Connecting a Set of Circles with Minimum Sum of Radii

Erin Wolf Chambers[1], Sándor P. Fekete[2], Hella-Franziska Hoffmann[2], Dimitri Marinakis[34], Joseph S.B. Mitchell[5], Venkatesh Srinivasan[4], Ulrike Stege[4], and Sue Whitesides[4]

[1] Department of Computer Science, Saint Louis University, USA. echambe5@slu.edu
[2] Algorithms Group, TU Braunschweig, Braunschweig, Germany. {s.fekete,h-f.hoffmann}@tu-bs.de
[3] Kinsol Research Inc., Duncan, BC, Canada. dmarinak@kinsolresearch.com
[4] Department of Computer Science, University of Victoria, Victoria, BC, Canada. {sue,stege,venkat}@uvic.ca
[5] Department of Applied Mathematics and Statistics, Stony Brook University, USA. jsbm@ams.stonybrook.edu

Abstract. We consider the problem of assigning radii to a given set of points in the plane, such that the resulting set of circles is connected, and the sum of radii is minimized. We show that the problem is polynomially solvable if a connectivity tree is given. If the connectivity tree is unknown, the problem is NP-hard if there are upper bounds on the radii and open otherwise. We give approximation guarantees for a variety of polynomial-time algorithms, describe upper and lower bounds (which are matching in some of the cases), provide polynomial-time approximation schemes, and conclude with experimental results and open problems.

Keywords: intersection graphs, connectivity problems, NP-hardness problems, approximation, upper and lower bounds

1 Introduction

We consider a natural geometric connectivity problem, arising from assigning ranges to a set of center points. In a general graph setting, we are given a weighted graph $G = (V, E)$. Each vertex $v \in V$ in the graph is assigned a radius r_v, and two vertices v and w are connected by an edge f_{vw} in the connectivity graph $H = (V, F)$, if the shortest-path distance $d(v, w)$ in G does not exceed the sum $r_v + r_w$ of their assigned radii. In a geometric setting, V is given as a set of points $P = \{p_1, \ldots, p_n\}$ in the plane, and the respective radii r_i correspond to circular ranges: two points p_i, p_j have an edge f_{ij} in the connectivity graph, if their circles 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 circles to some exponent α. The cases $\alpha = 2$ or 3 correspond to minimizing the overall power; an example for the case $\alpha = 1$ arises from scanning the corresponding ranges with a minimum required angular resolution, so that the scan time for each circle corresponds to its perimeter, and thus radius.

In the context of clustering, Doddi et al. [7], Charikar and Panigraphy [5], and Gibson et al. [9] consider the following problems. Given a set P of n points in a metric space, metric $d(i, j)$ and an integer k, partition P into a set of at most k clusters with minimum sum of a) cluster diameters, b) cluster radii. Thus, the most significant difference to our problem is the lack of a connectivity constraint. Doddi et al. [7] 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, transforming an instance into a min-cost set-cover problem instance yields 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 [5]. In a more geometric setting, Bilò et al. [3] provide approximation schemes for clustering problems.

Alt *et al.* [1] consider the closely related problem of selecting circle centers and radii such that a given set of points in the plane are covered by the circles. 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 circle coverage problem (with $\alpha = 1$) is also considered by Lev-Tov and Peleg [10] in the context of radio networks. Again, connectivity is not a requirement.

The work of Clementi *et al.* [6] focuses on connectivity. It 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 [8] for a collection of hardness results of different (directed) communication graphs.

Carmi et al. [4] 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 circles.

In this paper we present a variety of algorithmic aspects of the problem. In Section 2 we show that for a given connectivity tree, an optimal solution can be computed efficiently. Section 3 sketches a proof of NP-hardness for the problem when there is an upper bound on the radii. Section 4 provides a number of approximation results in case there is no upper bound on the radii. In Section 5 we present a PTAS for the general case, complemented by experiments in Section 6. A concluding discussion with open problems is provided in Section 7.

2 CRA for a Given Connectivity Tree

For a given connectivity tree, our problem is polynomially solvable, based on the following observation.

Lemma 1. *Given a connectivity tree T with at least three nodes. There exists an optimal range assignment for T with $r_i = 0$ for all leaves p_i of T.*

Proof. Assume an optimal range assignment for T has a leaf $p_i \in P$ with radius $r_i > 0$. The circle C_i around p_i with radius r_i intersects circle C_j around p_i's parent p_j with radius r_j. Extending C_j to $r_j := \mathrm{dist}(p_i, p_j)$ while setting $r_i := 0$ does not increase the solution value $R = \sum_{p_i \in P} r_i$. □

Direct consequences of Lemma 1 are the following.

Corollary 1. *There is an optimal range assignment satisfying Lemma 1 and $r_j > 0$ for all $p_j \in P$ of height 1 in T (i.e., each p_j is a parent of leaves only).*

Corollary 2. *Consider an optimal range assignment for T satisfying Lemma 1. Further let $p_j \in P$ be of height 1 in T. Then $r_j \geq \max_{p_i \text{ is child of } p_j}\{dist(p_i, p_j)\}$.*

These observations allow a solution by dynamic programming. The idea is to compute the values for subtrees, starting from the leaves. Details are omitted.

Theorem 1. *For a given connectivity tree, CRA is solvable in $O(n)$.*

3 Range Assignment for Bounded Radii

Without a connectivity tree, and assuming an upper bound of ρ on the radii, the problem becomes NP-hard. In this extended abstract, we sketch a proof of NP-hardness for the graph version of the problem; for the geometric version, a suitable embedding (based on PLANAR 3SAT) can be used.

Theorem 2. *With radii bounded by some constant ρ, the problem CRA is NP-hard in weighted graphs.*

See Figure 1 for the basic construction. The proof uses a reduction from 3SAT. Variables are represented by closed "loops" at distance ρ that have two feasible connected solutions: auxiliary points ensure that either the odd or the even points in a loop get radius ρ. (In the figure, those are shown as bold black or white dots. The additional small dots form equilateral triangles with a pair of black and white dots, ensuring that one point of each thick pair needs to be

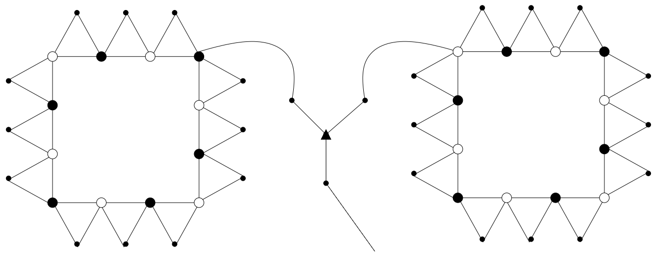

Fig. 1. Two variable gadgets connected to the same clause gadget. “True” and “False” vertices marked in bold white or black; auxiliary vertices are indicated by small dots; the clause vertex is indicated by a triangle. Connectivity edges are not shown.

chosen, so a minimum-cardinality choice consists of all black or all white within a variable.) Additional “connectivity” edges ensure that all variable gadgets are connected. Each clause is represented by a star-shaped set of four points that is covered by one circle of radius ρ from the center point. This circle is connected to the rest of the circles, if and only if one of the variable loop circles intersects it, which is the case if and only if there is a satisfying variable.

4 Solutions with a Bounded Number of Circles

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

Theorem 3. *A best k-circle solution may be off by a factor of $(1 + \frac{1}{2^{k+1}-1})$.*

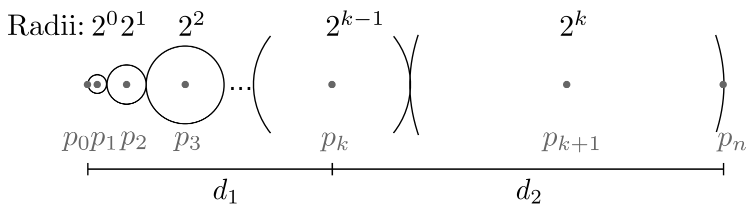

Fig. 2. A class of CRA instances that need $k + 1$ circles in an optimal solution.

Proof. Consider the example in Fig. 2. The provided solution r is optimal, as $R := \sum r_i = \frac{\text{dist}(p_0,p_n)}{2}$. Further, for any integer $k \geq 2$ we have $d_1 = 2 \cdot \sum_{i=0}^{k-2} 2^i + 2^{k-1} < 2 \cdot 2^k + 2^{k-1} = d_2$. So the radius r_{k+1} cannot be changed in an optimal solution. Inductively, we conclude that exactly $k+1$ circles are needed. Because we only consider integer distances, a best k-circle solution has cost $R_k \geq R + 1$, i.e., $\frac{R_k}{R} \geq 1 + \frac{1}{2^{k+1}-1}$. □

In the following we give some good approximation guarantees for CRA using one or two circles.

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

This follows directly from the definition of the connectivity graph which for any edge $e = p_u p_v$ in $\mathcal{P}$ requires $r_u + r_v \geq |e|$.

Theorem 4. *A best* 1*-circle solution for CRA is a* $\frac{3}{2}$*-approximation, even in the graph version of the problem.*

Proof. Consider a longest path $\mathcal{P} = (p_0, \ldots, p_k)$ of length $|\mathcal{P}| = d_{\mathcal{P}}(p_0, \ldots, p_k) := \sum_{i=0}^{k-1} |p_i p_{i+1}|$ 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 2. Let $\bar{d}_i := d_{\mathcal{P}}(p_i, \ldots, p_k)$ and $\bar{d}_{i+1} := d_{\mathcal{P}}(p_0, \ldots, p_{i+1})$. Then $\min\{\bar{d}_i, \bar{d}_{i+1}\} \leq \frac{\bar{d}_i + \bar{d}_{i+1}}{2} = \frac{d_{\mathcal{P}}(p_0,\ldots,p_i) + 2|e_m| + d_{\mathcal{P}}(p_{i+1},\ldots,p_k)}{2} = \frac{|\mathcal{P}|}{2} + \frac{|e_m|}{2} \leq R^* + \frac{R^*}{2} = \frac{3}{2}R^*$. So one circle with radius $\frac{3}{2}R^*$ around the point in P that is nearest to the middle of path $\mathcal{P}$ covers P, as otherwise there would be a longer path. □

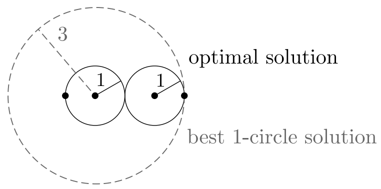

Fig. 3. A lower bound of $\frac{3}{2}$ for 1-circle solutions.

Fig. 3 shows that this bound is tight. Using two circles yields an even better approximation factor.

Theorem 5. *A best* 2*-circle solution for CRA is a* $\frac{4}{3}$*-approximation, even in the graph version of the problem.*

Proof. Let $\mathcal{P} = (p_0, \ldots, p_k)$ be a longest path of length $|\mathcal{P}| = d_{\mathcal{P}}(p_0, \ldots, p_k) := \sum_{i=0}^{k-1} |p_i p_{i+1}|$ in the connectivity graph of an optimal solution with radii r_i^*. Then $R^* := \sum r_i^* \geq \frac{1}{2}|\mathcal{P}|$. We distinguish two cases; see Fig. 4.

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-circle solution that is a $\frac{4}{3}$-approximation.

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

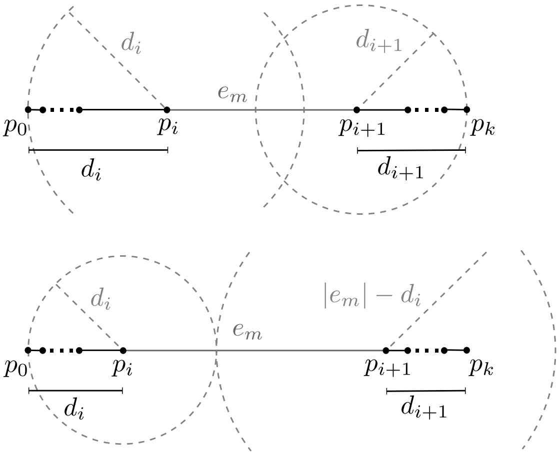

Fig. 4. The two $\frac{4}{3}$-approximate 2-circle solutions constructed in the proof of Theorem 5: (Top) case 2a; (bottom) case 2b.

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. 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 and $r_i + r_{i+1} = d_i + (|e_m| - d_i) = |e_m|$, which is the lower bound and thus the range assignment is optimal. □

If all points of P lie on a straight line, the approximation ratio for two circles can be improved.

Lemma 3. *Let P be a subset of a straight line. Then there is a non-overlapping optimal solution, i.e., one in which all circles have disjoint interior.*

Proof. An arbitrary optimal solution is modified as follows. For every two overlapping circles 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 circle 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 circle with radius l around p_{j+1} creates an optimal solution without overlap. □

Theorem 6. *Let P a subset of a straight line g. Then a best 2-circle solution for CRA is a $\frac{5}{4}$-approximation.*

Proof. According to Lemma 3 we are, w.l.o.g., given an optimal solution with non-overlapping circles. Let p_0 and p_n be the outermost intersection points of the

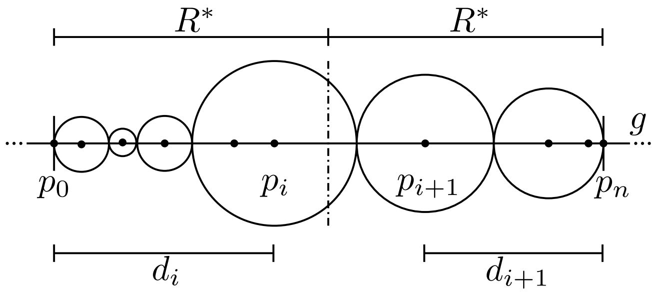

Fig. 5. A non-overlapping optimal solution.

optimal solution circles and g. W.l.o.g., 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 circle 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 $\overline{p_0p_n}$ 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. 5). Assume, $d_i \geq d_{i+1}$. We now give $\frac{5}{4}$-approximate solutions using one or two circles that cover $\overline{p_0p_n}$.

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 circle 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 circles to cover $\overline{p_0p_n}$ with $\frac{5}{4}R^*$.

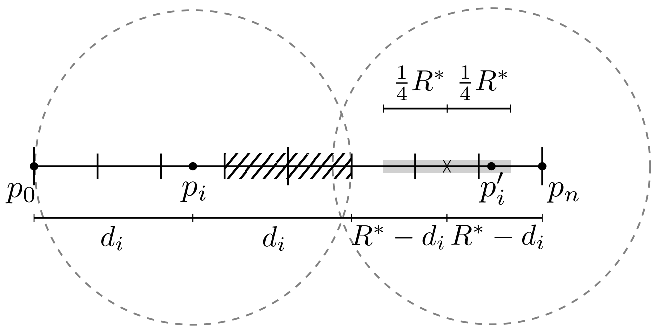

Fig. 6. A $\frac{5}{4}$-approximate 2-circle solution with $d_i < \frac{3}{4}R^*$. The cross marks the position of the optimal counterpart p_i^* to p_i and the grey area sketches A_i.

Case 2a. The point p_i could be a center point of an optimal two-circle 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. 6). 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}. Because 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 all points of P. Because p_i, the rightmost point on the left half of $\overline{p_0p_n}$, has a greater distance to A_i than to p_0, any circle 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 circle 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. □

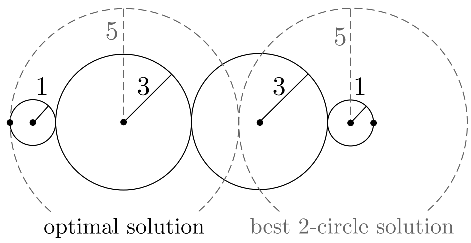

Fig. 7. A lower bound of $\frac{5}{4}$ for 2-circle solutions.

Fig. 7 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 6 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 circle 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 Polynomial-Time Approximation Schemes

We now consider the problem in which each of the n points of $P = \{p_1, \ldots, p_n\}$ has an associated upper bound, $\bar{r}_i$, on the radius r_i that can be assigned to p_i.

5.1 Unbounded Radii

We begin with the case in which $\bar{r}_i = \infty$, for each i. Consider an optimal solution, with radius r_i^* associated with input point p_i. We first prove a structure theorem that allows us to apply the m-guillotine method to obtain a PTAS. The following simple lemma shows that we can round up the radii of an optimal solution, at a small cost to the objective function:

Lemma 4. *Let* $R^* = \sum_i r_i^*$ *be the sum of radii in an optimal solution,* $\mathcal{D}^*$*. Then, for any fixed* $\epsilon > 0$*, there exists a set,* $\mathcal{D}_m$*, of* n *circles of radii* r_i *centered on points* p_i*, such that (a).* $r_i \in \mathcal{R} = \{D/mn, 2D/mn, \ldots, D\}$*, where* D *is the diameter of the input point set* P *and* $m = \lceil 2/\epsilon \rceil$*; and (b).* $\sum_i r_i \leq (1+\epsilon)R^*$*.*

Proof. Each of the n radii r_i^* can be increased by at most $D/mn \leq \epsilon D/2n$ at a total cost of at most $\epsilon D/2$. Since increasing the radii of the circles keeps the set of circles connected and since $R^* \geq D/2$, we obtain the result. □

Disks centered at the points P of radii in the set $\mathcal{R} = \{D/mn, 2D/mn, \ldots, D\}$ will be referred to as $\mathcal{R}_{\epsilon,P}$*-circles*, or $\mathcal{R}$*-circles*, for short, with the understanding that ϵ and P will be fixed throughout our discussion. Consider the arrangement of all $\mathcal{R}$-circles. We let $\mathcal{I}_x$ (resp., $\mathcal{I}_y$) denote the x-coordinates of the left/right (resp., y-coordinates of the top/bottom) extreme points of these circles. (Specifically, $\mathcal{I}_x = \{x_{p_i} \pm j(D/mn) : 1 \leq i \leq n, 0 \leq j \leq mn\bar{r}_i/D\}$ and $\mathcal{I}_y = \{y_{p_i} \pm j(D/mn) : 1 \leq i \leq n, 0 \leq j \leq mn\bar{r}_i/D\}$.)

We say that a set $\mathcal{D}$ of n $\mathcal{R}$-circles is m*-guillotine* if the bounding box, $BB(\mathcal{D})$, of $\mathcal{D}$ can be recursively partitioned into a rectangular subdivision by axis-parallel "m-perfect cuts" that are defined by coordinates $\mathcal{I}_x$ and $\mathcal{I}_y$, with the finest subdivision consisting of a partition into rectangular faces each of which has no circle of $\mathcal{D}$ strictly interior to it. An axis-parallel cut line ℓ is m*-perfect* with respect to $\mathcal{D}$ and a rectangle ρ if ℓ intersects at most $2m$ circles of $\mathcal{D}$ that have a nonempty intersection with ρ.

Key to our method is a structure theorem, which shows that we can transform an arbitrary set $\mathcal{D}$ of circles centered on points P, having a connected union and a sum of radii R, into an m-guillotine set of $\mathcal{R}$-circles, $\mathcal{D}_m$, having sum of radii at most $(1+\epsilon)R^*$. More specifically, we show (proof deferred to the full paper):

Theorem 7. *Let* $\mathcal{D}$ *be a set of circles of radii* r_i *centered at points* $p_i \in P$*, such that the union of the circles is connected. Then, for any fixed* $\epsilon > 0$*, there exists an* m*-guillotine set* $\mathcal{D}_m$ *of* n $\mathcal{R}$*-circles such that the union of the circles* $\mathcal{D}_m$ *is connected and the sum of the radii of circles of* $\mathcal{D}_m$ *is at most* $(1+(C/m))\sum_i r_i$*. Here,* $m = \lceil 1/\epsilon \rceil$ *and* C *is a constant.*

Proof. First, by Lemma 4, we can afford to round up the radii of the circles $\mathcal{D}$ to make them $\mathcal{R}$-circles.

Now, let ρ denote the axis-aligned bounding box of the resulting circles. If there are no circles strictly interior to ρ (i.e., if all circles intersecting ρ intersect (touch) the boundary, $\partial\rho$, of ρ), then we are done: the set of circles is trivially m-guillotine already. Thus, we assume that there is at least one circle of $\mathcal{D}$ strictly interior to ρ.

If there exists an m-perfect cut of ρ, by a horizontal or vertical line intersecting at most $2m$ circles within ρ and intersecting at least one circle that lies interior to ρ, then we partition ρ with it and recurse on the two boxes on each side of the cut. We observe that if such an m-perfect cut exists, then one exists with the property that the defining coordinate of the cut is from the set $\mathcal{I}_x$ (for vertical cuts) or $\mathcal{I}_y$ (for horizontal cuts), since we can translate the cut between two consecutive such coordinates without changing the set of circles it intersects. (More precisely, we can translate a cut to be infinitesimally close to one of the discrete coordinates without changing its combinatorial type. To address this technicality, we can either augment the sets $\mathcal{I}_x$ and $\mathcal{I}_y$ with the midpoints of the intervals between consecutive coordinates, or we can replace each coordinate z with three coordinates – z and z^+ (infinitesimally greater than z) and z^- (infinitesimally smaller than z).)

Thus, we now assume that no m-perfect cut exists for partitioning ρ, and that ρ contains in its interior at least one circle. For a vertical line, ℓ_x, through coordinate x, let $f(x)$ denote the length of the *m-span* of ℓ_x with respect to $\mathcal{D}$ and ρ: $f(x) = 0$ if, within ρ, ℓ_x intersects at most $2m$ circles of $\mathcal{D}$; otherwise, if ℓ_x intersects $K > 2m$ circles of $\mathcal{D}$ within ρ, then $f(x)$ is the distance (along ℓ_x) from the first point, a_m, where ℓ_x enters into the mth circle, going from the top boundary of ρ downwards along ℓ_x, to the first point, b_m, where ℓ_x enters into the mth circle, going from the bottom boundary of ρ upwards along ℓ_x. Because $K > 2m$, we know that a_m is above b_m and that there must be at least one circle of $\mathcal{D}$ whose intersection with ℓ_x is a proper subset of the *bridge segment* $a_m b_m$. We similarly define $g(y)$ to be the length of the (horizontal) bridge segment along a horizontal cut ℓ_y through coordinate y.

We think of $f(x)$ and $g(y)$ as the cost of augmentation for the network $\mathcal{N}_\rho$ consisting of the union of circles, truncated within ρ, that are the boundaries of the circles $\mathcal{D}$; by adding segments (bridges) of length $f(x)$ (resp., $g(y)$), a vertical cut ℓ_x (resp., horizontal cut ℓ_y) can be made m-perfect. We claim that we can charge off the lengths of the bridges that would suffice to augment the network $\mathcal{N}_\rho$ to make it m-guillotine, in the usual sense of an m-guillotine network (subdivision), as in [11]. Specifically, we argue that we can select cuts for which the bridge lengths (m-spans) can be charged off to the total length of the network (sum of the circle circumferences, which is $O(\sum_i r_i)$), showing that the sum of the bridge lengths is at most $(C/m)\sum_i r_i$.

We partition each circle (bounding the circles $\mathcal{D}$) into four 90-degree arcs: two "vertical arcs" (with angular ranges (-45,45) and (135,225)) and two "horizontal arcs" (with angular ranges (45,135) and (225,315)). We define the "chargeable length" of a vertical cut ℓ_x to be the "m-dark" length of $\ell_x \cap \rho$. Specifically, a subsegment ab of $\ell_x \cap \rho$ is said to be m-dark with respect to $\mathcal{N}$ if for any $p \in ab$, the rightwards and leftwards rays from p each cross at least m vertical arcs of $\mathcal{N}_\rho$ before exiting ρ. If we cut ρ along ℓ_x, then the m-dark portion of the cut can be charged off to the left/right sides of the vertical arcs of $\mathcal{N}_\rho$ lying to the right/left of ℓ_x, distributing the charge to be $(1/m)$th to each of the m arcs first hit.

Our charging scheme is based on the observation, following the method of [11], that there must exist a "favorable" vertical cut ℓ_x or horizontal cut ℓ_y for ρ such that the chargeable length of the cut is at least as long as the cost ($f(x)$ or $g(y)$) of the cut. The existence of a favorable cut follows from the observation that $\int_{x\in\rho} f(x)dx = \int_{y\in\rho} h(y)dy$, where $h(y)$ is the chargeable length associated with the horizontal cut ℓ_y; thus, assuming, without loss of generality, that $\int_{x\in\rho} f(x)dx \geq \int_{y\in\rho} g(y)dy$, we see that there must exist a value y^* where $g(y^*) \leq h(y^*)$, which defines a favorable horizontal cut ℓ_{y^*} for which the chargeable length exceeds the length of the m-span. (In case $\int_{x\in\rho} f(x)dx < \int_{y\in\rho} g(y)dy$, there exists a favorable vertical cut.) Further, we claim that there must exist a favorable cut corresponding to the coordinate sets $\mathcal{I}_x$, $\mathcal{I}_y$: the chargeable length of cut ℓ_{y^*} does not change as we perturb y^* to the nearest coordinate in the set $\mathcal{I}_y$, while, by convexity of the circular arcs, the length of the m-span associated with a cut is locally minimized at endpoints of the intervals defined by the points of $\mathcal{I}_y$.

Once a favorable cut is found with respect to rectangle ρ, the cut partitions the problem into two subrectangles, and the argument is recursively applied to each. Since each circular arc of length λ is charged for length at most $\lambda/2m$ on each of its two sides, we get that the overall length of all m-spans that are associated with favorable cuts constructed recursively in converting the network $\mathcal{N}_{BB(\mathcal{D})}$ to an m-guillotine network is at most $(C/m)\sum_i r_i$, for a constant C.

Finally, we claim that the circles $\mathcal{D}$ can be converted to an m-guillotine set $\mathcal{D}_m$, having sum of radii at most $(C/m)\sum_i r_i$: Associated with each m-span $a_m b_m$ that is added to the network of circular arcs in order to make the network m-guillotine, we enlarge the radius of the circle defining one of the m-span endpoints (say, a_m), by at most $|a_m b_m|$, so that this circle now covers the entire m-span segment. Since we have enlarged one circle in a set of circles whose union is connected, the union remains connected. Thus, with a total increase in radii of at most $(C/m)\sum_i r_i$, we end up with an m-guillotine set $\mathcal{D}_m$ of circles, proving our structure theorem. □

We now give an algorithm to compute a minimum-cost (sum of radii) m-guillotine set of $\mathcal{R}$-circles 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 circles", which are $\mathcal{R}$-circles intersecting the boundary, $\partial\rho$, of ρ, with at most $2m$ circles specified per side of ρ; and, (ii) a connection pattern, specifying which subsets of the portal circles are required to be connected within ρ. There are a polynomial number of subproblems, for any fixed m. For a given subproblem, the dynamic program optimizes over all (polynomial number of) possible cuts (horizontal at $\mathcal{I}_y$-coordinates or vertical at $\mathcal{I}_x$-coordinates), and choices of up to $2m$ $\mathcal{R}$-circles intersecting the cut bridge, along with all possible compatible connection patterns for each side of the cut. The result is an optimal m-guillotine set of $\mathcal{R}$-circles such that their union is connected and the sum of the radii is minimum possible for m-guillotine sets of $\mathcal{R}$-circles. Since we know, from the structure theorem, that an optimal set of circles centered at points P can be converted into an m-guillotine set of $\mathcal{R}$-circles centered at points of P, whose union is connected, and we have computed an optimal such structure, we know that the circles obtained by our dynamic programming algorithm yield an approximation to an optimal set of circles. In summary, we have shown the following result:

Theorem 8. *There is a PTAS for the min-sum radius connected circle problem with unbounded circle radii.*

5.2 Bounded Radii

We now address the case of bounded radii, in which circle i has a maximum allowable radius, $\bar{r}_i < \infty$. The PTAS given above relied on circle radii being arbitrarily large, so that we could increase the radius of a single circle to cover the entire m-span segment. A different argument is needed for the case of bounded radii.

We obtain a PTAS for the bounded radius case, if we make an additional assumption: that for any segment pq there exists a connected set of circles, centered at points of $p_i \in P$ and having radii $r_i \leq \bar{r}_i$, such that p and q each lie within the union of the circles and the sum of the radii of the circles is $O(|pq|)$.

Here, we only give a sketch of the method, indicating how it differs from the unbounded radius case. The PTAS method proceeds as above in the unbounded radius case, except that we now modify the proof of the structure theorem by replacing each m-span bridge $a_m b_m$ by a shortest connected path of $\mathcal{R}$-circles. We know, from our additional assumption, that the sum of the radii along such a shortest path is $O(|a_m b_m|)$, allowing the charging scheme to proceed as before. The dynamic programming algorithm changes some as well, since now the subproblem specification must include the "bridging circle-path", which is specified only by its first and last circle (those associated with the bridge endpoints a_m and b_m); the path itself, which may have complexity $\Omega(n)$, is implicitly specified, since it is the shortest path (which we can assume to be unique, since we can specify a lexicographic rule to break ties).

In summary, we have

Theorem 9. *There is a PTAS for the min-sum radius connected circle problem with bounded circle radii, assuming that for any segment pq, with p and q within feasible circles, there exists a (connected) path of feasible circles whose radii are $O(|pq|)$.*

6 Experimental Results

It is curious that even in the worst case, a one-circle solution is close to being optimal. This is supported by experimental evidence. In order to generate random problem instances, we considered different numbers of points uniformly distributed in a 2D circular region. For each trial considering a single distribution of points, we enumerated all possible spanning trees using the method described in [2], and recorded the optimal value with the algorithm mentioned in Section 2. This we compared with the best one-circle solution; as shown in Fig. 8, the latter seems to be an excellent heuristic choice. These results were obtained in several hours using an i7 PC.

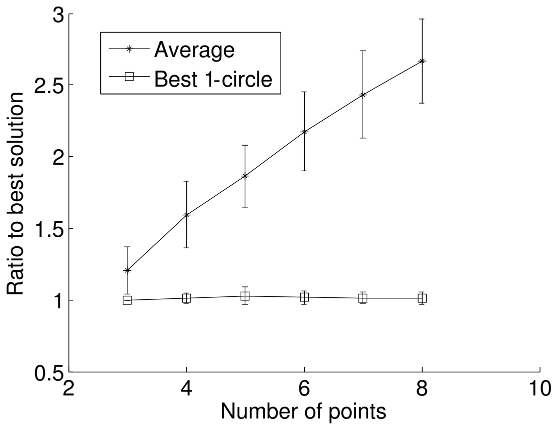

Fig. 8. Ratios of the average over all enumerated trees and of the best 1-circle tree to the optimal $\sum r_i$. Results were averaged over 100 trials for each number.

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-circle solution (and even better of solutions with higher, but limited numbers of circles), and the difficulty of constructing solutions for which the one-circle 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 circles only needed triangle inequality, i.e., they did not explicitly make use of geometry.

One possible way may be to use methods from linear programming: modeling the objective function and the variables by linear methods 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-circle solutions (shown in Fig. 7) was found by solving appropriate LPs.

Other open problems are concerned with the worst-case performance of heuristics using a bounded number of circles. We showed that two circles 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 circles, 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-circle solution approximates the optimal value arbitrarily well for large k, i.e., $\lim_{k\to\infty} f(k) = 1$.

Acknowledgments. A short version of this extended abstract appears in the informal, non-competitive European Workshop on Computational Geometry. This work was started during the 2009 Bellairs Workshop on Computational Geometry. We thank all other participants for contributing to the great atmosphere.

References

1. H. Alt, E. M. Arkin, H. Brönnimann, J. Erickson, S. P. Fekete, C. Knauer, J. Lenchner, J. S. B. Mitchell, and K. Whittlesey. Minimum-cost coverage of point sets by disks. In *Proc. 22nd ACM Symp. Comp. Geom. (SoCG)*, pages 449–458, 2006.
2. D. Avis and K. Fukuda. Reverse search for enumeration. *Disc. Appl. Math.*, 65(1-3):21–46, 1996.
3. V. Bilò, I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Geometric clustering to minimize the sum of cluster sizes. In *In Proc. 13th European Symp. Algorithms, Vol 3669 of LNCS*, pages 460–471, 2005.
4. P. Carmi, M. J. Katz, and J. S. B. Mitchell. The minimum-area spanning tree problem. *Comput. Geom. Theory Appl.*, 35:218–225, October 2006.
5. M. Charikar and R. Panigrahy. Clustering to minimize the sum of cluster diameters. *J. Comput. Syst. Sci.*, 68:417–441, March 2004.
6. A. E. Clementi, P. Penna, and R. Silvestri. On the power assignment problem in radio networks. *Mobile Networks and Applications*, 9(2):125–140, 2004.
7. S. Doddi, M. V. Marathe, S. S. Ravi, D. S. Taylor, and P. Widmayer. Approximation algorithms for clustering to minimize the sum of diameters. *Nordic J. of Computing*, 7:185–203, September 2000.
8. B. Fuchs. On the hardness of range assignment problems. *Networks*, 52(4):183–195, 2008.
9. M. Gibson, G. Kanade, E. Krohn, I. A. Pirwani, and K. Varadarajan. On clustering to minimize the sum of radii. In *Proc. 19th ACM-SIAM Symp. Disc. Alg. (SODA)*, pages 819–825, 2008.
10. N. Lev-Tov and D. Peleg. Polynomial time approximation schemes for base station coverage with minimum total radii. *Computer Networks*, 47(4):489–501, 2005.
11. 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):1298–1309, 1999.