

Connecting a set of circles with minimum sum of radii



Erin W. Chambers^a, Sándor P. Fekete^{b,*}, Hella-Franziska Hoffmann^c,
Dimitri Marinakis^d, Joseph S.B. Mitchell^e, Venkatesh Srinivasan^f, Ulrike Stege^f,
Sue Whitesides^f

^a Department of Computer Science, Saint Louis University, USA

^b Department of Computer Science, TU Braunschweig, Germany

^c David R. Cheriton School of Computer Science, University of Waterloo, Canada

^d Kinsol Research Inc., Duncan, BC, Canada

^e Department of Applied Mathematics and Statistics, Stony Brook University, USA

^f Department of Computer Science, University of Victoria, Canada

ARTICLE INFO

Article history:

Received 6 March 2015

Accepted 21 May 2017

Available online 7 June 2017

Dedicated to our dear friend and colleague, Ferran Hurtado, who has inspired us to pose, to ponder, to question, to solve, to celebrate, and to enjoy life to the fullest. Thank you for inspiring this generation, and many generations to come!

Keywords:

Intersection graphs

Connectivity problems

NP-hardness problems

Approximation

Upper and lower bounds

ABSTRACT

We consider the problem of assigning radii to a given set of points in the plane, such that the resulting set of disks is connected, and the sum of radii is minimized. We prove that the problem is NP-hard in planar weighted graphs if there are upper bounds on the radii and sketch a similar proof for planar point sets. For the case when there are no upper bounds on the radii, the complexity is open; we give a polynomial-time approximation scheme. We also give constant-factor approximation guarantees for solutions with a bounded number of disks; these results are supported by lower bounds, which are shown to be tight in some of the cases. Finally, we show that the problem is polynomially solvable if a connectivity tree is given, and we conclude with some experimental results.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Problems of connectivity are among the most fundamental ones for many types of networks. Typically, these arise in a geometric setting, e.g., when considering the relation between the location of transmitters, the range of their transmissions, and their ability to connect. As a result, important aspects include the underlying geometry, the nature of connectivity, and the study of corresponding cost functions. Thus, connectivity problems bring together graph theory with computational geometry, a combination that was always dear to Ferran Hurtado; e.g., see [1] for a study of connectivity that mixes graphs and geometry.

* Corresponding author.

E-mail addresses: echambe5@slu.edu (E.W. Chambers), s.fekete@tu-bs.de (S.P. Fekete), hrhoffma@uwaterloo.ca (H.-F. Hoffmann), dmarinak@kinsolresearch.com (D. Marinakis), jsbm@ams.sunysb.edu (J.S.B. Mitchell), srinivas@uvic.ca (V. Srinivasan), ustege@uvic.ca (U. Stege), sue@uvic.ca (S. Whitesides).

<http://dx.doi.org/10.1016/j.comgeo.2017.06.002>

0925-7721/© 2017 Elsevier B.V. All rights reserved.

In this paper, we consider a natural connectivity problem, with a focus on the geometric aspects, arising from assigning ranges to a set of points, such that the resulting disk intersection graph is connected. More precisely, we are given a set of points $P = \{p_1, \dots, p_n\}$ in the plane. Each point p_i is assigned a range r_i , inducing a disk of radius r_i . Two points p_i, p_j are adjacent in the connectivity graph H , if their disks intersect. The CONNECTED RANGE ASSIGNMENT PROBLEM (CRA) requires an assignment of radii to P , such that the objective function $R = \sum_i r_i^\alpha$, $\alpha = 1$ is minimized, subject to the constraint that H is connected.

Problems of this type have been considered before and have natural motivations from fields including networks, robotics, and data analysis, where ranges have to be assigned to a set of devices, and the total cost is given by an objective function that considers the sum of the radii of disks to some exponent α . The cases $\alpha = 2$ or 3 correspond to minimizing the overall power. The motivation for the case $\alpha = 1$ arises from scanning the corresponding ranges with a minimum required angular resolution, so that the scan time for each disk corresponds to its perimeter, and thus radius.

1.1. Related work

There is a large body of literature on algorithmic methods for range assignment in wireless sensor and ad-hoc networks; see Calinescu et al. [6], Calinescu and Wan [7], Carmi and Katz [9], Carmi et al. [11], Caragiannis et al. [8], Lloyd et al. [20], Wan et al. [23] for a (small) selection of aspects and variants.

In the context of clustering, Doddi et al. [15], Charikar and Panigraphy [13], and Gibson et al. [17] consider the following problems. Given a set P of n points in a metric space with metric $d(i, j)$ and an integer k , partition P into a set of at most k clusters with minimum sum of either (a) cluster diameters, or (b) cluster radii. Thus, the most significant difference with our problem is the lack of a connectivity constraint. Doddi et al. [15] provide approximation results for (a). They present a polynomial-time algorithm, which returns $O(k)$ clusters that are $O(\log(\frac{n}{k}))$ -approximate. For a fixed k , they transform an instance of their problem into a min-cost set-cover problem instance yielding a polynomial-time 2-approximation. They also show that the existence of a $(2 - \epsilon)$ -approximation would imply $P = NP$. In addition, they prove that the problem in weighted graphs without triangle inequality cannot be efficiently approximated within any factor, unless $P = NP$. Note that every solution to (b) is a 2-approximation for (a). Thus, the approximation results can be applied to case (a) as well. A greedy logarithmic approximation and a primal-dual based constant factor approximation for minimum sum of cluster radii is provided by Charikar and Panigraphy [13].

Another often considered setting is the one in which the coverage of a given set of base stations is required. In this setting, Alt et al. [2] consider a closely related problem of selecting disk centers and radii such that a given set of points in the plane are covered by the disks. Like our work, they focus on minimizing an objective function based on $\sum_i r_i^\alpha$ and produce results specific to various values of α . The minimum sum of radii disk coverage problem (with $\alpha = 1$) is also considered by Lev-Tov and Peleg [18] in the context of radio networks. Again, connectivity is not a requirement. In a more geometric setting, Bilò et al. [5] provide approximation schemes for the minimum size k -clustering problem that requires dividing the set of centers into at most k clusters with minimum cluster cost.

A lot of work has also been done on radii/range assignment problems which require the special connectivity of “communication”. The work of Clementi et al. [14] considers minimal assignments of transmission power to devices in a wireless network such that the network stays connected. In that context, the objective function typically considers an $\alpha > 1$ based on models of radio wave propagation. Furthermore, in the type of problem considered by Clementi et al. the connectivity graph is directed; i.e. the power assigned to a specific device affects its transmission range, but not its reception range. This is in contrast to our work in which we consider an undirected connectivity graph. See Fuchs [16] for a collection of hardness results of different (directed) communication graphs. Carmi et al. [10] prove that an Euclidean minimum spanning tree is a constant-factor approximation for a variety of problems including the *Minimum-Area Connected Disk Graph* problem, which equals our problem with the different objective of minimizing the *area* of the *union* of disks, while we consider minimizing the *sum* of the *radii* (or perimeters) of all disks. This can apply to a hybrid robot or sensor network system such as that described in Marinakis et al. [21]: Consider a situation in which a mobile robot is required to visit the region of each static network component, e.g., for environmental monitoring purposes, and requires constant one-way information from at least one static network component at all times; e.g., for navigational purposes. In this case, a reasonable objective is to reduce the area covered by the network while maintaining connectivity.

1.2. Our work

In this paper we investigate a variety of algorithmic aspects of the CRA problem. In Section 2, we show that for a given connectivity tree, an optimal solution can be computed efficiently. Section 3 gives a proof of NP-hardness for the problem when there is an upper bound on the radii, with full details for the case of planar weighted graphs and a sketch for planar point sets. Section 4 provides a number of approximation results for solutions with bounded number of disks. In Section 5, we present a PTAS for the case in which there are no upper bounds on the radii of the disks. These theoretical results are complemented by experimental results in Section 6. A concluding discussion with open problems is provided in Section 7.




Fig. 1. One step in Algorithm 1 for a tree T of height h_T (left). The tree \tilde{T} of height $h_T - 1$ resulting from the remaining gaps is shown on the right.

2. CRA for a given connectivity tree

In some settings, a specific set of connectivity edges is required instead of *any* connectivity tree. In this section, we will show that the Connected Range Assignment problem is solvable in polynomial time, when we are given a specific connectivity tree. First we give a formal definition of this variant.

In the *Connected Range Assignment Problem for a Given Connectivity Tree (CRACT)*, we are given a finite set P , a metric d and a spanning tree T for the vertex set P . We are required to find an assignment r_i of radii for point p_i in P such that $R = \sum_i r_i$ is minimized subject to the constraint that if (p_i, p_j) is an edge in T , then the two corresponding disks C_i and C_j (with radii r_i and r_j respectively) intersect.

We call a solution for Problem CRACT an *optimal range assignment* for T . We first state some structural properties of optimal range assignments and then, based on those, provide a polynomial-time algorithm for Problem CRACT. In the following, we always consider the case where $|P|$ is at least 2 and assume that T is rooted at some internal node p_r of T .

Lemma 1. *Given a connectivity tree T with at least three nodes, there exists an optimal range assignment for T with $r_l = 0$ for all leaves p_l of T .*

Proof. Consider a CRACT instance (P, d, T) . Assume an optimal range assignment r for T has a leaf $p_l \in P$ with $r_l > 0$. Then, the disk C_l around p_l with radius r_l must intersect the disk C_u around p_l 's parent p_u with radius r_u . We extend r_u to $r_u := r_u + r_l$ and set $r_l := 0$. This does not increase $R = \sum_{p_i \in P} r_i$ and maintains a solution for T . \square

The following lemma extends Lemma 1. It shows that there is an optimal range assignment for which it is possible to compute the radius assigned to any point p of height h , given the radius assigned to its children.

Lemma 2. *Given a CRACT instance (P, d, T) , there is an optimal range assignment r for T that satisfies the following conditions:*

1. $r_l = 0$ for all leaves p_l of T ;
2. For any node $p_u \in P$ of height h , $1 \leq h \leq h_T$, in T ,

$$r_u = \max_{p_c \text{ is child of } p_u} d(p_c, p_u) - r_c.$$

Proof. To prove that r satisfying the two conditions above is an optimal range assignment, we show that any optimal range assignment can be gradually converted into assignment r in a sequence of $h_T + 1$ steps as follows: At the end of step i , all nodes p_u in T with height h , $0 \leq h \leq i - 1$, will be assigned the radius r_u . Furthermore, the sum of radii at the end of step i is at most the sum of radii at the end of step $i - 1$.

[Step 1] Perform the transformation as described in the proof of Lemma 1. It is easily seen that the condition (1) is true at the end of step 1.

[Step i] Let p_u be a node of height $i - 1$ in T . Further let $d_{\max}(p_u) := \max d(p_c, p_u)$ over all children p_c of p_u . Then we have $r_u \geq d_{\max}(p_u)$. If $r_u > d_{\max}(p_u)$, we shrink C_u until it just covers its farthest child and extend the circle of p_u 's unique parent (height $\geq i$) by the same amount. This ensures that condition (2) is satisfied for the point p_u . Furthermore, this maintains all connections induced by C_u before shrinking and does not increase the solution value R . \square

Using the observed optimal solution properties, we next describe an algorithm which will result in an optimal range assignment for T (see Fig. 1).

Theorem 3. *CRACT can be solved in $O(n)$ time.*

Algorithm 1: Algorithm solving *CRCT*.**Input :** A *CRCT* instance (P, d, T) , where $T = (P, E_T)$ of height h_T with weighted edges: $d : E_T \rightarrow \mathbb{R}_+$.**Output:** An optimal range assignment $r : P \rightarrow \mathbb{R}_+$ for T

```

 $h := 0;$ 
forall leaves  $p_l$  in  $T$  do  $r_l := 0;$ 
while  $h \leq h_T$  do
  forall nodes  $p_u$  of height  $h$  in  $T$  do
     $r_u := \max_{p_c \text{ child of } p_u} \{d(p_c, p_u) - r_c\};$ 
  end
   $h := h + 1;$ 
end
return  $r$ 

```

Proof. 1) [Correctness]. In the range assignment r computed by Algorithm 1, a node p_u at height h , $0 < h \leq h_T$, in T is assigned a radius r_u where $r_u := \max_{p_c \text{ child of } p_u} \{d(p_c, p_u) - r_c\}$ and the leaves p_l are assigned the radius $r_l = 0$. The correctness of the algorithm now follows from Lemma 2.

2) [Runtime]. As the height of any node in T is unique for a fixed root, the algorithm assigns radii to each node in the tree exactly once. For every node p of height $h \geq 1$, the computation of the required radius needs time proportional to its degree. During this computation, each edge in the tree (of which we have exactly $n - 1$) is considered exactly once in total for the computation. Hence, the algorithm determines an optimal range assignment in $O(n)$ steps. Note that even if the height of the nodes cannot be gained directly from the data structure, it can be computed via one iteration over the tree, which again only requires linear time. \square

3. NP-hardness for bounded radii

In this section, we show that, if an upper bound of β on the radii is also specified, the problem CRA becomes NP-hard using a reduction from 3-SAT. More formally, we are given a finite set P of objects, the pairwise distances of the elements in P measured by metric d and an upper bound β on the radii. The goal is to find an assignment of radii to P , such that the objective function $R = \sum_i r_i$ is minimized, subject to the constraint that H is connected and $r_i \leq \beta$ for all i .

Definition 1 (3-SAT). In logic, a *literal* is a single (possibly negated) variable. A disjunction of three literals is called a *clause* γ_i . Given a Boolean formula $\mathcal{F} = \gamma_1 \wedge \gamma_2 \wedge \dots \wedge \gamma_c$ that is a conjunction of c clauses formed by a set of v variables, the 3-Satisfiability problem (3-SAT) is to decide whether there exists an assignment of truth values to each of the v variables, such that \mathcal{F} evaluates to true. The *variable-clause incidence graph* consists of a vertex for each variable and each clause; an edge exists between a clause vertex representing a clause c_j and a variable vertex representing a variable x_i , iff variable x_i occurs in clause c_j . A 3-SAT instance is *planar* if the clause-variable incidence graph is planar.

It was shown by Lichtenstein [19] that 3-SAT is NP-complete, even when restricted to planar instances. This is the basis for many NP-completeness reductions for planar graphs and for point sets. In the following, we give a proof of NP-hardness of the problem CRA with bounded radii in the graph setting.

Theorem 4. *With radii bounded by some constant β , the problem CRA is NP-hard for planar weighted graphs.*

Proof. See Fig. 2 for the main ideas behind the construction. The proof uses a reduction from PLANAR 3SAT. Variables are represented by closed “loops” at distance β . Additional “connectivity” edges ensure that all variable gadgets are connected. Each clause is represented by a star-shaped set of four points which is connected to each corresponding variable loop over a β -edge. We claim that there is an optimal solution for the constructed CRA instance “ $I_{CRA}(I_{3Sat})$ ” with cost $R = v \cdot 6\beta + c \cdot \beta$ if and only if there is a satisfying variable assignment for the corresponding 3SAT instance I_{3Sat} with v variables and c clauses.

“ \Rightarrow ”: To verify that claim we first state some properties of solutions to an $I_{CRA}(I_{3Sat})$ CRA instance:

1. *There are exactly two different feasible connected radii assignments of value 6β for each variable loop.*
To make this clear see the triangle in Fig. 3 which illustrates that only “non-fractional” solutions provide the required sum of radii. More precisely, if more than one disk is used, the sum of radii either exceeds β or one point remains uncovered, as the sum of each two radii must be greater or equal to β . For symmetric reasons this argument yields two valid “integral” 6β solutions for each entire variable loop.
2. *The four nodes of a clause can be connected to the rest with cost at most β only if there is at least one satisfying variable.*
Since every clause star center p_c can only be covered by disks around points that are not part of the variable loops, the only way to connect the star centers to the remainder with cost at most β for each clause, is to use exactly one




Fig. 2. A graph representation for the 3SAT instance $(\bar{x}_1 \vee x_2 \vee \bar{x}_3) \wedge (x_1 \vee \bar{x}_2 \vee \bar{x}_4)$. All displayed edges have length β . “True” and “false” vertices are marked in bold white or black, respectively; auxiliary vertices are indicated by small dots. Clause vertices are indicated by small triangles, while the subgraphs for variables are labeled by variable names. Marked in red are the vertices assigned with a radius of β , as well as the corresponding connecting edges. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. The triangular parts of the variable gadgets. A coverage by multiple disks (left) always costs more than a coverage with one disk (right).




Fig. 4. We cover each clause star by a β -disk centered at one of the inner nodes.

β -disk centered at one of the four corresponding clause nodes. Using p_c as the β -disk center we must have a satisfying variable loop that closes the gap (of length β) between the clause and the remainder of the graph. Moreover, if we use a different node of the clause as the β -disk center, the remaining two uncovered star nodes must also be covered by disks from the variable loops; see Fig. 4 for illustration.

From the first property we know that the loop construction directly provides a lower bound of $v \cdot 6\beta$ for the costs of the total variable loop coverage. Thus, only costs of $c \cdot \beta$ remain for the coverage of the clauses with total costs $R = v \cdot 6\beta + c \cdot \beta$.




Fig. 5. A class of CRA instances that need $k + 1$ disks in an optimal solution.

Combined with the second property we conclude that there is a solution to $I_{CRA}(I_{3Sat})$ with cost $R = v \cdot 6\beta + c \cdot \beta$ only if there is a satisfying variable assignment for I_{3Sat} .

“ \Leftarrow ”: Given a solution to I_{3Sat} . We have seen that any satisfying variable assignment directly encodes a variable loop coverage in $I_{CRA}(I_{3Sat})$ with cost $v \cdot (6\beta)$. If there is a variable that is not needed (i.e., one for which the value does not impact satisfiability of the clauses), it can be either treated as a normal “true” or “false” variable and is automatically connected to another variable loop via the “connectivity edges”. The clause stars can each be covered with a β -disk centered at the clause star center p_c . Thereby, they are all connected to the remainder as we have at least one connecting (satisfying) variable loop for each clause. This construction costs $c \cdot \beta$. In total, for every solution to I_{3Sat} we get a solution with the desired costs for $I_{CRA}(I_{3Sat})$. \square

The main reason why this proof cannot be extended to the case of unbounded radii is that we consider the objective function $Q_1 = \sum r_i$ which does not prefer smaller disks as any Q_α with $\alpha > 1$ does. That is, although we can construct some local loops that allow only two different integral optimal solutions, we are unable to see a way to connect these gadgets guaranteeing the absence of an optimal global 1-disk solution.

For the case of bounded radii, we can also show an NP-hardness proof for a geometric setting; again the proof uses a reduction to planar 3-Satisfiability. For this, we construct variable gadgets that each consist of disk-shaped point clusters. At the boundary of these clusters we uniformly arrange points such that exactly every third point is used as a disk center in an optimal solution. Chains of β -distant points ensure interior connectivity of the variable gadgets. (This is completely analogous to the construction for the proof of NP-completeness of packing squares into a polygonal domain presented as Theorem 1 in the paper by Baur and Fekete [4].) The clause gadgets are constructed as in the graph case. Additional point chains connect them to the corresponding variable loops.

4. Solutions with a limited number of circles

A natural class of solutions arises when only a limited number of k disks may have positive radius. In this section we show that these k -disk solutions already yield good approximations; we start by giving a class of lower bounds.

We first observe that the number of disks needed for an optimal solution cannot be bounded by a constant number in general.

Theorem 5. *For any integer $k \geq 1$, there is an instance of $n = k + 3$ points in the plane that need exactly $k + 1$ disks to be covered optimally.*

Proof. Consider the example in Fig. 5 with $n = k + 3$ points. The provided solution R^* with $k + 1$ disks is optimal, as $R^* := \sum r_i^* = \frac{\text{dist}(p_0, p_n)}{2}$. To show that the best k -disk solution is not optimal, we make the following observation. For any integer $k \geq 1$, we have $d_1 = d(p_0, p_k) = 2 \cdot \sum_{i=0}^{k-2} 2^i + 2^{k-1} < 2 \cdot 2^k + 2^{k-1} = d(p_k, p_n) = d_2$. Therefore, any k -disk solution that does not use point p_{k+1} cannot be optimal, as the sum of the radii would be strictly greater than R^* . On the other hand, if the k -disk solution contains a disk centered at p_{k+1} , it must be of radius r_{k+1}^* in order to be optimal. This results in a set of $n - 1$ points to be connected optimally using $k - 1$ disks. Hence, using induction on k , we conclude that exactly $k + 1$ disks are needed. \square

In the following we establish upper and lower bounds on the approximation ratios of 1- and 2-disk solutions for CRA.

Lemma 6. *Let \mathcal{P} be a longest (simple) path in a connectivity graph corresponding to an optimal solution R^* , and let e_m be an edge in \mathcal{P} containing the midpoint of \mathcal{P} . Then $R^* = \sum r_i^* \geq \max\{\frac{1}{2}|\mathcal{P}|, |e_m|\}$.*

This follows directly from the following simple property of the connectivity graph: In this graph, for any edge $e = p_u p_v$ with length $|e|$ in \mathcal{P} , $r_u + r_v \geq |e| = d(p_u, p_v)$.

Theorem 7. *A best 1-disk solution for CRA is a $\frac{3}{2}$ -approximation, even in the graph version of the problem.*




Fig. 6. The presence of an uncovered point p' (red) implies the existence of a longer path $\mathcal{P}' \cup (p_j, \dots, p_k)$ (sketched by the curved red line). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 7. A lower bound of $\frac{3}{2}$ for 1-disk solutions.




Fig. 8. Two $\frac{4}{3}$ -approximate 2-disk solutions (light gray, dashed): (Top) **Case 2a**, when $|e_m| < d_i + d_{i+1}$; (bottom) **Case 2b**, where $|e_m| > d_i + d_{i+1}$.

Proof. Consider a longest path $\mathcal{P} = (p_0, \dots, p_k)$ of length $|\mathcal{P}| = d_{\mathcal{P}}(p_0, \dots, p_k) := \sum_{i=0}^{k-1} |p_i p_{i+1}|$ in the connectivity graph of an optimal solution. Let $R^* := \sum r_i^*$ be the cost of the optimal solution, and $e_m = p_i p_{i+1}$ as in **Lemma 6**. Let $\bar{d}_i := d_{\mathcal{P}}(p_i, \dots, p_k)$ and $\bar{d}_{i+1} := d_{\mathcal{P}}(p_0, \dots, p_{i+1})$. Then $d_{\min} := \min\{\bar{d}_i, \bar{d}_{i+1}\} \leq \frac{\bar{d}_i + \bar{d}_{i+1}}{2} = \frac{d_{\mathcal{P}}(p_0, \dots, p_i) + 2|e_m| + d_{\mathcal{P}}(p_{i+1}, \dots, p_k)}{2} = \frac{|\mathcal{P}|}{2} + \frac{|e_m|}{2} \leq R^* + \frac{R^*}{2} = \frac{3}{2}R^*$. So a single disk C with a radius of $d_{\min} \leq \frac{3}{2}R^*$ around the point in \mathcal{P} that is nearest to the middle of path \mathcal{P} covers \mathcal{P} . Without loss of generality, let this center point of C be p_{i+1} , then $d_{\min} = \bar{d}_{i+1}$. We claim that the chosen disk C also covers all points in \mathcal{P} . This can be seen by assuming the existence of a point $p' \in \mathcal{P}$ that is not covered by C ; see **Fig. 6**. Because the optimal solution must induce a connected connectivity graph, there is a path \mathcal{P}' from p' to a node p_j in \mathcal{P} that has no edge in common with \mathcal{P} , i.e. $V(\mathcal{P}') \cap V(\mathcal{P}) = \{p_j\}$. Without loss of generality, let p_j lie between p_0 and p_i in \mathcal{P} . If p' is not covered by C , the path \mathcal{P}' must be longer than the path segment (p_0, \dots, p_j) of \mathcal{P} . We get a path $\mathcal{P}' \cup (p_j, \dots, p_k)$ with length greater than $|\mathcal{P}|$, which contradicts the maximality of \mathcal{P} . \square

Fig. 7 shows that this bound is tight. There is a 2-disk solution of cost 2 while the best 1-disk solution has cost 3. We now show that using two disks yields an even better approximation factor.

Theorem 8. A best 2-disk solution for CRA is a $\frac{4}{3}$ -approximation, even in the graph version of the problem.

Proof. Let $\mathcal{P} = (p_0, \dots, p_k)$ be a longest path of length $|\mathcal{P}| = d_{\mathcal{P}}(p_0, \dots, p_k) := \sum_{i=0}^{k-1} |p_i p_{i+1}|$ in the connectivity graph of an optimal solution with radii r_i^* . Then by **Lemma 6**, $R^* := \sum r_i^* \geq \frac{1}{2}|\mathcal{P}|$. We distinguish two cases; see **Fig. 8**.

Case 1. There is a point x on \mathcal{P} at a distance of at least $\frac{1}{3}|\mathcal{P}|$ from both endpoints. Then there is a 1-disk solution that is a $\frac{4}{3}$ -approximation.




Fig. 9. A non-overlapping optimal solution.

Case 2. There is no such point x . Then two disks are needed. One of them is placed at a point in the first third of \mathcal{P} , and the other disk is placed at a point in the last third of \mathcal{P} . Let $e_m = p_i p_{i+1}$ be defined as in Lemma 6. Further, let $d_i := d_{\mathcal{P}}(p_0, \dots, p_i)$, and let $d_{i+1} := d_{\mathcal{P}}(p_{i+1}, \dots, p_k)$. Then $|e_m| = |\mathcal{P}| - d_i - d_{i+1}$ and $d_i, d_{i+1} < \frac{1}{3}|\mathcal{P}|$.

Case 2a. If $|e_m| < \frac{1}{2}|\mathcal{P}|$ then $d_i + d_{i+1} = |\mathcal{P}| - |e_m| > \frac{1}{2}|\mathcal{P}| > |e_m|$. Set $r_i := d_i$ and $r_{i+1} := d_{i+1}$, then the path is covered. Analogously to the proof of Theorem 7, we conclude that there is no point $p' \in P$ uncovered by the chosen disks, as there will otherwise be a longer path than \mathcal{P} . Since $d_i, d_{i+1} < \frac{1}{3}|\mathcal{P}|$ we have $r_i + r_{i+1} = d_i + d_{i+1} < \frac{2}{3}|\mathcal{P}| \leq \frac{4}{3}R^*$ and the claim holds.

Case 2b. Otherwise, if $|e_m| \geq \frac{1}{2}|\mathcal{P}|$ then $d_i + d_{i+1} \leq \frac{1}{2}|\mathcal{P}| \leq |e_m|$. Assume $d_i \geq d_{i+1}$. Choose $r_i := d_i$ and $r_{i+1} := |e_m| - d_i$. As $d_{i+1} \leq |e_m| - d_i$ the path \mathcal{P} is covered. Again the existence of an uncovered point contradicts the maximality of \mathcal{P} . Finally, $r_i + r_{i+1} = d_i + (|e_m| - d_i) = |e_m|$, which is the lower bound and thus the range assignment is optimal. \square

If all points of P lie on a straight line, the approximation ratio for two disks can be improved. We first observe an interesting property of optimal solutions in this case.

Lemma 9. Let P be a set of points on a straight line. Then there is a non-overlapping optimal solution, i.e., one in which all disks have disjoint interior.

Proof. An arbitrary optimal solution is modified as follows. For every two overlapping disks C_i and C_{i+1} with centers p_i and p_{i+1} , we decrease r_{i+1} , such that $r_i + r_{i+1} = \text{dist}(p_i, p_{i+1})$, and increase the radius of C_{i+2} by the same amount. This can be iterated, until there is at most one overlap at the outermost disk C_j (with C_{j-1}). Then there must be a point p_{j+1} on the boundary of C_j ; otherwise we could shrink C_j contradicting optimality. Decreasing C_j 's radius r_j by the overlap l and adding a new disk with radius l around p_{j+1} creates an optimal solution without overlap. \square

Theorem 10. Let P be a set of points on a straight line g . Then, a best 2-disk solution for CRA is a $\frac{5}{4}R^*$ -approximation.

Proof. According to Lemma 9, we are, without loss of generality, given an optimal solution R^* with non-overlapping disks. Let p_0 and p_n be the outermost intersection points of the optimal solution disks and g . Without loss of generality, we may further assume $p_0, p_n \in P$ and $R^* := \sum r_i^* = \frac{\text{dist}(p_0, p_n)}{2}$ (otherwise, we can add the outermost intersection point of the outermost disk and g to P , which may only improve the approximation ratio). Let p_i denote the rightmost point in P left to the middle of $\overline{p_0 p_n}$ and let p_{i+1} its neighbor on the other half. Further, let $d_i := \text{dist}(p_0, p_i)$, $d_{i+1} := \text{dist}(p_{i+1}, p_n)$ (see Fig. 9). Assume, $d_i \geq d_{i+1}$. We now give $\frac{5}{4}$ -approximate solutions using one or two disks that cover $\overline{p_0 p_n}$.

Case 1. If $\frac{3}{4}R^* \leq d_i$ then $\frac{5}{4}R^* \geq 2R^* - d_i = \text{dist}(p_i, p_n)$. Thus, the solution consisting of exactly one disk with radius $2R^* - d_i$ centered at p_i is sufficient.

Case 2. If $\frac{3}{4}R^* > d_i \geq d_{i+1}$ we need two disks to cover $\overline{p_0 p_n}$ with $\frac{5}{4}R^*$.

Case 2a. The point p_i could be a center point of an optimal two-disk solution if there was a point p_i^* with $\text{dist}(C_i, p_i^*) = \text{dist}(p_i^*, p_n) = R^* - d_i$. So in case there is a $p_i' \in P$ that lies in a $\frac{1}{4}R^*$ -neighborhood of such an optimal p_i^* we get $\text{dist}(C_i, p_i'), \text{dist}(p_i', p_n) \leq R^* - d_i + \frac{1}{4}R^*$ (see Fig. 10). Thus, $r(p_i) := d_i$, $r(p_i') := R^* - d_i + \frac{1}{4}R^*$ provides a $\frac{5}{4}$ -approximate solution.

Case 2b. Analogously to Case 2a, there is a point $p_{i+1}' \in P$ within a $\frac{1}{4}R^*$ -range of an optimal counterpart to p_{i+1} . Then we can take $r(p_{i+1}) := d_{i+1}$, $r(p_{i+1}') := R^* - d_{i+1} + \frac{1}{4}R^*$ as a $\frac{5}{4}$ -approximate solution.

Case 2c. Assume that there is neither such a p_i' nor such a p_{i+1}' . Since d_i, d_{i+1} are in $(\frac{1}{4}R^*, \frac{3}{4}R^*)$, we have $\frac{1}{4}R^* < R^* - d_j < \frac{3}{4}R^*$ for $j = i, i+1$, which implies that there are two disjoint areas A_i, A_{i+1} , each with diameter equal to $\frac{1}{2}R^*$ and excluding




Fig. 10. A $\frac{5}{4}$ -approximate 2-disk solution with $d_i < \frac{3}{4}R^*$. The cross marks the position of the optimal counterpart p_i^* to p_i and the gray area sketches A_i . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 11. A lower bound of $\frac{5}{4}$ for 2-disk solutions.

all points of P . Since p_i , the rightmost point on the left half of $\overline{p_0 p_n}$, has a greater distance to A_i than to p_0 , any disk around a point on the left could only cover parts of both A_i and A_{i+1} if it has a greater radius than its distance to p_0 . This contradicts the assumption that p_0 is a leftmost point of a disk in an optimal solution. The same applies to the right-hand side. Thus, $A_i \cup A_{i+1}$ must contain at least one point of P , and therefore one of the previous cases leads to a $\frac{5}{4}$ -approximation. \square

Fig. 11 shows that the bound is tight. We believe that this is also the worst case when points are *not* on a line. Indeed, the solutions constructed in the proof of Theorem 10 cover a longest path \mathcal{P} in an optimal solution for a general P . If this longest path consists of at most three edges, $p_i (= p'_{i+1})$ and $p_{i+1} (= p'_i)$ can be chosen as disk centers, covering all of P . However, if \mathcal{P} consists of at least four edges, a solution for the diameter may produce two internal non-adjacent center points that do not necessarily cover all of P .

5. A polynomial-time approximation scheme for CRA

We have seen (Theorem 4, Section 3) that the problem CRA is NP-hard when there are upper bounds given on the possible radii, r_i , of the input points $p_i \in P$, even if these upper bounds are all the same value β . We do not yet know the complexity of exactly solving the problem CRA in the case in which the radii r_i are not bounded above. In this section, we give a polynomial-time approximation scheme (PTAS) for CRA for this case ($\beta = \infty$): for any fixed $\epsilon > 0$, we compute, in polynomial time, a set of radii, r_i , so that the disks of radii r_i centered at the respective input points p_i form a connected set of disks, and the sum of radii $\sum_i r_i$ is at most $(1 + \epsilon)$ times optimal.

To design a PTAS for this problem, we first prove a structure theorem that allows us to apply the m -guillotine method ([22]) and transform any solution into an “ m -guillotine solution” (defined precisely below) with a small increase in the sum of radii. This theorem permits the subsequent use of dynamic programming to obtain a PTAS.

5.1. Structural results

For our first structural result, we need some definitions. Consider a set \mathcal{D} of disks, centered at the points $P = \{p_1, p_2, \dots, p_n\}$. Let $\delta = \text{diam}(P)$ denote the diameter of the point set; i.e., δ is the maximum of the (Euclidean) distances between pairs of points of P . If the radii of the disks \mathcal{D} are all from the discrete set $\mathcal{R} = \{\delta/mn, 2\delta/mn, \dots, \delta\}$, then we will say that \mathcal{D} is a set of $\mathcal{R}_{m,P}$ -disks, or \mathcal{R} -disks, for short, with the understanding that $m = O(1/\epsilon)$ and P will be fixed throughout our discussion. Now consider a set \mathcal{D} of \mathcal{R} -disks, centered on the points P . We let \mathcal{I}_x (resp., \mathcal{I}_y) denote the set of x -coordinates (resp., y -coordinates) of the set of all coordinate-extreme points (leftmost, rightmost, topmost, and bottommost) of the disks \mathcal{D} . Specifically, $\mathcal{I}_x = \{x(p_i) \pm j(\delta/mn) : 1 \leq i \leq n, 0 \leq j \leq mn\}$ and $\mathcal{I}_y = \{y(p_i) \pm j(\delta/mn) : 1 \leq i \leq n, 0 \leq j \leq mn\}$, where $x(\cdot)$ and $y(\cdot)$ denote x - and y -coordinates.

We begin with a simple lemma that shows that we can round up the radii of a feasible solution, to make it a set of \mathcal{R} -disks, at a small cost to the objective function:

Lemma 11. Let $\sum_i r_i$ be the sum of radii in a feasible (connected union) solution, \mathcal{D} . Then, for any fixed $\epsilon > 0$, there exists a set, \mathcal{D}_m , of n disks of radii r'_i centered on points p_i , such that (a). $r'_i \in \mathcal{R} = \{\delta/mn, 2\delta/mn, \dots, \delta\}$, where $\delta = \text{diam}(P)$ is the diameter of the input point set P and $m = \lceil 2/\epsilon \rceil$; and (b). $\sum_i r'_i \leq (1 + \epsilon) \sum_i r_i$.




Fig. 12. An example of an m -guillotine set of disks. The vertical red cut ℓ is the first cut (labeled “1”) in the recursive decomposition of the rectangle ρ . The blue cuts occur in the recursive decomposition; each is labeled with an integer (“2” through “5”) indicating the level in the recursive decomposition. Each cut intersects the interior of at most 2 disks that are interior to the rectangle corresponding to the cut. All cuts are at coordinates \mathcal{I}_x or \mathcal{I}_y . In the final decomposition of ρ , no disk lies interior to a rectangular face (some do lie within a face, in contact with the boundary). Thus, the set of disks is 1-guillotine, with constants $c_1 = 2$, $c_2 = 0$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Proof. Each of the n radii r_i can be increased by at most $\delta/mn \leq \epsilon\delta/2n$ at a total cost (increase in the sum of all radii) of at most $\epsilon\delta/2$. Since increasing the radii of the disks keeps the set of disks connected, and since $\sum_i r_i \geq \delta/2$, we obtain the result. \square

Next, we state and prove a simple observation about the structure of an optimal solution:

Lemma 12. Let \mathcal{D} be a set of circular disks of radii r_i centered at points $p_i \in P$, such that the union of the disks is connected. Then, there exists a set \mathcal{D}' of circular disks centered at points $p_i \in P$, such that the union of the disks \mathcal{D}' is connected, the sum of the radii of disks \mathcal{D}' is at most $\sum_i r_i$, and no point in the plane lies within more than 6 disks of \mathcal{D}' . In particular, in an optimal solution, no point in the plane lies within more than 6 disks of the set of optimal disks. Further, if \mathcal{D} are \mathcal{R} -disks, then there exists a set \mathcal{D}' of \mathcal{R} -disks of total radii $\sum_i r_i$ having the property that no point in the plane lies within more than 6 disks of \mathcal{D}' .

Proof. If some point p lies within more than 6 disks of \mathcal{D} , then we know that two of them, say C_1 and C_2 , have centers, p_1 and p_2 , within a 60-degree cone with apex p . Thus, the distance $|p_1 p_2|$ between the two centers is less than the larger, say r_1 (the radius of C_1), of the two radii (r_1 and $r_2 \leq r_1$), implying that center p_2 lies within the disk C_1 . Thus, if we enlarge the disk C_1 to have radius $r_2 + |p_1 p_2| \leq r_1 + r_2$, the enlarged disk completely covers disk C_2 , allowing C_2 to be deleted (i.e., allowing the radius of C_2 to be shrunk to 0). This modification yields a set \mathcal{D}' of disks that are connected (since the union of the disks of \mathcal{D}' contains the union of the disks of \mathcal{D}), whose total sum of radii is at most that of the set \mathcal{D} .

If the disks \mathcal{D} are \mathcal{R} -disks, then the same argument applies, but we now replace the radius r_1 with the enlarged radius $r_1 + r_2$ ($\geq r_2 + |p_1 p_2|$), while shrinking r_2 to zero. Since the disks \mathcal{D} are \mathcal{R} -disks, we know that $r_1, r_2 \in \mathcal{R} = \{\delta/mn, 2\delta/mn, \dots, \delta\}$, and therefore $r_1 + r_2 \in \mathcal{R}$. Thus, disks \mathcal{D}' are \mathcal{R} -disks. \square

Our main structural result requires further definitions. Let \mathcal{D} be a set of n disks; we consider the disks to be *closed* (include their bounding circles). Let ρ be an axis-parallel rectangle. An axis-parallel line ℓ that intersects ρ is said to be a *cut* of ρ . A cut ℓ is said to be m -perfect with respect to \mathcal{D} and the rectangle ρ if ℓ intersects the interiors of at most $c_1 m + c_2$ disks of \mathcal{D} that lie interior to ρ , for fixed constants (independent of m, ϵ) c_1 and c_2 .

We say that a set \mathcal{D} of \mathcal{R} -disks is m -guillotine with respect to (axis-aligned) rectangle ρ if either (1) no disk of \mathcal{D} lies interior to ρ , or (2) there exists an (axis-parallel) cut ℓ , defined by coordinates \mathcal{I}_x and \mathcal{I}_y , that is m -perfect with respect to \mathcal{D} and ρ , such that \mathcal{D} is m -guillotine with respect to both of the rectangles into which ρ is partitioned by the cut ℓ . We say that a set \mathcal{D} of \mathcal{R} -disks is m -guillotine if \mathcal{D} is m -guillotine with respect to the bounding box, $BB(\mathcal{D})$, of \mathcal{D} . See Fig. 12 for an example of a set of disks that is 1-guillotine.

We are now ready to state and prove our structure theorem, which shows that we can transform an arbitrary set \mathcal{D} of circular disks centered on points P , having a connected union and a sum of radii $R = \sum_i r_i$, into an m -guillotine set of \mathcal{R} -disks, \mathcal{D}_m , having sum of radii at most $(1 + \epsilon)R$ and having a connected union. More specifically, we show:




Fig. 13. An example of an m -span (bold green) for a set of edges E , shown in blue. Here, $m = 4$, and the cut ℓ has $\xi = 14$ endpoints of segments of E along it. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Theorem 13. Let \mathcal{D} be a set of n circular disks of radii r_i centered at points $p_i \in P$, such that the union of the disks is connected. Then, for any fixed $\epsilon > 0$, there exists an m -guillotine set \mathcal{D}_m of at most n \mathcal{R} -disks such that the union of the circular disks \mathcal{D}_m is connected and the sum of the radii of disks of \mathcal{D}_m is at most $(1 + (C/m)) \sum_i r_i$. Here, $m = \lceil 2/\epsilon \rceil$ and C is a constant.

Proof. First, by Lemma 12, we can assume that no point in the plane lies in more than 6 of the input disks.

Consider the set of bounding boxes (squares) of the disks \mathcal{D} . For each disk of radius r_i , we consider the 4 edges of its bounding square, together with the vertical and horizontal diameter segments of the disk. For each disk, the resulting arrangement of horizontal/vertical line segments forms an arrangement (a square with a “+” inside it, partitioning the square into four subsquares) of 12 line segments, each of length r_i . We let E denote this set of $12n$ line axis-parallel line segments. Let $BB(E)$ denote the axis-aligned bounding box of E , which is the same as $BB(\mathcal{D})$, the bounding box of the disks. The edges E have their endpoint coordinates among \mathcal{I}_x and \mathcal{I}_y , and they form a network of horizontal/vertical segments within $BB(E)$, of total length $\lambda(E) = 12 \sum_i r_i$.

We utilize the concept of an “ m -span” from [22]. Let ρ be an axis-aligned rectangle. The intersection, $\ell \cap (E \cap \rho)$, of a horizontal/vertical cut line ℓ with $E \cap \rho$ consists of a discrete (possibly empty) set of line segments and singleton intersection points (which we can consider to be zero-length line segments). Let the endpoints of these segments, be denoted by q_1, \dots, q_ξ , in order along ℓ . We define the m -span, $\sigma_m(\ell)$, of ℓ (with respect to ρ) to be the empty set if $\xi \leq 2(m-1)$, and to be the (possibly zero-length) line segment $q_m q_{\xi-m+1}$ otherwise. Refer to Fig. 13 for an example.

By the m -guillotine charging argument [22], we know that the edge set E can be augmented by a set of m -span horizontal/vertical segments (bridges), of total length $O(\lambda(E)/m) = O((1/m) \sum_i r_i)$, so that the resulting edge set E' is m -guillotine with respect to $BB(E)$ in the usual sense defined in [22]: there exists a horizontal/vertical cut for $BB(E)$ such that the m -span of E' is contained in E' , and, recursively, the set E' is m -guillotine with respect to the subrectangles of ρ on either side of the cut. The base case of the recursion is specified by defining E' to be m -guillotine with respect to rectangle ρ if no vertex (endpoint) of an edge of E' lies interior to ρ .

We say that an axis-parallel segment $ab \subset \ell$ fully crosses a disk D if ab intersects the interior of D and $ab \cap D = \ell \cap D$; i.e., ab crosses two opposite (top/bottom or left/right) semicircles that comprise the boundary of D , so that as one traverses the segment ab , one enters D exactly once and then exits D exactly once. Note that if $ab \subset \ell$ fully crosses D , then the projection of the center of D onto ℓ lies on the segment ab . Note too that if segment ab fully crosses a disk D , then ab properly crosses the (horizontal/vertical) diameter segment of D that is perpendicular to ab .

If $E' = E$ (i.e., no m -span segments had to be added to E , as E is already m -guillotine as an edge set), then the disks \mathcal{D} are already m -guillotine, since the recursive cutting that exhibits that E is m -guillotine as a network also serves to establish that the set \mathcal{D} of disks is m -guillotine as a set of disks. Specifically, if the m -span along a cut ℓ of a rectangle ρ is contained in the edge set E , then there are at most $O(m)$ points within ρ where ℓ crosses edges of E , and, in particular, at most $O(m)$ points where ℓ properly crosses diameter segments of disks. This implies that the segment $\ell \cap \rho$ fully crosses at most $O(m)$ disks, implying that $\ell \cap \rho$ intersects $O(m)$ disks (Lemma 12 implies that each endpoint of $\ell \cap \rho$ intersects at most 6 disks), making ℓ an m -perfect cut (for appropriate choices of constants c_1 and c_2).

If $E' \neq E$, then consider an m -span segment, ab , that was added to E along some cut ℓ in the recursive partitioning of ρ . Let $a'b' = \ell \cap \rho$, with a the endpoint of ab that is closest to a' (i.e., the order along ℓ is (a', a, b, b')). We then distinguish two cases:




Fig. 14. Illustration of the proof of Claim 1. The vertical cut ℓ intersects the rectangle ρ in the segment $a'b'$. The m -span ($m = 4$), ab , is shown in thick green. The disk D_L , centered at p_L , of radius r_L , is shown, as is another disk D , centered at p , with radius r . The claim is that any point $q \in D$ must be covered by an enlarged disk, of radius $r_L + 2|ab|$, centered on p_L . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

- (a) If the m -span ab does not fully cross a disk of \mathcal{D} , then we know that $a'b'$ intersects only $O(m)$ disks of \mathcal{D} , by the following reasoning. By the definition of m -span for E , the segments $a'a$ and bb' each have at most m points along them where the segment crosses (orthogonally) an edge of E corresponding to a diameter segment of a disk; thus, each of $a'a$ and bb' fully cross at most m disks of \mathcal{D} . Since we are assuming that sb does not fully cross any disk, we know that the total number of disks that are fully crossed by $a'b'$ is $O(m)$. Points a' and b' each lie within at most 6 disks, by Lemma 12. Any disk that intersects $a'b'$ must either contain a' or b' or be fully crossed by $a'b'$ (meaning that, when going from a' to b' along the segment, we both enter and leave the disk). Thus, in total there are $O(m)$ disks of \mathcal{D} that are intersected by $a'b' = \ell \cap \rho$, implying that ℓ is m -perfect with respect to \mathcal{D} and ρ .
- (b) If the m -span ab does fully cross at least one disk, then such a disk has the property that the projection of its center point onto ℓ lies on the segment ab . Assume, without loss of generality, that ℓ (and thus ab) is vertical. Then, among disks that ab fully crosses, if there are any with center point to the left of ℓ , let D_L (centered at p_L , with radius r_L) be a disk with leftmost center point p_L . Similarly, if there are fully crossed disks with center point to the right of ℓ , let D_R (centered at p_R , with radius r_R) be a disk with rightmost center point p_R .

Claim 1. *If there are any disks of \mathcal{D} that are fully crossed by ab and have center point to the left of ab , then the disk of radius $r_L + 2|ab|$ centered at p_L covers any disk $D \in \mathcal{D}$, with center p and radius r , that is fully crossed by ab and has p to the left of ab . Similarly, if there are any disks of \mathcal{D} that are fully crossed by ab and have center point to the right of ab , then the disk of radius $r_R + 2|ab|$ centered at p_R covers any disk of \mathcal{D} that is fully crossed by ab and has its center to the right of ab .*

Proof. Assume that there is at least one disk of \mathcal{D} that is fully crossed by ab and has its center to the left of ab . Then, we let disk D_L , with center p_L and radius r_L , be one such disk whose center point is leftmost. We know that the center point p_L projects onto ab ; let h_L be the distance from p_L to ab (i.e., the (horizontal) distance from p_L to its projection on the line ℓ). Let $D \in \mathcal{D}$ be a disk, with center p and radius r , that is fully crossed by ab , with center point p to the left of ab , and let h be the (horizontal) distance from p to ab . By our choice of D_L , we know that $h_L \geq h$. Let $q \in D$ be any point in disk D . Refer to Fig. 14. Then, by the triangle inequality,

$$|p_L q| \leq |p_L p| + |p q|.$$

Also by the triangle inequality, since the points p_L and p project onto ab at two points whose y -coordinates differ by at most $|ab|$, we get

$$|p_L p| \leq |x(p_L) - x(p)| + |y(p_L) - y(p)| \leq (h_L - h) + |ab|,$$

where $x(\cdot)$ and $y(\cdot)$ denote x - and y -coordinates. Thus, using the fact that $|p q| \leq r$ (since $q \in D$), we get that

$$|p_L q| \leq |p_L p| + |p q| \leq (h_L - h) + |ab| + r.$$

Again using the triangle inequality, we know that $|pa| \leq h + |ab|$; thus, $h \geq |pa| - |ab|$. Since disk D is fully crossed by ab , so that a is not contained in D , we know that $|pa| \geq r$, implying that

$$r - h \leq r - (|pa| - |ab|) \leq |ab|.$$

Putting this in the above inequality, we get

$$|p_L q| \leq (h_L - h) + |ab| + r \leq h_L + |ab| + |ab| \leq r_L + 2|ab|,$$

implying that any $q \in D$ is at distance at most $r_L + 2|ab|$ from the center point p_L , as we claimed.

The proof for disks with center points to the right of ab is symmetric. \square

We now modify the set \mathcal{D} of disks as follows. If there are disks fully crossed by ab with center to the left of ab , we select one, D_L , having leftmost center point, and increase its radius to $r_L + 2|ab|$. By Claim 1, we know that this new enlarged disk covers all disks $D_i \in \mathcal{D}$ that were fully crossed by ab , with center points p_i to the left of ab . This implies that we can remove all such disks D_i (i.e., shrink r_i to zero), while maintaining connectivity of the set of disks. If there are disks fully crossed by ab whose center points are to the right of ab , we similarly enlarge one of them, D_R , whose center point p_R is rightmost, while removing the others (shrinking their radii to zero). The net change to the sum of the radii of all disks is that it goes up by at most $O(|ab|)$.

As in case (a) above, we know that the segments $a'a$ and bb' each intersect $O(m)$ disks. After the modification above, we know that ab intersects at most $O(1)$ disks (at most 2 (enlarged) disks of radii $r_L + 2|ab|$ and $r_R + 2|ab|$, and at most 6 disks containing a and 6 disks containing b). Thus, in total $a'b' = \ell \cap \rho$ intersects at most $O(m)$ disks of \mathcal{D} , so we know that ℓ is m -perfect with respect to \mathcal{D} and ρ .

Applying the above modification of the disk radii to all cuts in the recursive set of cuts associated with the m -guillotine edge set E' , we obtain a set of disks \mathcal{D}_m . Since the sum of all m -span lengths $|ab|$ is charged off to the length $\lambda(E) = O(\sum_i r_i)$ of the edge set E , totaling at most $O(1/m)$ times the sum $\sum_i r_i$, we have shown that the sum of the radii of disks of \mathcal{D}_m is at most $(1 + (C/m)) \sum_i r_i$.

Finally, by Lemma 11, we can round up the radii of the resulting disks \mathcal{D}_m to make them \mathcal{R} -disks, while increasing the sum of radii by at most a factor $1 + O(1/m)$. \square

5.2. Dynamic programming

We now give an algorithm to compute a minimum-cost (sum of radii) m -guillotine set of \mathcal{R} -disks whose union is connected. The algorithm is based on dynamic programming. A subproblem is specified by a rectangle, ρ , with x - and y -coordinates among the sets \mathcal{I}_x and \mathcal{I}_y , respectively, of discrete coordinates. The subproblem includes specification of boundary information, for each of the four sides of ρ . Specifically, the boundary information includes:

- (i) $O(m)$ “portal disks”, which are \mathcal{R} -disks intersecting the boundary, $\partial\rho$, of ρ , with at most $O(m)$ disks specified per side of ρ ; and,
- (ii) a connection pattern, specifying which subsets of the portal disks are required to be connected within ρ . (Thus, the connection pattern is specified by giving a partition of the set of portal disks; while the number of partitions is exponential in m , it is constant for fixed $\epsilon = \lceil 2/m \rceil$.)

There are a polynomial number of subproblems (specifically, $n^{O(m)}$), for any fixed m . For a given subproblem, the dynamic program optimizes over all (polynomial number of) possible cuts ℓ (horizontal at \mathcal{I}_y -coordinates or vertical at \mathcal{I}_x -coordinates), and choices of up to $O(m)$ \mathcal{R} -disks intersecting the cut segment $\ell \cap \rho$, along with all possible compatible connection patterns for each side of the cut. (The optimal value of the objective function for each of the two corresponding subproblems on each side of each candidate cut ℓ have been precomputed and tabulated, since we fill in the table of data in order of increasing (combinatorial) size of the rectangles ρ .) The result is an optimal m -guillotine set of \mathcal{R} -disks such that their union is connected and the sum of the radii is minimum possible for m -guillotine sets of \mathcal{R} -disks. Since we know, from the structure theorem, that an optimal set of disks centered at points P can be converted into an m -guillotine set of \mathcal{R} -disks centered at points of P , whose union is connected, and we have computed an optimal such structure, we know that the disks obtained by our dynamic programming algorithm yield an approximation to an optimal set of disks. In summary, we have shown the following result:

Theorem 14. *For any $\epsilon > 0$, there is an approximation algorithm for CRA running in time $n^{O(1/\epsilon)}$ that produces an approximate solution within factor $(1 + \epsilon)$ of the optimal.*

6. Experimental results

It is interesting that even in the worst case, a one-disk solution is close to being optimal. This is supported by experimental evidence. Using the MATLAB programming language, we analyzed the difficulty associated with finding the optimal assignment of radii in randomized trials. In order to generate random problem instances, we considered different numbers of points, uniformly distributed in a two-dimensional circular region. For each trial considering a single distribution

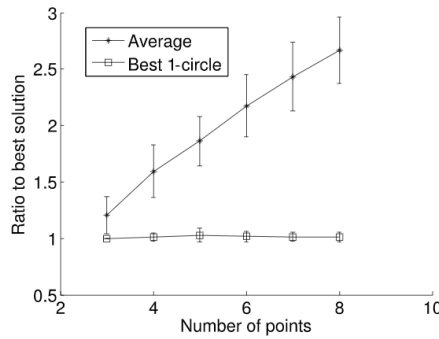


Fig. 15. Ratios of the average over all enumerated trees and of the best 1-disk tree to the optimal $\sum r_i$. Results were averaged over 100 trials for each number.

of points, we enumerated all possible spanning trees using the method described in [3] and recorded the optimal value with the algorithm described in Section 2. Additionally, for each trial, we noted the sum of R values obtained from the best one-disk solution. Fig. 15 shows that the ratio of the best one-disk solution to the best solution remains relatively flat as the number of points increases, as opposed to the average solution, suggesting that the one-disk solution is an excellent heuristic choice. These results were obtained in less than 6 hours using an i7 PC.

7. Conclusion

A number of open problems remain. One of the most puzzling is the issue of complexity in the absence of upper bounds on the radii. The strong performance of the one-disk solution (and even better of solutions with a higher, but limited, number of disks), and the difficulty of constructing solutions for which the one-disk solution is not optimal strongly hint at the possibility of the problem being polynomially solvable. Another indication is that our positive results for one or two disks only needed triangle inequality, i.e., they did not explicitly make use of geometry.

One possible approach to this problem may be to use methods from linear programming. Modeling the objective function and the variables is straightforward; describing the connectivity of a spanning tree by linear cut constraints is also well known. However, even though separating over the exponentially many cut constraints is polynomially solvable (and hence optimizing over the resulting polytope), the overall polytope is not necessarily integral. On the other hand, we have been unable to prove NP-hardness without upper bounds on the radii, even in the more controlled context of graph-induced distances. Note that some results were obtained by means of linear programming: the tight lower bound for 2-disk solutions (shown in Fig. 11) was found by solving appropriate LPs.

Other open problems are concerned with the worst-case performance of heuristics using a bounded number of disks. We showed that two disks suffice for a $\frac{4}{3}$ -approximation in general, and a $\frac{5}{4}$ -approximation on a line; we conjecture that the general performance guarantee can be improved to $\frac{5}{4}$, matching the existing lower bound. Obviously, the same can be studied for k disks, for any fixed k ; at this point, the best lower bounds we have are $\frac{7}{6}$ for $k=3$ and $1 + \frac{1}{2^{k+1}}$ for general k . We also conjecture that the worst-case ratio $f(k)$ of a best k -disk solution approximates the optimal value arbitrarily well for large k , i.e., $\lim_{k \rightarrow \infty} f(k) = 1$.

We have given a PTAS for CRA in two dimensions with unbounded radii, a problem with unknown complexity – there might be a polynomial-time exact solution. Does there exist a PTAS for the general case of CRA with bounded radii, a problem we have shown to be NP-hard? In the conference version of the paper [12], we sketched an approach for how to modify our PTAS for unbounded radii to address a special case of the CRA problem with bounded radii, if an additional assumption is made, that for any segment pq , with p and q within feasible disks, there exists a (connected) path of feasible disks whose sum of radii is $O(|pq|)$. It would be interesting to obtain a PTAS for the CRA for the general case of bounded radii.

Acknowledgments

We thank Ferran for being a great inspiration to all of us. A preliminary version of this work appears in the Algorithms and Data Structures Symposium (WADS), 2011 [12]. This work was started during the 2009 McGill/INRIA/University of Victoria Bellairs Workshop on Computational Geometry. We thank all other participants for contributing to the great atmosphere. This work has been partially supported by the National Science Foundation (grants CCF-1054779 and IIS-1319573, Erin Chambers; grants CCF-1018388 and CCF-1526406, Joseph Mitchell), by the Binational Science Foundation (BSF 2010074, Joseph Mitchell), Sandia National Labs (Joseph Mitchell), an NSERC Postdoctoral Fellowship (Dimitri Marinakis), and three individual NSERC Discovery grants (one each for Venkatesh Srinivasan, Ulrike Stege, and Sue Whitesides). We thank two anonymous reviewers who helped to improve the presentation of this paper.

References

- [1] M. Abellanas, A. García Olaverri, F. Hurtado, J. Tejel, J. Urrutia, Augmenting the connectivity of geometric graphs, *Comput. Geom.* 40 (3) (2008) 220–230.
- [2] H. Alt, E.M. Arkin, H. Brönnimann, J. Erickson, S.P. Fekete, C. Knauer, J. Lenchner, J.S.B. Mitchell, K. Whittlesey, Minimum-cost coverage of point sets by disks, in: *Proc. 22nd ACM Symp. Comp. Geom. (SoCG)*, 2006, pp. 449–458.
- [3] D. Avis, K. Fukuda, Reverse search for enumeration, *Discrete Appl. Math.* 65 (1–3) (1996) 21–46.
- [4] C. Baur, S.P. Fekete, Approximation of geometric dispersion problems, *Algorithmica* 30 (3) (2001) 451–470.
- [5] V. Bilò, I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, Geometric clustering to minimize the sum of cluster sizes, in: *Proc. 13th European Symp. Algorithms*, in: *Lect. Notes Comput. Sci.*, vol. 3669, 2005, pp. 460–471.
- [6] G. Calinescu, I.I. Mandoiu, P.-J. Wan, A. Zelikovsky, Selecting forwarding neighbors in wireless ad hoc networks, *Mob. Netw. Appl.* 9 (2004) 101–111.
- [7] G. Calinescu, P.J. Wan, Range assignment for high connectivity in wireless ad hoc networks, in: *Proc. 2nd International Conference on AD-HOC Networks and Wireless (AdHoc-NOW'03)*, 2003, pp. 235–246.
- [8] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, Energy-efficient wireless network design, *Theory Comput. Syst.* 39 (2006) 593–617.
- [9] P. Carmi, M.J. Katz, Power assignment in radio networks with two power levels, in: *Proc. 9th Scandinavian Workshop on Algorithm Theory*, 2004, pp. 431–441.
- [10] P. Carmi, M.J. Katz, J.S.B. Mitchell, The minimum-area spanning tree problem, *Comput. Geom. Theory Appl.* 35 (October 2006) 218–225.
- [11] P. Carmi, M.J. Katz, M. Segal, H. Shpungin, Fault-tolerant power assignment and backbone in wireless networks, in: *Proc. IEEE Internat. Workshop on Foundations and Algorithms for Wireless Networking*, 2006, pp. 80–84.
- [12] E. Chambers, S.P. Fekete, H.-F. Hoffman, D. Marinakis, J.S.B. Mitchell, V. Srinivasan, U. Stege, S. Whitesides, Connecting a set of circles with minimum sum of radii, in: *Algorithms and Data Structures – 12th International Symposium (WADS2001)*, in: *Lect. Notes Comput. Sci.*, vol. 6844, Springer, 2011, pp. 183–194.
- [13] M. Charikar, R. Panigrahy, Clustering to minimize the sum of cluster diameters, *J. Comput. Syst. Sci.* 68 (March 2004) 417–441.
- [14] A.E. Clementi, P. Penna, R. Silvestri, On the power assignment problem in radio networks, *Mob. Netw. Appl.* 9 (2) (2004) 125–140.
- [15] S. Doddi, M.V. Marathe, S.S. Ravi, D.S. Taylor, P. Widmayer, Approximation algorithms for clustering to minimize the sum of diameters, *Nord. J. Comput.* 7 (September 2000) 185–203.
- [16] B. Fuchs, On the hardness of range assignment problems, *Networks* 52 (4) (2008) 183–195.
- [17] M. Gibson, G. Kanade, E. Krohn, I.A. Pirwani, K. Varadarajan, On clustering to minimize the sum of radii, in: *Proc. 19th ACM-SIAM Symp. Disc. Alg. (SODA)*, 2008, pp. 819–825.
- [18] N. Lev-Tov, D. Peleg, Polynomial time approximation schemes for base station coverage with minimum total radii, *Comput. Netw.* 47 (4) (2005) 489–501.
- [19] D. Lichtenstein, Planar formulae and their uses, *SIAM J. Comput.* 11 (2) (1982) 329–343.
- [20] E.L. Lloyd, R. Liu, M.V. Marathe, R. Ramanathan, S.S. Ravi, Algorithmic aspects of topology control problems for ad hoc networks, *Mob. Netw. Appl.* 10 (2005) 19–34.
- [21] D. Marinakis, D. Meger, I. Rekleitis, G. Dudek, Hybrid inference for sensor network localization using a mobile robot, in: *AAAI National Conference on Artificial Intelligence*, Vancouver, Canada, July 2007, pp. 1089–1094.
- [22] J.S.B. Mitchell, Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, *k*-MST, and related problems, *SIAM J. Comput.* 28 (4) (1999) 1298–1309.
- [23] P.-J. Wan, G. Calinescu, X.-Y. Li, O. Frieder, Minimum-energy broadcast routing in static ad hoc wireless networks, *Wirel. Netw.* 8 (2002) 607–617.