



UvA-DARE (Digital Academic Repository)

Optimization and approximation on systems of geometric objects

van Leeuwen, E.J.

[Link to publication](#)

Citation for published version (APA):

van Leeuwen, E. J. (2009). Optimization and approximation on systems of geometric objects

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Part III

Approximating Geometric Coverage Problems

Overview

One of the most fundamental and best-known optimization problems is Minimum Set Cover. It has many applications, for instance in wireless network planning, as discussed in Section 1.2.2 and below. Given this particular application, it is natural to consider Minimum Set Cover in a geometric setting. Hence this part of the thesis is devoted to the approximability of the geometric version of Minimum Set Cover, as well as of several of its variants.

Problems

For sake of completeness, we start by defining Minimum Set Cover. Throughout this part, let \mathbb{U} be a universe, $\mathcal{P} \subseteq \mathbb{U}$ a set of *elements*, and \mathcal{S} a set of subsets of \mathbb{U} .

Definition III.1 *The minimum set cover problem is to find a smallest set $C \subseteq \mathcal{S}$ such that each element of \mathcal{P} is contained in (covered by) a set in C .*

Minimum Set Cover where $\mathbb{U} = \mathbb{R}^d$ for some $d > 0$ will be called *Geometric Set Cover*. We will mostly discuss the case where $\mathbb{U} = \mathbb{R}^2$ and the sets in \mathcal{S} are induced by simple geometric objects, such as disks or squares.

A first variation of Minimum Set Cover is its weighted case, where the sets in \mathcal{S} are given a weight and we look for a cover of \mathcal{S} of minimum total weight. We can extend further in this direction by assuming that each element u of \mathcal{P} has a *profit* $p(u)$, each set \mathcal{S}_i of \mathcal{S} a *cost* $c(\mathcal{S}_i)$, and that we are given a budget B . This leads to the following problem.

Definition III.2 *The budgeted maximum coverage problem is find some set $C \subseteq \mathcal{S}$ of total cost at most B that maximizes the total profit of the elements covered by C .*

The geometric version is called *Geometric Budgeted Maximum Coverage*.

Two further variants of Minimum Set Cover that we will be particularly interested in are Unique Coverage and Minimum Membership Set Cover. In the former problem, one is given a collection of sets of elements from some universe and aims to select sets that maximize the number of elements contained in precisely one selected set. In the latter problem, the goal is to cover all elements of the universe while minimizing the maximum number of sets in which any element is contained. We discuss and formally define both problems below.

The unique coverage problem was proposed by Demaine et al. [83] and is mainly motivated by wireless network planning. Providers of wireless communication networks provide service for their customers. This can be achieved

by placing a number of base stations that cover customer locations, which is a geometric set cover problem. However, if too many base stations cover a certain customer location, the resulting interference might cause this customer to receive no service at all. Ideally, each customer is serviced by exactly one base station and service is provided to as many customers as possible.

Definition III.3 *Given a set $C \subseteq \mathcal{S}$, an element $u \in \mathcal{P}$ is uniquely covered by C if there is precisely one $s \in C$ containing u . The (maximum) unique coverage problem is to find a set $C \subseteq \mathcal{S}$ that maximizes the number of uniquely covered elements of \mathcal{P} .*

Of course, it might be more profitable to provide service to certain customers. Furthermore, placing a base station is costly and providers generally have a limited budget. This gives rise to the budgeted unique coverage problem.

Definition III.4 *The budgeted unique coverage problem is find a set $C \subseteq \mathcal{S}$ of total cost at most B that maximizes the total profit of the elements uniquely covered by C .*

In practice, mobile devices can distinguish between signals from different base stations. However, this capability is limited and decreases with the number of base stations in range. Demaine et al. [83] model this by *satisfactions* $s_0 = 0, s_1 \geq s_2 \geq \dots \geq 0$, where an element (customer) u yields *satisfaction-modulated profit* $s_i \cdot p(u)$ if it receives service from exactly i base stations.

Definition III.5 *The budgeted low-coverage problem is to find a set $C \subseteq \mathcal{S}$ of total cost at most B that maximizes the satisfaction-modulated profit of the elements covered by C .*

Another way to handle the limited capability of mobile devices to distinguish signals from different base stations, is to minimize the number of signals a device receives.

Definition III.6 *For a set $C \subseteq \mathcal{S}$, the membership $\text{mem}_C(u)$ of an element $u \in \mathcal{P}$ is equal to the number of sets in C that contain u . The maximum membership is $\text{mem}_C(\mathcal{P}) = \max_{u \in \mathcal{P}} \text{mem}_C(u)$. Then the minimum membership set cover problem is to find a set $C \subseteq \mathcal{S}$ that covers all elements in \mathcal{P} and minimizes $\text{mem}_C(\mathcal{P})$.*

As mentioned before, we want to study these problems in their geometric version because of the close connection to wireless network planning. For instance, consider the case when the universe is the plane \mathbb{R}^2 , \mathcal{P} is a set of points corresponding to customer locations, and each $s \in \mathcal{S}$ is a geometric object modeling the broadcasting range of the corresponding base station. If all base stations are equivalent and we ignore obstacles to the signal, these geometric objects are unit disks.

One can make the problem more realistic by assuming that the base stations may have different broadcasting ranges and that they are hindered by obstacles, but the overlap of the broadcasting ranges of the potential base station locations is bounded. The latter assumption is reasonable, as in practice there are usually very few spots where a base station can or may be placed. We model this by a set of fat objects where any point in the plane is overlapped by a bounded number of objects, i.e. a set of fat objects of bounded ply.

Previous Work

Geometric Set Cover

Minimum Set Cover can be approximated within $1 + \ln |\mathcal{S}|$ by a greedy algorithm, even in the weighted case [156, 197, 66]. Hochbaum [149] gives a survey on these and other approximation algorithms for Minimum Set Cover. The greedy algorithm is also optimal. That is, Minimum Set Cover has no polynomial-time algorithm attaining an approximation ratio of $(1 - \epsilon) \ln |\mathcal{S}|$ for any $\epsilon > 0$, unless $\text{NP} \subset \text{DTIME}(n^{O(\log \log n)})$ [108].

Geometric Set Cover is NP-hard on unit squares and on unit disks [110, 157], even if the point set \mathcal{P} corresponds to the centers of the squares or disks. By reducing from Minimum Dominating Set on unit disk/square graphs of bounded density, the NP-hardness continues to hold if the density is 1 [67].

The best approximation ratio for Geometric Set Cover is steadily being improved. Brönnimann and Goodrich [45] gave the first constant-factor approximation algorithm for Geometric Set Cover on unit disks, attaining an unspecified approximation ratio. Călinescu et al. [49] strengthened this result to a 108-approximation algorithm. Narayanappa and Vojtěchovský [218] built on the ideas of this algorithm to give a 72-approximation algorithm. The algorithm by Carmi, Katz, and Lev-Tov [52] has approximation ratio 38.

Recently a ptas for Geometric Set Cover on general disks was discovered. By a simple transformation, this problem is equivalent to the geometric version of *Minimum Hitting Set* (where a smallest subset of points hitting each object must be found) on three-dimensional half-spaces. Mustafa and Ray [217] give a ptas for this problem, as well as for *Geometric Hitting Set* on pseudo-disks.

The above algorithms are not known to be applicable to the weighted case. Ambühl et al. [13] give a 72-approximation algorithm for Weighted Geometric Set Cover on unit disks. A 2-approximation algorithm on unit squares is given by Mihalák [209].

Lev-Tov and Peleg [191] give a ptas for a special case of Geometric Set Cover on unit disks where the set \mathcal{P} is a subset of the set of disk centers. Liao and Hu [193] consider the case where all points of \mathcal{P} lie on the corners of a constant-size grid and give a ptas that extends to the weighted case.

Clarkson and Varadarajan [68] present a constant-factor approximation algorithm on pseudo-disks, attaining an unspecified approximation ratio.

Glaßer, Reitwießner, and Schmitz [125] consider a multi-objective version of Geometric Set Cover on unit disks, thus providing a trade-off between minimizing the number of selected disks and maximizing the number of covered points. They give a polynomial-time approximation scheme, meaning that for any (fixed) $\epsilon > 0$ an ϵ -optimal Pareto curve is output in polynomial time, i.e. for every dominating solution on the Pareto curve, the returned curve has an ϵ -approximate solution. It should be noted that Glaßer, Reitwießner, and Schmitz restrict to the case where the selected disks must have constant precision. Several inapproximability results are also given.

In three dimensions, Laue [188] showed that Geometric Set Cover has a constant-factor approximation algorithm on translated copies of a fixed polytope. The polytope need not be convex or fat.

Budgeted Maximum Coverage has a $(1 - \frac{1}{e})$ -approximation algorithm in both the unit cost [264, 151, 149] and the general case [165]. Khuller, Moss, and Naor [165] proved that no polynomial-time algorithm can obtain an approximation ratio better than $(1 - \frac{1}{e})$, unless $\text{NP} \subset \text{DTIME}(n^{\text{O}(\log \log n)})$.

As far as we know, the budgeted version of Minimum Set Cover has not been considered yet on unit squares or unit disks. Geometric Budgeted Maximum Coverage on unit disks or unit squares, even if the density is 1 and the point set \mathcal{P} corresponds to the centers of the disks or squares, can easily be shown to be NP-hard by reduction from Geometric Set Cover.

Geometric Unique Coverage and Membership Set Cover

Demaine et al. [83] formulated the unique coverage problem and studied it in its general setting. They present a polynomial-time $\Omega(1/\log \rho) = \Omega(1/\log n)$ approximation algorithm using a greedy method, where n is the number of elements and ρ is one plus the ratio of the maximum number and the minimum number of sets in which an element is contained. They give hardness results to show that this algorithm is (near-)optimal. Demaine et al. give a gap-preserving reduction from a variant of Balanced Binary Independent Set, so any (in)approximability result for this problem holds for Unique Coverage as well. Hence, for any $\epsilon > 0$, it is hard to approximate Unique Coverage within ratio $\Omega(1/\log^{\sigma(\epsilon)} n)$, assuming that $\text{NP} \not\subseteq \text{BPTIME}(2^{n^\epsilon})$, where $\sigma(\epsilon)$ is some constant dependent on ϵ . Under the hypothesis that refuting random instances of 3SAT is hard on average, one can strengthen this to $\Omega(1/\log^{1/3-\epsilon} n)$ for any $\epsilon > 0$. By making a (plausible) assumption on the hardness of Balanced Binary Independent Set, a further strengthening to $\Omega(1/\log n)$ is possible.

The *unique hitting set* problem, where one tries to select elements to uniquely hit as many sets as possible, is equivalent to Unique Coverage. If sets have cardinality at most k , Guruswami and Trevisan [134] give an $\Omega(1/\log k)$ -approximation algorithm. They also consider the more general *1-in- k -SAT* problem and give a $1/e$ -approximation algorithm on satisfiable instances.

Moser, Raman, and Sikdar [216] showed that if Unique Coverage is parameterized by the number k of elements to cover uniquely, it is in FPT. However,

Budgeted Unique Coverage parameterized by k and the budget B is not in FPT, unless $P=NP$. If parameterized by B and the profits and costs are integer, the problem is not in FPT unless $FPT=W[1]$.

As far as we know, the unique coverage problem as is and its extensions have not been studied in a geometric setting, although several related problems have. If one tries to maximize the number of points that are uniquely covered subject to the constraint that the selected objects are disjoint, this is a maximum-weight independent set problem, which has a ptas on fat objects of arbitrary ply (see Chapter 7). For a grid-based version of Unique Coverage on unit disks, where the disks are restricted to lie at grid points, a ptas has been announced by Lev-Tov and Peleg [190]. The eptas for fat objects of bounded ply given in Chapter 10 is significantly more general than this result. Recently, Glaßer, Reitwießner, and Schmitz [125] considered a multi-objective version of Geometric Unique Coverage on unit disks, where the selected disks must have constant precision. Their approximation scheme outputs an ϵ -optimal Pareto curve in polynomial time for any fixed $\epsilon > 0$.

Minimum Membership Set Cover has a $O(\ln n)$ -approximation algorithm, but no polynomial-time $(1 - \epsilon) \ln n$ -approximation algorithm for any $\epsilon > 0$, unless $NP \subset DTIME(n^{O(\log \log n)})$ [186]. The minimum membership set cover problem has not been studied yet in its geometric setting.

Further Variations

There are many variations of Geometric Set Cover. For instance, suppose that the point set and the centers of the disks are given, but we are free to choose the radii of the disks. Lev-Tov and Peleg [191] give a ptas that minimizes the sum of the radii. Bilò et al. [34] provide a more general scheme that can also be applied if at most k radii may be nonzero. This problem, called *k-clustering*, is also considered by Alt et al. [11].

The related *k-center* clustering problem, where the maximum radius has to be minimized, is also well-studied and has a ptas in many cases. A ptas for the two-dimensional case was given by Agarwal and Procopiuc [3]. Bádoiu, Har-Peled, and Indyk [20] improved this ptas and generalized to arbitrary dimension. Agarwal and Procopiuc [3] give a nice overview of earlier results.

If the radius of the disks is fixed, but the positions of the centers of the disks may be chosen freely, we obtain the *geometric covering* problem. This problem is NP-hard on unit disks and on unit squares [110, 157] and a ptas is known both on unit disks [150] and on unit squares [128]. The variation where at most k disks may be selected also has a ptas [113].

When we are given both a set of points that should be covered by the objects and a set of points that should not be covered, we obtain a generalization of the geometric covering problem called the *class cover* problem. Cannon and Cowen [50] give a ptas on unit disks. Efrat et al. [96] give an approximation algorithm on ellipses, where the eccentricity of the ellipses may also be chosen freely. The approximation ratio is logarithmic in the optimum.

A closely related problem is *Geometric Piercing*, where given a set of objects, a minimum set of points should be found piercing (hitting) all objects. This problem has a ptas on fat objects [57].

Glaßer, Reith, and Vollmer [124] consider a maximum coverage problem where, for each (unit) disk and each point of \mathcal{P} in the disk, we are given a ‘signal strength’. Then a point of \mathcal{P} is ‘supplied’ if it is covered by a disk whose signal strength for this point is higher than the sum of the signal strengths of other selected disks covering this point. If the given set of unit disks has constant precision, Glaßer, Reith, and Vollmer give a ptas for both the case where at most k disks can be selected and the number of supplied points must be maximized, and the case where at least l points must be supplied and the number of selected disks is minimized.

A problem similar to Geometric Unique Coverage is the problem to find a subset of a given set of objects that maximizes the total area that is uniquely covered. Chen et al. [63] provide a ptas for this problem on unit disks.