

A Constant-Factor Approximation for Multi-Covering with Disks

Santanu Bhowmick
Department of Computer
Science
University of Iowa
Iowa City, USA
santanu-
bhowmick@uiowa.edu

Kasturi Varadarajan
Department of Computer
Science
University of Iowa
Iowa City, USA
kasturi-
varadarajan@uiowa.edu

Shi-Ke Xue
Department of Computer
Science
Massachusetts Institute of
Technology
Cambridge, USA
shikexue@mit.edu

ABSTRACT

We consider variants of the following multi-covering problem with disks. We are given two point sets Y (servers) and X (clients) in the plane, and a coverage function $\kappa : X \rightarrow \mathbb{N}$. Centered at each server is a single disk whose radius we are free to set. The requirement is that each client $x \in X$ be covered by at least $\kappa(x)$ of the server disks. The objective function we wish to minimize is the sum of the areas of the disks.

We present a polynomial time algorithm for this problem achieving an $O(1)$ approximation.

Categories and Subject Descriptors

F.2.2 [Nonnumerical Algorithms and Problems]: Geometric Problems and Computations

General Terms

Algorithms, Theory

Keywords

Approximation Algorithm, Geometric Set Cover

1. INTRODUCTION

We begin with the statement of the problem studied in this article. We are given two point sets Y (servers) and X (clients) in the plane, and a coverage function $\kappa : X \rightarrow \mathbb{N}$. An assignment $r : Y \rightarrow \mathbb{R}^+$ of radii to the points in Y corresponds to “building” a disk of radius r_y centered at each $y \in Y$. For an integer $j \geq 0$, let us say that a point $x \in X$ is j -covered under the assignment if x is contained in at least j of the disks, i.e.

$$|\{y \in Y \mid \|y - x\|_2 \leq r_y\}| \geq j$$

The goal is to find an assignment that $\kappa(x)$ -covers each point $x \in X$ and minimizes $\pi * \sum_{y \in Y} r_y^2$, the sum of the areas of the disks.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SoCG'13, June 17-20, 2013, Rio de Janeiro, Brazil.

Copyright 2013 ACM 978-1-4503-2031-3/13/06 ...\$15.00.

We call this the *non-uniform minimum-cost multi cover* problem (non-uniform MCMC problem).

Abu-Affash et al [1] consider the version of this problem where $\kappa(x) = k$, $\forall x \in X$, where $k > 0$ is some given integer. We will refer to this as the *uniform MCMC* problem, or following [1], simply as the MCMC problem. We are interested in designing a polynomial time algorithm that outputs a feasible solution (a κ -cover) whose cost is at most some factor $f \geq 1$ times the cost of an optimal solution. We call such an algorithm an f -approximation, and it is implicit that the algorithm is actually polynomial-time.

Related Work.

The (uniform) MCMC problem was considered in two recent papers, motivated by fault-tolerant sensor network design that optimizes energy consumption. Abu-Affash et al. [1] gave an $O(k)$ approximation for the problem using mainly geometric ideas. Bar-Yehuda and Rawitz [3] gave another algorithm that achieves the same approximation factor of $O(k)$, using an analysis based on the local ratio technique. The central question that we investigate in this article is whether an approximation guarantee that is independent of k is possible.

There is a considerable amount of work on clustering and covering problems related to the MCMC problem, and we refer the reader to the previous papers for a detailed survey [1, 3]. Here, we offer a view of some of that work from the standpoint of techniques that may be applicable to the problem at hand. For the case $k = 1$ of the problem, constant factor approximations can be obtained using approaches based on linear programming, and in particular, the primal-dual method [7, 10]. The $O(k)$ approximation of Bar-Yehuda and Rawitz [3] for $k > 1$ can be situated in this line of work.

There has been some recent work on the geometric set multi-covering problem [8, 2]. In particular, the recent work of Bansal and Pruhs [2] addresses the following problem. We are given a set of points in the plane, a set of disks each with an arbitrary non-negative weight, and an integer k . The goal is to pick a subset of the disks so that each of the given points is covered at least k times. The objective function we want to minimize is the sum of the weights of the chosen disks. Bansal and Pruhs [2] give an $O(1)$ approximation for the problem, building on techniques developed for the case $k = 1$ [13, 6].

It would seem that the problem considered in this paper can be reduced to the problem solved by Bansal and Pruhs: for each $y \in Y$ and $x \in X$, add a disk centered at y with radius $\|x - y\|_2$, and let X be the set of points that need to be covered. The reason this reduction does not work is that we have to add an additional

constraint saying that we can use only one disk centered at each $y \in Y$. Notice that this additional constraint is not an issue for the case $k = 1$, since here if the returned solution uses two disks centered at the same $y \in Y$, we can simply discard the smaller one.

In the geometric set cover problems considered by [8, 13, 6, 2], the input disks are “immutable”, and the complexity of the problem stems from the combinatorial geometry of the disks. For the MCMC application, it would be more fruitful to consider geometric set multi-cover problems where the algorithm is allowed to slightly enlarge the input disks. This version of covering with $k = 1$ is considered by Har-Peled and Lee [11]. For $k > 1$, however, we still have the above-mentioned difficulty of reducing MCMC to a set multi-cover problem.

The case $k = 1$ of our MCMC problem actually admits a polynomial time approximation scheme (PTAS) using dynamic programming on top of randomly shifted quad-trees [9, 5]. This was shown by the work of Bilo et al. [4], following the work of Lev-Tov and Peleg [12] for a related problem. The difficulty with extending these results for $k = 1$ to general k is that the “density” of the solution grows with k , and therefore the number of sub-problems that the dynamic program needs to solve becomes exponential in k . It is conceivable that further discretization tricks [11] can be employed to get around this difficulty, but we have not succeeded in this effort. On the other hand, we are also not aware of any hardness result that rules out a PTAS.

Our Results.

Using a simple geometric approach, we obtain an $O(1)$ approximation for the uniform MCMC problem. Thus our main contribution is the demonstration that we can obtain an approximation bound that is independent of k for the uniform MCMC problem.

Our technique is closely related to the geometric approach of Abu-Affash et al. [1] and builds on their ideas. For each $x \in X$, consider the disk $\delta(y^k(x), \|y^k(x) - x\|_2)$ of radius $\|y^k(x) - x\|_2$ centered at $y^k(x)$, the k -th nearest neighbor of x in Y . The work in [1] revolves around a carefully chosen subset of this family of disks. This subset of disks, which are pairwise disjoint, are called *primary disks* by [1]. They show that the sum of the areas of the primary disks is a lower bound on the cost of any solution, and that the cost of the solution output by their algorithm is at most k times this.

Our approach also uses primary disks, but differently. Our algorithm recursively computes a $(k - 1)$ -cover and then extends it to a k cover. We relate the cost of this extension to the cost of the primary disks. To upper bound the cost of our solution, we decompose the optimal k -cover into a $(k - 1)$ -cover and a residual set whose cost upper bounds the cost of the primary disks. We expect that these insights will prove useful, particularly in the context of some non-trivial extensions of our work that we mention in the concluding section.

Our algorithm and approximation guarantee of $O(1)$ works for the non-uniform MCMC problem as well. We therefore present our work in this slightly more general setting.

Organization.

In Section 2, we describe our algorithm for the non-uniform MCMC problem and its analysis. For clarity, the reader may find it useful to keep the uniform MCMC problem in mind during a first perusal. In Section 3, we discuss variants of the problem that our approach addresses. We also mention other intriguing variants that our work leaves unresolved.

2. COMPUTING A COVERING FOR THE NON-UNIFORM MCMC PROBLEM

For convenience, we solve the variant of the non-uniform MCMC problem where we have l_∞ disks rather than l_2 disks. Our input is two point sets Y and X in \mathbb{R}^2 and a coverage function $\kappa : X \rightarrow \mathbb{N} \cup \{0\}$. (It will be useful to allow $\kappa(x)$ to be 0 for some $x \in X$.) We also assume that $\kappa(x) \leq |Y|$ for each $x \in X$, for otherwise there is no feasible solution.

We describe an algorithm for assigning a *radius* $r_y \geq 0$ for each $y \in Y$, with the guarantee that for each $x \in X$, there are at least $\kappa(x)$ points $y \in Y$ such that the l_∞ disk of radius r_y centered at y contains x . In other words the guarantee is that for each $x \in X$,

$$|\{y \in Y \mid \|x - y\|_\infty \leq r_y\}| \geq \kappa(x)$$

Our objective is to minimize the sum of the areas of the disks, that is, $4 * \sum_{y \in Y} r_y^2$. For this optimization problem, we will show that our algorithm outputs an $O(1)$ approximation. Clearly, this also gives an $O(1)$ approximation for the original problem, where distances are measured in the l_2 norm.

Terminology.

Before proceeding, we define some of the terms that will be used extensively in the remainder of this section. We will use $\|\cdot\|$ to denote the l_∞ norm.

For each $x \in X$, fix an ordering of the points in Y that is non-decreasing in terms of l_∞ distance to x . For $1 \leq j \leq |Y|$, let $y^j(x)$ denote the j -th point in this ordering. In other words, $y^j(x)$ is the j -th closest point in Y to x . For brevity, we denote $y^{\kappa(x)}(x)$ by $y^\kappa(x)$.

Given an assignment of radius r_y to each $y \in Y$, we will say that a point $x \in X$ is *j -covered* if at least j disks cover it, that is,

$$|\{y \in Y \mid \|x - y\| \leq r_y\}| \geq j.$$

We will sometimes say that x is κ -covered to mean that it is $\kappa(x)$ -covered. Similarly, if we have an assignment of radii to each $y \in Y$ such that for a set of points $P \subseteq X$, every point $x \in P$ is covered by at least $\kappa(x)$ disks, we say that P is κ -covered.

Let $\delta(p, r)$ denote the l_∞ disk of radius r centered at p .

2.1 The Algorithm

The procedure $\text{Cover}(X, Y, \kappa)$ computes an assignment of radii so that each point $x \in X$ is $\kappa(x)$ -covered. This algorithm is recursive, and in the base case we have $\kappa(x) = 0$ for each $x \in X$. In the base case, the radius r_y is assigned to 0 for each $y \in Y$. Otherwise, we define

$$\kappa'(x) = \max\{0, \kappa(x) - 1\}, \text{ for each } x \in X,$$

and recursively call $\text{Cover}(X, Y, \kappa')$ to compute an assignment that $\kappa'(x)$ -covers each $x \in X$. We then compute $X' \subseteq X$, the set of points that are not $\kappa(x)$ -covered. The goal of the while-loop is to increase some of the r_y to ensure that each $x \in X'$ is also $\kappa(x)$ -covered.

To do this, we find the point $\bar{x} \in X'$ that maximizes the distance to its $\kappa(\bar{x})$ -th nearest neighbor in Y . We add \bar{x} and some “nearby” points in X' to a cluster $\text{XC}_{\bar{x}}$. More precisely, $\text{XC}_{\bar{x}}$ consists of all $x' \in X'$ such that $\delta(y^{\kappa(\bar{x})}(\bar{x}), \|x - y^{\kappa(\bar{x})}(\bar{x})\|)$ and $\delta(y^{\kappa(x')}(x'), \|x' - y^{\kappa(x')}(x')\|)$ intersect. The set $\text{YC}_{\bar{x}}$ contains, for each $x \in \text{XC}_{\bar{x}}$, the $\kappa(x)$ nearest neighbors of x in Y . For purposes of analysis, we add \bar{x} to a set \bar{X} as well. Figure 1 demonstrates how \bar{x} is chosen, along with corresponding members of $\text{XC}_{\bar{x}}$.

Algorithm 1 Cover(X, Y, κ)

- 1: **if** $\forall x \in X, \kappa(x) = 0$ **then**
- 2: Assign $r_y \leftarrow 0$ for each $y \in Y$, and return.
- 3: Define $\kappa'(x)$ as follows:

$$\forall x \in X, \kappa'(x) = \begin{cases} 0, & \text{if } \kappa(x) = 0 \\ \kappa(x) - 1, & \text{if } \kappa(x) > 0 \end{cases}$$

- 4: Recursively call Cover(X, Y, κ').
- 5: Let $X' = \{x \in X \mid x \text{ is not } \kappa(x)\text{-covered}\}$
- 6: Let $\bar{X} \leftarrow \emptyset$.
- 7: **while** $X' \neq \emptyset$ **do**
- 8: Let $\bar{x} \leftarrow \operatorname{argmax}_{x' \in X'} \|x' - y^\kappa(x')\|$
- 9: $\bar{X} \leftarrow \bar{X} \cup \{\bar{x}\}$.
- 10: Let $\text{XC}_{\bar{x}} \leftarrow \emptyset, \text{YC}_{\bar{x}} \leftarrow \emptyset$.
- 11: **for all** $x' \in X'$ **do**
- 12: **if** $\delta(y^\kappa(\bar{x}), \|x - y^\kappa(\bar{x})\|)$ and $\delta(y^\kappa(x'), \|x' - y^\kappa(x')\|)$ intersect **then**
- 13: $\text{XC}_{\bar{x}} \leftarrow \text{XC}_{\bar{x}} \cup \{x'\}$.
- 14: $\text{YC}_{\bar{x}} \leftarrow \text{YC}_{\bar{x}} \cup \{y^1(x'), y^2(x'), \dots, y^\kappa(x')\}$.
- 15: Let $\text{YC}'_{\bar{x}} \subseteq \text{YC}_{\bar{x}}$ be a set of at most four points such that

$$\bigcap_{y \in \text{YC}'_{\bar{x}}} \delta(y, r_y) = \bigcap_{y \in \text{YC}_{\bar{x}}} \delta(y, r_y).$$

- 16: For each $y \in \text{YC}'_{\bar{x}}$, increase r_y by the smallest amount that ensures $\text{XC}_{\bar{x}} \subseteq \delta(y, r_y)$.
 - 17: Remove from X' any points x that are $\kappa(x)$ -covered.
-




Figure 1: Snapshot of Cover(X, Y, κ) with $\kappa(x) = 1$ for each $x \in X$. At this stage, $r_y = 0$ for each $y \in Y$. The distance between b and its closest server $y^1(b)$ is maximum, hence it is picked as \bar{x} in the first iteration of the while loop. The primary disk $\delta(y^1(b), \|b - y^1(b)\|)$ intersects the disks $\delta(y^1(a), \|a - y^1(a)\|)$ and $\delta(y^1(c), \|c - y^1(c)\|)$, causing a and c to be included in $\text{XC}_{\bar{x}}$ in addition to b .

Next, we identify a set $\text{YC}'_{\bar{x}} \subseteq \text{YC}_{\bar{x}}$ of at most 4 points such that

$$\bigcap_{y \in \text{YC}'_{\bar{x}}} \delta(y, r_y) = \bigcap_{y \in \text{YC}_{\bar{x}}} \delta(y, r_y).$$

Why does such a $\text{YC}'_{\bar{x}}$ exist? If, on the one hand, the intersection of disks $\bigcap_{y \in \text{YC}_{\bar{x}}} \delta(y, r_y)$ is empty, then Helly's Theorem tells us that there are three disks (or maybe even two) whose intersection is empty. On the other hand, if the intersection $\bigcap_{y \in \text{YC}_{\bar{x}}} \delta(y, r_y)$ is non-empty, then it is a rectangle (as these are l_∞ disks) and therefore equal to the intersection of four of the disks.

We enlarge the radius r_y of each $y \in \text{YC}'_{\bar{x}}$ by the minimum amount needed to ensure that $\text{XC}_{\bar{x}} \subseteq \delta(y, r_y)$. We argue that after this each point in $\text{XC}_{\bar{x}}$ is κ -covered. To see why, consider any $x' \in \text{XC}_{\bar{x}}$. Notice that $|\text{YC}_{\bar{x}}| \geq \kappa(x')$, since the $\kappa(x')$ nearest neighbors of x' are included in $\text{YC}_{\bar{x}}$. Thus before the enlargement, x' does not belong to $\bigcap_{y \in \text{YC}_{\bar{x}}} \delta(y, r_y)$. (Recall that no point in $\text{XC}_{\bar{x}}$ was κ -covered.) Therefore, x' does not belong to $\bigcap_{y \in \text{YC}'_{\bar{x}}} \delta(y, r_y)$. It follows that there is at least one $y \in \text{YC}'_{\bar{x}}$ such that $\delta(y, r_y)$ did not contain x' before the enlargement. As a consequence of the enlargement, $\delta(y, r_y)$ does contain x' . Since x' was $(\kappa(x') - 1)$ -covered before the enlargement, it is now $\kappa(x')$ -covered.

After increasing r_y for $y \in \text{YC}'_{\bar{x}}$ as stated, we discard from X' all points that are now κ -covered. The discarded set contains $\text{XC}_{\bar{x}}$ which has at least the one point \bar{x} ; so the iteration of the while loop makes progress. We go back and iterate the while loop with the new X' .

It is clear that when Cover(X, Y, κ) terminates, each point $x \in X$ is $\kappa(x)$ -covered.

2.2 Approximation Ratio

To bound the approximation ratio, we begin by obtaining a handle on the increase in the cost (the objective function) in going from a κ' -cover to a κ -cover. For this, we need to focus on the set \bar{X} , to which the while loop adds the point \bar{x} that it chooses in each iteration.

LEMMA 1. *The increase in the objective function $4 * \sum_{y \in Y} r_y^2$*

from the time $\text{Cover}(X, Y, \kappa')$ completes to the time $\text{Cover}(X, Y, \kappa)$ completes is $O(\sum_{x' \in \bar{X}} \|x' - y^\kappa(x')\|^2)$.

PROOF. Let us fix an $\bar{x} \in \bar{X}$, and focus on the iteration when \bar{x} was added to \bar{X} . Notice that there is exactly one such iteration, since \bar{x} is removed from X' in the iteration it gets added to \bar{X} .

We will argue that the increase in cost during that iteration is $O(\|\bar{x} - y^\kappa(\bar{x})\|^2)$, thus completing the proof. The increase in cost is due to the increase in the radii of r_y for $y \in \text{YC}'_{\bar{x}}$, as no other radius is increased. For each such y , it suffices to show that if r_y is increased during that iteration, then the increase in $4r_y^2$ is $O(\|\bar{x} - y^\kappa(\bar{x})\|^2)$.

To show this, we make two claims.

CLAIM 1. For any $x' \in \text{XC}_{\bar{x}}$, we have

$$\|y^\kappa(\bar{x}) - x'\| \leq 3 * \|y^\kappa(\bar{x}) - \bar{x}\|$$

PROOF. Recall that x' is in $\text{XC}_{\bar{x}}$ because the disks $\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|)$ and $\delta(y^\kappa(x'), \|x' - y^\kappa(x')\|)$ intersect. Since \bar{x} is the point in X' that maximizes the distance to the $\kappa(\bar{x})$ -th nearest point in Y , we have $\|x' - y^\kappa(x')\| \leq \|\bar{x} - y^\kappa(\bar{x})\|$. Thus,

$$\begin{aligned} \|y^\kappa(\bar{x}) - y^\kappa(x')\| &\leq \|x' - y^\kappa(x')\| + \|x - y^\kappa(\bar{x})\| \\ &\leq 2 * \|\bar{x} - y^\kappa(\bar{x})\| \end{aligned}$$

from which

$$\begin{aligned} \|y^\kappa(\bar{x}) - x'\| &\leq \|y^\kappa(\bar{x}) - y^\kappa(x')\| + \|y^\kappa(x') - x'\| \\ &\leq 3 * \|y^\kappa(\bar{x}) - \bar{x}\| \end{aligned}$$

□

CLAIM 2. For any $y' \in \text{YC}'_{\bar{x}}$, we have

$$\|y^\kappa(\bar{x}) - y'\| \leq 4 * \|y^\kappa(\bar{x}) - \bar{x}\|$$

PROOF. The point y' belongs to $\text{YC}'_{\bar{x}}$ because there is an $x' \in \text{XC}_{\bar{x}}$ for which $y' = y^j(x')$ for some $1 \leq j \leq \kappa(x')$. With this x' , we have

$$\begin{aligned} \|y^\kappa(\bar{x}) - y'\| &\leq \|y^\kappa(\bar{x}) - x'\| + \|x' - y^\kappa(x')\| \\ &\leq 3 * \|y^\kappa(\bar{x}) - \bar{x}\| + \|y^\kappa(\bar{x}) - \bar{x}\|. \end{aligned}$$

□

Fix a $y \in \text{YC}'_{\bar{x}}$. If r_y was increased in this iteration, it now equals $\|y - x'\|$ for some $x' \in \text{XC}_{\bar{x}}$. By the above two claims,

$$\begin{aligned} \|y - x'\| &\leq \|y - y^\kappa(\bar{x})\| + \|y^\kappa(\bar{x}) - x'\| \\ &\leq 7 * \|y^\kappa(\bar{x}) - \bar{x}\|. \end{aligned}$$

Thus the increase in the quantity $4 * \sum_{y \in Y} r_y^2$ (in the iteration of the while loop under consideration) is $O(\|y^\kappa(\bar{x}) - \bar{x}\|^2)$, finishing the proof of the lemma. □

Following [1], let us refer to the set of disks $\{\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|) \mid \bar{x} \in \bar{X}\}$ as the set of *primary disks*. The above lemma has shown that the increase in cost in going from a κ' -cover to a κ -cover is bounded by a constant times the sum of the areas of the primary disks.

We now observe another property of the primary disks, one that is also used by [1].

LEMMA 2. Any two primary disks are disjoint.

PROOF. Consider the iteration when an \bar{x} is added to \bar{X} . At this point, if there is a point $x' \in X'$ such that the primary disk $\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|)$ and the disk $\delta(y^\kappa(x'), \|x' - y^\kappa(x')\|)$ intersect, then x' gets added to $\text{XC}_{\bar{x}}$ and subsequently gets removed from X' . It follows that there is no possibility that x' gets added to \bar{X} in a subsequent iteration. □

We now bound the approximation ratio of our algorithm.

THEOREM 1. Let $r' : Y \rightarrow \mathbb{R}^+$ be any assignment of radii to the points in Y under which each point $x \in X$ is $\kappa(x)$ -covered. Then the cost of the output of $\text{Cover}(X, Y, \kappa)$ is at most $c * 4 * \sum_{y \in Y} r_y'^2$, where $c > 0$ is an absolute constant.

PROOF. Our proof is by induction on $\max_{x \in X} \kappa(x)$. For the base case, where $\kappa(x) = 0$ for each $x \in X$, the claim in the theorem clearly holds for any $c > 0$.

Let $D = \{\delta(y, r_y') \mid y \in Y\}$ be the set of disks corresponding to the assignment r' . We can view such an assignment as a collection of disks, and vice versa. In particular, we say that D κ -covers each point in X . Our proof strategy is to show that there is a subset $D_\kappa \subseteq D$ such that

1. The cost increase incurred by $\text{Cover}(X, Y, \kappa)$ in going from the κ' -cover to the κ -cover is at most c times the sum of the areas of the disks in D_κ .
2. The set of disks $D \setminus D_\kappa$ κ' -covers any point $x \in X$.

By the induction hypothesis, the cost of the κ' -cover computed by $\text{Cover}(X, Y, \kappa')$ is at most c times the sum of the areas of the disks in $D \setminus D_\kappa$. As the increase in cost incurred by $\text{Cover}(X, Y, \kappa)$ in turning the κ' -cover to a κ -cover is at most c times the sum of the areas of the disks in D_κ , the theorem follows.

We first compute D_κ from D , and then show that it has the above two properties. Consider the set of primary disks $\{\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|) \mid \bar{x} \in \bar{X}\}$ computed by our algorithm in the process of going from a κ' -cover to a κ -cover. Let $\text{largest}(\bar{x})$ be the largest disk from D that contains \bar{x} . Since at least $\kappa(\bar{x})$ disks from D contain \bar{x} , we know that the radius of $\text{largest}(\bar{x})$ is at least $\|\bar{x} - y^\kappa(\bar{x})\|$, the radius of the primary disk $\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|)$. Let

$$D'_\kappa = \{\text{largest}(\bar{x}) \mid \bar{x} \in \bar{X}\}.$$

Sort the disks in D'_κ by decreasing (non-increasing) radii. Let $B \leftarrow \emptyset$ initially. For each disk $d \in D'_\kappa$ in the sorted order, performing the following operation: add d to B if d does not intersect any disk already in B .

Let D_κ be the set B at the end of this computation. Since no two disks in D_κ intersect, and D κ -covers any point in X , it follows that $D \setminus D_\kappa$ κ' -covers any point in X . This establishes Property 2 of D_κ .

To show that the cost increase incurred by $\text{Cover}(X, Y, \kappa)$ in going from the κ' -cover to the κ -cover is at most c times the sum of the areas of the disks in D_κ (Property 1), it suffices, by Lemma 1, to show that the sum of the areas of the primary disks is at most c' times the sum of the areas of the disks in D_κ . Here, $c' > 0$ is an absolute constant.

For this, let us *charge* each primary disk $\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|)$ to some disk in D_κ that is (a) at least as large as $\text{largest}(\bar{x})$, and (b) intersects $\text{largest}(\bar{x})$. Notice that such a disk in D_κ does indeed exist.

Fix a disk $\delta(y, r_y') \in D_\kappa$ and consider a primary disk $\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|)$ that is charged to it. Note that $\delta(y, r_y')$ intersects $\text{largest}(\bar{x})$ and is at least as large; and $\text{largest}(\bar{x})$ intersects $\delta(y^\kappa(\bar{x}), \|\bar{x} - y^\kappa(\bar{x})\|)$.




Figure 2: The disk $\delta(y, r'_y)$ in D_κ intersects the disk $\text{largest}(\bar{x})$, which is at least as large as the primary disk $\delta(y^\kappa(\bar{x}), \|y^\kappa(\bar{x}) - \bar{x}\|)$. All 3 disks are contained inside the disk $\delta(y, 5r'_y)$, which is not drawn to scale in this figure.

$\|y^\kappa(\bar{x}) - \bar{x}\|$ and is at least as large. Thus the primary disk $\delta(y^\kappa(\bar{x}), \|y^\kappa(\bar{x}) - \bar{x}\|)$ is contained in the larger disk $\delta(y, 5r'_y)$. This is illustrated in figure 2.

Since the primary disks are pair-wise disjoint (Lemma 2), and the primary disks charged to $\delta(y, r'_y)$ are contained in $\delta(y, 5r'_y)$, it follows that the sum of the areas of the primary disks charged to $\delta(y, r'_y)$ is at most the area of $\delta(y, 5r'_y)$, which equals 25 times the area of $\delta(y, r'_y)$. We conclude that the sum of the areas of the primary disks is at most c' times the sum of the areas of the disks in D_κ , where $c' = 25$.

This establishes Property 1, and completes the proof of the theorem. \square

We conclude with a statement of the main result of this article.

THEOREM 2. *Given point sets X and Y in the plane and $\kappa : X \rightarrow \{1, 2, \dots, |Y|\}$, the algorithm $\text{Cover}(X, Y, \kappa)$ runs in polynomial time and computes a κ -cover of X with cost at most $O(1)$ times that of the optimal κ -cover.*

3. CONCLUSIONS

Our result generalizes to give $O(1)$ approximations for the following variants of MCMC.

1. X and Y are points in \mathbb{R}^d , where d is any constant, and we want to minimize the sum of the d -dimensional volumes of the balls. The constant in the approximation factor has an exponential dependence on d .
2. X and Y are points in the plane, and we want to find an assignment $r : Y \rightarrow \mathbb{R}^+$ under which each point $x \in X$ is $\kappa(x)$ -covered. The quantity we wish to minimize is $\sum_{y \in Y} r_y^\alpha$, where $\alpha \geq 2$ is a constant.

The algorithm and analysis generalize along the expected lines. The only new observation we need is towards the end of the proof of Theorem 1. After observing that the primary

disks charged to $\delta(y, r'_y)$ are contained in $\delta(y, 5r'_y)$, and are disjoint, we conclude that the sum of the squares of the radii of the primary disks charged to $\delta(y, r'_y)$ is at most the square of the radius of $\delta(y, 5r'_y)$. This however implies that the sum of the α -th powers of the radii of the primary disks charged to $\delta(y, r'_y)$ is at most the α -th power of the radius of $\delta(y, 5r'_y)$. Notice that we need $\alpha \geq 2$ for this.

It would be interesting to generalize our result to obtain $O(1)$ approximations for the following variants of the uniform MCMC.

1. X , Y , and k are given as in MCMC, but the objective we wish to minimize is $\sum_{y \in Y} r_y$.
2. X and Y are points in an arbitrary metric space, and we are given an integer k . We want to find an assignment $r : Y \rightarrow \mathbb{R}^+$ under which each point in X is k -covered. The objective we wish to minimize is (say) $\sum_{y \in Y} r_y$.

When $k = 1$, well known algorithms based on the primal dual method give an $O(1)$ approximation [7, 10]. These results and the work reported here leads us to imagine that an $O(1)$ approximation is possible for general $k > 1$.

For both the problems mentioned above, the recent work of [3] gives an $O(k)$ approximation. For the non-uniform versions, their approximation guarantee is $O(\max_{x \in X} \kappa(x))$.

Acknowledgements. This material is based on work supported by the National Science Foundation under Grant No. 0915543.

4. REFERENCES

- [1] K. A. Affash, P. Carmi, M. J. Katz, and G. Morgenstern. Multi cover of a polygon minimizing the sum of areas. *International Journal of Computational Geometry and Applications*, 21(6):685–698, 2011.
- [2] N. Bansal and K. Pruhs. Weighted geometric set multi-cover via quasi-uniform sampling. In *Proceedings of the European Symposium on Algorithms (ESA)*, 2012, 145–156.
- [3] R. Bar-Yehuda and D. Rawitz. A note on multicovering with disks. *Computational Geometry*, 46(3):394–399, 2013.
- [4] V. Bilò, I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Geometric clustering to minimize the sum of cluster sizes. In *Proceedings of the European Symposium on Algorithms (ESA)*, 2005, 460–471.
- [5] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *Journal of Algorithms*, 46(2):178–189, 2003.
- [6] T. M. Chan, E. Grant, J. Könemann, and M. Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2012, 1576–1585.
- [7] M. Charikar and R. Panigrahy. Clustering to minimize the sum of cluster diameters. *Journal of Computer and System Sciences*, 68(2): 417–441, 2004.
- [8] C. Chekuri, K. L. Clarkson, and S. Har-Peled. On the set multi-cover problem in geometric settings. In *Proceedings of the Symposium on Computational Geometry (SoCG)*, 2009, 341–350.
- [9] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM Journal on Computing*, 34(6):1302–1323, 2005.

- [10] A. Freund and D. Rawitz. Combinatorial interpretations of dual fitting and primal fitting. *Proceedings of the International Workshop on Approximation and Online Algorithms (WAOA)*, 2003, 137–150.
- [11] S. Har-Peled and M. Lee. Weighted geometric set cover problems revisited. *Journal of Computational Geometry*, 3(1):65–85, 2012.
- [12] 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.
- [13] K. Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2010, 641–648.