminus \bigcup \mathcal{B}$ until no more disks can be added. This is needed to ensure the following lemma:

Lemma 1. *The nearest neighbor graph G of $\mathcal{B}$ is planar and has maximum degree 6.*

Proof. Let $C \in \mathcal{B}$ be an arbitrary unit disk, let $C' \in \mathcal{B}$ be the nearest neighbor of C in G, and let $\mathcal{D} = \{D_1, \ldots, D_k\} \subseteq \mathcal{B}$ be the neighbors of C in G for which C is their nearest neighbor. If $k \leq 5$ then the degree bound obviously holds, thus we only have to consider the case when $k \geq 6$. For each disk D_i, $1 \leq i \leq k$, place a unit disk D_i' with center on $g(c, d_i)$ such that $|c, d_i'| = 2$, i.e., D_i' touches C. From the definition of the nearest-neighbor graph it follows that every point on $g(c, d_i)$ is closer to C or D_i than to any other unit disk in $\mathcal{B} \setminus \{D_i\}$. As a result the set $D_1', \ldots, D_k'$ and C' has to be disjoint, as illustrated in Fig. 4a. Using a simple packing argument it follows that $k = 6$ and that $C' \in \mathcal{D}$, thus the degree bound stated in the lemma holds.

Finally, G is planar since no two edges in a nearest-neighbor graph can intersect. ☐

Note that G is 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 only one outgoing edge.

From now on we will call $\{C, D\} \subseteq \mathcal{B}$ a *nearest pair* if $\{C, D\}$ is an edge in G, i.e., either D is the nearest disk in $\mathcal{B}$ to C or C is the nearest disk in $\mathcal{B}$ to D. For every nearest pair $\{C, D\}$ we define $\mathcal{A}(C, D)$ to be $C \cup D \cup T_{2/3}(C, D)$. As the nearest pair $\{C, D\}$ is a potential candidate to become a matching pair, we want to ensure that we can use $\mathcal{A}(C, D)$ to pack three 2/3-disks in it such that all the packed 2/3-disks are pairwise disjoint. Thus, we have to prove:

(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 $\mathcal{A}(C, D) \cap \mathcal{A}(E, F) = \emptyset$.

Note that we do not have to care whether, e.g., $\mathcal{A}(C, D)$ intersects $\mathcal{A}(C, E)$ because the matching will choose 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, thus, (i) is fulfilled. The remaining part of the paper will focus on proving (ii).

We split the proof into two parts. The first part shows that $T_{2/3}(C, D)$ does not intersect any disk other than C and D. The second part 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.

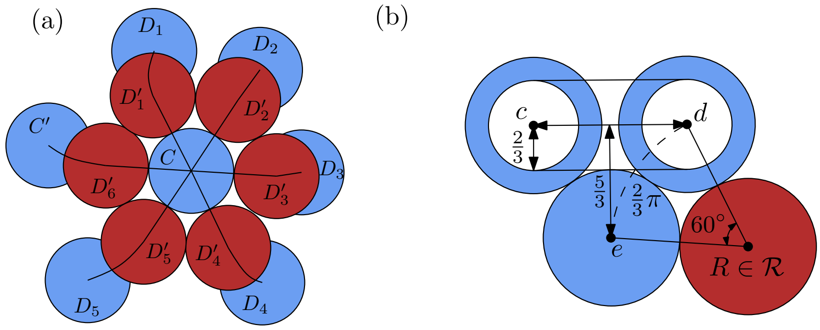

Fig. 4. (a) The nearest neighbor graph of $\mathcal{B}$ is planar and has maximum degree 6. (b) Illustrating the proof of Lemma 2.

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

Proof. Let $T'_{2/3}(C, D)$ be the Minkowski sum of a 2/3-disk and the line segment cd, see Fig. 4b. 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 $|c, d|$ 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 $|c, e| = |d, e| = 2$. 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 $|c, d|$ to be at least $\frac{2}{3}\sqrt{11}$. □

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

Proof. For simplification we assume that D and E touch, as illustrated in Fig. 4b. The curve $g(D, E)$ attains its longest length 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 an arc of radius 2 and 60°, thus its length is at most $\frac{1}{6} \cdot 2 \cdot 2\pi = \frac{2}{3}\pi$. □

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 $T_{2/3}(C, D)$.

W.l.o.g. let C be the nearest disk in $\mathcal{R}$ to D. The proof is done by contradiction, i.e., assume that there is a disk $E \in \mathcal{B}$ that intersects $T_{2/3}(C, D)$.

First, we move a unit disk from the position of D on the center-point curve $g(C, D)$ to the first position in which it hits E, denote the disk in this position by $\overline{D}$, see Fig. 4b where $D = \overline{D}$ holds. Note that $\overline{D}$ does not necessarily lie entirely within $\mathcal{F}_1^{\otimes}$. However, according to Lemma 2 ($C, \overline{D}$ and E are disjoint), the Euclidean distance between c and $\overline{d}$ is at least $\frac{2}{3}\sqrt{11}$. We prove that the geodesic $g(\overline{d}, e)$ within $\mathcal{F}_1^{+}$ is of length less than $\frac{2}{3}\sqrt{11}$. This contradicts C being the nearest neighbor of D.

We have to consider two cases for the upper bound on the length of $g(\overline{d}, e)$. Case 1: the geodesic is located exclusively in $\mathcal{F}_1$. In this case $g(\overline{d}, e)$ attains its longest length if there is an obstacle disk $R \in \mathcal{R}$ that touches $\overline{D}$ and E and $g(\overline{d}, e)$ is an arc of radius 2 and 60°. Lemma 3 yields an upper bound of $\frac{2}{3}\pi \approx 2.09 < 2.21 \approx \frac{2}{3}\sqrt{11}$ for this case. Case 2: the geodesic takes a shortcut. In this case, the length of $g(\overline{d}, e)$ is even less than 2.09. $\square$

Lemma 4 settles that no other disks apart from C and D intersect $T_{2/3}(C, D)$. We still have to show that any two $\frac{2}{3}$-tunnels $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ do not intersect.

Theorem 1. *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. Assume that $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ intersect. First, we exclude the scenario in which the geodesics $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ intersect as well. Note that this comprises the case in which $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ take the same shortcut. Let i be one of the intersection points of $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$. By exchanging the four parts of $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ having endpoints c, d, e, f and i, it is easy to construct a curve that at least contradicts one of $g_{2/3}(C, D)$ and $g_{2/3}(E, F)$ to be a geodesic and thus one of the pairs $\{C, D\}$ or $\{E, F\}$ to be nearest. Thus, we can assume that only $T_{2/3}(C, D)$ and $T_{2/3}(E, F)$ intersect. Obviously, it is enough to prove the theorem for the case in which the tunnels intersect only in a single point p, see Fig. 5a. Thus, we can w.l.o.g. assume that neither $g_{2/3}(C, D)$ nor $g_{2/3}(E, F)$ takes a shortcut (comparing with one of the disks of $\mathcal{D}(S)$ instead of C, D, E or F for the used shortcut S). We will again show that $\{C, D\}$ and $\{E, F\}$ cannot be the nearest pairs at the same time.

We observe that at least one of the disks $\{C, D, E, F\}$ intersects the unit disk P 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. W.l.o.g. let C be a disk that intersects P.

Let p_{CD} be the point on $g_{2/3}(C, D)$ such that $|p, p_{CD}| = 2/3|$, as shown in Fig. 5a. Define p_{EF} correspondingly. We 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 in which at least one vicinity of p_{CD} and p_{EF} is a straight-line is easier and can be treated with 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 Fig. 5a. 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, see Fig. 5a.

As a consequence of Lemma 4 we get that each geodesic $g_{2/3}$ starts with a straight-line segment of length at least $\frac{1}{3}\sqrt{11} \approx 1.11$. The curvature of $g_{2/3}(C, D)$ in p_{CD} infers that $|c, p_{CD}| \geq 1.11$ 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. W.l.o.g. we assume that C and E lie to the left of the y-axis and D and F lie to the right, see Fig. 5a.

Note that we have to take care which relationship inferred that $\{C, D\}$ and $\{E, F\}$ are 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.

The proofs of (i),(ii) and (iii) are in the appendix. □

5 The set $\mathcal{B}_{2/3}$

After computing $\mathcal{B}$ and the nearest neighbor graph $G = (\mathcal{B}, E)$, we compute a matching in G. Let $m' = |\mathcal{B}|$ be the number of unit disks in $\mathcal{B}$. Recall that G is planar and has degree at most 6. We show that we can find a matching in which the number of matched disks is at least $\frac{1}{6} \cdot m'$. Observe that G can consist of more than one connected component. We look at each connected component separately. Let C be a connected component and let c be the number of disks that it contains. Clearly, C contains a spanning tree of degree at most 6. It is easy to see that there is a matching in C that matches at least $\frac{1}{6} \cdot c$ disks. Doing this for each connected component yields a matching in G that contains at least $\frac{1}{6} \cdot m'$ matched disks.

According to Theorem 1 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 we pack one 2/3-disk in each disk. Let set of all disks packed as above be $\mathcal{B}_{2/3}$. By construction, there are no interferences between these sets belonging to different connected components of $\mathcal{F}_1^+$. The cardinality of $\mathcal{B}_{2/3}$ is at least $\frac{1}{6} \cdot \frac{3}{2} \cdot m' + \frac{5}{6} \cdot m' = \frac{13}{12} \cdot m'$. Since the cardinality of $\mathcal{B}$ is at least $\frac{12}{13} \cdot m$, the set $\mathcal{B}_{2/3}$ contains at least m 2/3-disks and we can conclude with the following theorem:

Theorem 2. *Algorithm* DiskPacking *is a polynomial-time* 2/3*-approximation for the problem* ApproxSize.

6 Conclusion

Naturally, our result is purely of theoretic interest. The bottleneck for the running time is the application of Hochbaum and Maas' PTAS with approximation factor $(1 - 1/13)$. To obtain an algorithm with better running time, it seems unavoidable to use a completely different approach. For future work it would also be desirable to narrow the gap between the approximation factor of 2/3 of our algorithm and the inapproximability result of 13/14 of Baur and Fekete [1]. We conjecture that, unless $\mathcal{P} = \mathcal{NP}$, the lower bound of 2/3 is indeed tight.

References

1. C. Baur and S.P. Fekete. Approximation of Geometric Dispersion Problems. *Algorithmica*, 30:450–470, 2001.
2. P. Cappanera. A survey on obnoxious facility location problems. Technical Report TR-99-11, University of Pisa, 1999.
3. E.-C. Chang, S. W. Choi, D. Kwon, H. Park and C.-K. Yap. Shortest path amidst disc obstacles is computable. In proceedings of *Symposium on Computational Geometry*, 2005.
4. R.J. Fowler and M.S. Paterson and S.L. Tanimoto. Optimal packing and covering in the plane are NP-compete. *Information Processing Letters*, 12:133–137, 1981.
5. Z. Füredi. The densest packing of equal circles into a parallel strip. *Discrete & Computational Geometry*, 6:95–106, 1991.
6. D. Hochbaum and W. Maas. Approximation Schemes for Covering and Packing Problems in Image Processing and VLSI. *Journal of the ACM*, 32:130–136, 1985.
7. K. Kedem and R. Livne and J .Pach and M .Sharir. On the union of Jordan regions and collision-free translational motion admidst polygonal obstacles. *Discrete & Computational Geometry*, 1:59–71, 1986.
8. C. Maranas, C. Floudas and P. Pardalos. New results in the packing of equal circles in a square. *Discrete Mathematics*, 128:187–293, 1995.
9. N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30(4):852–865, 1983.
10. C. A. Rogers. Packing and Covering. *Cambridge University Press*, 1964.
11. G. F. Tóth. Packing and Covering. In Handbook of Discrete and Computational Geometry, 2nd Edition J. E. Goodman and J. O'Rourke, editors CRC Press LLC, 2004.
12. C. Zong and J. Talbot. Sphere Packings. Springer-Verlag, 1999.

Appendix

Theorem 1. *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. We still have to prove:

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

(i): To prove that $\operatorname{dist}(C, E) < \operatorname{dist}(E, F)$ 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. 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 Fig. 5b. The proof of (i) can then be completed by showing that $\operatorname{dist}(C, \overline{E}) < \operatorname{dist}(\overline{E}, F)$.

Can a unit disk pass between C and S and thus not intersect C. This could only be achieved by maximizing $|c, s|$. Recall that C intersects P, thus, $|c, s|$ is maximized if C touches R and P, i.e. C takes its left and topmost position, as shown in Fig. 5b. 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 calculations. Hence, it holds that $|c, s| \leq 3.86$ which in turn yields that no unit disk can pass between C and S since this would require $|c, s| \geq 4$.

Next, we try to minimize the distance $\operatorname{dist}(\overline{E}, F)$ in order to get F to be closer to $\overline{E}$ than C. For this, $\overline{E}$ should take its rightmost position. It is attained if C is as far as possible from S, i.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 $\operatorname{dist}(\overline{E}, F) \geq 1.29 + 1.11 = 2.40$ as $|p_{EF}, F| \geq 1.11$ holds. This settles (i) since $\operatorname{dist}(C, \overline{E}) \leq 2.09$, according to Lemma 3.

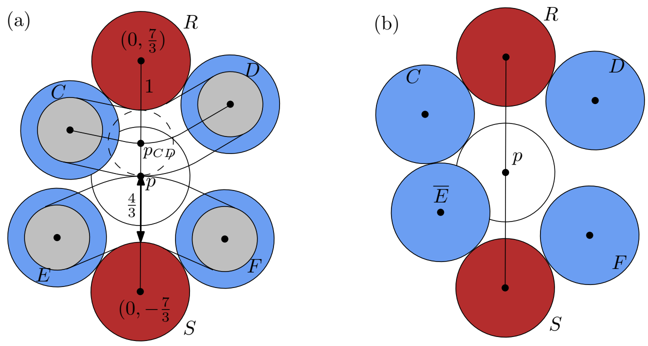

Fig. 5. Illustration for the proof of (a) Theorem 1 and (b) case (i).

(ii): We establish a lower bound on $|c, d|$ and an upper bound on $|c, e|$, which will show that $\operatorname{dist}(C, E) < \operatorname{dist}(C, D)$ holds in all cases. For the lower bound on $|c, d|$ 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 $|c, d|$; 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 that D is not allowed to intersect the tunnel $T(E, F)$. If it did 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 P.) However, $\text{dist}(D, E) < \text{dist}(E, F)$ 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 Fig. 6a. As we know all the side lengths in $\triangle rsd$ we can apply the cosines theorem to compute $\angle srd = \beta$ to be $\arccos \frac{11}{21} \approx 58.41°$. Using β we can compute the left and bottommost coordinates of d to be $(1.70, 1.29)$.

For the right and bottommost position of C we only use the following: Let $\overline{f}$ be the rightmost point on $g(E, F)$ such that $\overline{F}$ touches either C or E. We use that E has to be touched because otherwise C is closer to F than E. We compute the right and bottommost position of C if $\overline{F}$ touches C and E at the same time, see Fig. 6b. Note, that this even yields a position in which C is actually closer to F than E (w.r.t. metric d). As $\triangle ce\overline{f}$ and $\triangle es\overline{f}$ are equilateral, $|c, s|$ is $2\sqrt{3}$, twice the height in an equilateral triangle of side length 2. As we now know all the side lengths in $\triangle rcs$ we can apply the cosines theorem to compute $\angle crs = \alpha$ to be $\arccos \frac{31}{42} \approx 42.43°$. By means of α we compute the right and bottommost coordinates of c to be $(-1.35, 0.86)$. Now a lower bound on $|c, d|$ is the Euclidean distance between $(1.70, 1.29)$ and $(-1.35, 0.86)$ which is 3.08. By Pythagoras' theorem we can also compute the coordinates of $\overline{f}$ to be $(0.25, -0.35)$ — we will need them in the proof of (iii).

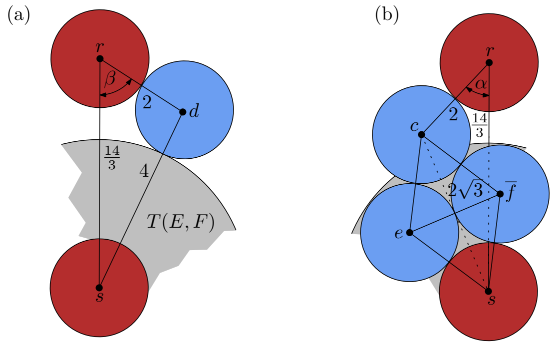

Fig. 6. Illustration of the proof of the lower bound on $|c, d|$ in case (ii).

In order to prove (ii) it now remains to show an upper bound on $|c, e|$ which is less than 3.08 because $g(C, D)$ for the lower bound on $|c, d|$ is an arc anyway. We try to push C and E as far away from each other as possible, under the restriction that E can still be the nearest neighbor of F. It is clear that C has to take its left and topmost position which we already know is $(-1.62, 1.17)$, while E should take its left and bottommost position. Again, we only use that on the rightmost point $\overline{f}$ on $g(E, F)$ such that either C or E is touched by $\overline{F}$, it must be E that is touched. First, we compute the rightmost point $\overline{f}$ on $g(E, F)$ where $\overline{F}$ 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 $\overline{f}$ by Pythagoras' theorem. It holds that $\overline{f}$ is $(-0.33, -0.36)$, see Figure 7a. Now, E takes its left and bottommost position if it touches S and $\overline{F}$. By Pythagoras again we compute e for this position of E to be $(-1.87, -1.64)$. This finally yields the upper bound on $|c, e|$ of 2.82 and we are done with (ii).

(iii): We use the lower bound on $|c, d|$ that was derived in (ii). Thus, we only have to show an upper bound on $|d, f|$ which is less than 3.08. 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 $|c, d|$. 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 $\overline{D}$, and show that $|\overline{d}, f| < 3.08$ holds. Then, $\text{dist}(D, F) < \text{dist}(C, D)$ also holds.

In order for F to take its right and bottommost position, C also has to take its right and bottommost position; look at the disk P' that touches C and S and lies to the right of cs, see Figure 7b. For the right and bottommost position of C which is $(-1.35, 0.86)$, we already computed the position of P' to be $(0.25, -0.35)$. As P' moves left and downwards if the position of C is to the left and/or above $(-1.35, 0.86)$, this implies that D does not intersect P' as $|d, \overline{f}| \approx 2.19 > 2$ holds. 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. This is the reason why F also takes its right and bottommost position if C takes its right and bottommost position. By Pythagoras' theorem we compute this position of F to be $(1.84, -1.55)$. This finally settles the proof because then the upper bound on $|d, f|$ is $2.84 < 3.08$. □

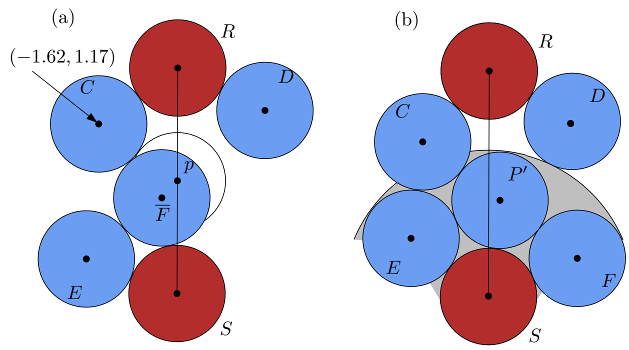

Fig. 7. (a) The upper bound on $|c, e|$ in case (ii). (b) Illustrating the proof of case (iii).