How to Water Carrots:
Geometric Coverage Problems for Point Sets

Helmut Alt* Esther M. Arkin† Hervé Brönnimann‡ Jeff Erickson§
Sándor P. Fekete¶ Christian Knauer* Jonathan Lenchner‖ Joseph S. B. Mitchell†
Sue Whitesides** Kim Whittlesey††

The problem. We study a new geometric optimization problem which arises in carrot crop management [10], as well as wireless network design and various other facililty location problems. The task is to select a number of locations t_j for the sprinklers, and assign a radius r_j to each sprinkler. A carrot p_i is covered iff it is within range of some transmission point t_{j_i}, i.e., $d(t_{j_i}, p_i) \leq r_{j_i}$. The resulting cost per sprinkler is some known function f, such as $f(r) = r^{\alpha}$. In the context of wireless network design, sprinklers are base stations, the radius represents the transmission range, and carrots are demand points. We adopt the neutral terminology *server/radius/client* in the rest of the paper. The goal is to minimize the total cost, $\sum_j f(r_j)$, over all placements of at most k servers that cover the set Y of clients.

Our problem is part of a vast family of clustering problem, among which are the *k-center* problem in which one optimizes $\max_j r_j$, the *k-median* in which the cost is $\sum_i d(p_i, t_{j_i})$, and the *k-clustering* which optimizes the sum of all inter-distances between the points in a same cluster. Our problem was named *min-size k-clustering* by Bilò et al. [4], who note that these problems are usually known to be NP-hard and many polynomial-time approximation schemes (PTAS) have been proposed, including for geometric instances of these problems [1, 8]. Clustering for minimizing the sum of radii was studied for points in metric spaces by Charikar and Panigrahy [5], who present a constant approximate solution using at most k clusters.

In the context of watering carrots, the cost function is simply the power of the sprinkler which is a linear function of the radius. When modeling the energy required for wireless transmission, it is common to assume a superlinear ($\alpha > 1$) dependence of the cost on the radius: a quadratic dependence ($\alpha = 2$) models the total area of the served region, and in fact, physically accurate simulation often requires superquadratic dependence ($\alpha > 2$). Nevertheless, a linear dependence ($\alpha = 1$) is often assumed, as in [7]. In addition, there may be special settings where linearity is indeed appropriate [9]: consider for instance a system that beams a narrow angle in a direction which can adapt to the configuration of the network.

In one of the problem variations which we call the *discrete* problem, it is customary to assume a set X of m potential locations for the servers. For the case of a linear cost function f, the problem has been considered recently by Lev-Tov and Peleg [6, 7] (although not for carrots), who give an $O((n+m)^3)$ algorithm when the clients and servers all lie on a given line, and a linear-time 4-approximation with their "closest center" (CC) algorithm. They also give a PTAS for the two-dimensional case when the clients and servers can lie anywhere in the plane. Bilò et al. [4] show that the problem is NP-hard in the plane for the case $\alpha \geq 2$, either when the sets X and Y are given and k is left unspecified ($k = n$), or when k is fixed but then $X = Y$. They observe that the problem is solvable in polynomial time in 1D and give a PTAS in 2D for any $\alpha \geq 1$.

Our contribution. In the 1D case, we extend Lev-Tog and Peleg's results to a linear-time 3-approximation by using a "closest center with growth" (CCG) algorithm and a near-linear-time 2-approximation that uses a "greedy growth" (GG) algorithm. Intuitively, greedy

---

[*] Institute of Computer Science, Freie Universität Berlin, D-14195 Berlin, Germany. `{alt, knauer}@inf.fu-berlin.de`.

[†] Department of Mathematics and Computer Science, Stony Brook University, Stony Brook, NY 11794, USA. `{estie, jsbm}@ams.sunysb.edu`.

[‡] Computer and Information Science, Polytechnic University, Brooklyn, NY 11201, USA. `hbr@poly.edu`.

[§] Department of Computer Science, University of Illinois, Urbana, IL 61801, USA. `jeffe@cs.uiuc.edu`.

[¶] Department of Mathematical Optimization, Braunschweig University of Technology, D-38106 Braunschweig, Germany. `s.fekete@tu-bs.de`.

[‖] IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, USA. `lenchner@us.ibm.com`.

[**] School of Computer Science, McGill University, Montreal, Quebec H3A 2A7, Canada. `sue@cs.mcgill.ca`.

[††] Department of Mathematics, University of Illinois, Urbana, IL 61801, USA. `kwhittle@math.uiuc.edu`.

growth works as follows: start with a disk with center at each server all of radius zero; amongst all clients, find the one which requires the least radial disk growth to capture it; repeat until all clients are covered.

In the 2D case, we generalize the min-size clustering problem in two directions. On the one hand, we consider less restrictive server placement policies. For instance, if we only restrict the servers to lie on a given fixed line, we can give a dynamic programming algorithm that solves the problem exactly and present faster constant factor approximation algorithms. Namely, given n clients in the plane, we can compute in $O(n^2 \log n)$ time a covering by circles (in any fixed L_p metric) centered on the x-axis, such that the sum of the radii is minimized. Handling superlinear non-decreasing cost functions can be done at a runtime cost of $O(n^4 \log n)$.

A practical example where servers are restricted to lie along a line is when the sprinklers are along a water-supplying linear pipe (although, to quote Reid [10], "process carrots are often grown on ridges or on beds with several single rows, whereas table carrots may be grown on beds with several double rows").

If the servers are restricted to lie on a horizontal line, but the location of this line may be chosen freely, then we show that the optimal position is actually not computable by radicals. Our approach is similar to the approach used by Bajaj on the unsolvability of the Fermat-Weber problem and other geometric optimization problems [2, 3].

Nevertheless, we also give constant approximation algorithms for selecting a line whose slope is given but not its position. Namely, our algorithm "square-greedy" (SG) gives in time $O(n \log n)$ a factor 3 approximation to the square covering problem with any linear or superlinear cost function. A small variation "square-greedy with growth" (SGG) gives a 2-approximation for a linear cost function, also in time $O(n \log n)$.

We also give a PTAS for the problem of finding an optimally positioned line and the optimal covering by circles centered on that line such that the sum of the α powers of the radii is minimized. If the line slope is given, there exists a PTAS of runtime $O(n^3 \log n)$ in the linear case ($\alpha = 1$) and $O(n^4 \log n)$ for higher values of α. If the line slope is unconstrained, given $\varepsilon > 0$, we can find in $O(n^5/\varepsilon^2 \log n)$ time a line and a set of discs centered on that line that cover a set of clients at cost at most $(1+\varepsilon)$ times the optimal.

Another generalization we consider combines the problem of finding a short tour and placing servers on it. The cost function we consider is a linear combination of the tour cost and the transmission costs as above. Namely, we consider the problem of placing a set X of k servers along a tour T to cover a set $Y = \{p_1, \dots, p_n\}$ of n clients, and whose cost is $\text{length(T)} + C \sum r_i^{\alpha}$. This problem, which we call *minimum cost covering tour* (MCCT), has not been studied before. The problem arises when a valuable or sensitive resource is distributed. There is a trade-off between the cost of broadcasting from a central location (thus wasting material or risking interception) and the cost of travelling to broadcast more locally, thereby reducing broadcast costs but incurring travel costs.

All our results for MCCT concern the linear cost case. We show that MCCT with linear cost is NP-hard if the ratio C is part of the input. On the one hand, we show that if the relative cost of the tour is much larger than the transmission cost (exactly, if $C \leq 4$), then the optimum solution is a single server placed at the circumcenter. At the other extreme (when C is very small), the optimum solution is a TSP among the clients. For any fixed value of $C > 4$ and linear cost, we present a PTAS for MCCT, namely in time $O(n^{O(1/\varepsilon)})$ we can find a placement of the servers to cover all the clients whose cost, $\text{length(T)} + C \sum r_i$, is within $(1+\varepsilon)$ of the optimal.

References

[1] S. Arora, P. Raghavan, and S. Rao. Approximation schemes for euclidean k-medians and related problems. In *STOC '98: Proc. of the 30th Annual ACM Symp. Theory Comput.*, pp. 106–113, 1998.

[2] C. Bajaj. Proving geometric algorithm non-solvability: An application of factoring polynomials. *J. Symbolic Comput.*, 2:99–102, 1986.

[3] C. Bajaj. The algebraic degree of geometric optimization problems. *Discrete Comput. Geom.*, 3:177–191, 1988.

[4] V. Bilò, I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Geometric clustering to minimize the sum of cluster sizes. In *Proc. 13th European Symp. Algorithms*, 2005.

[5] M. Charikar and R. Panigrahy. Clustering to minimize the sum of cluster diameters. *J. Comput. Syst. Sci.*, 68(2):417–441, 2004.

[6] N. Lev-Tov and D. Peleg. Exact algorithms and approximation schemes for base station placement problems. In *SWAT 2002*, volume 2368 of *Springer LNCS*, pages 90–99, 2002.

[7] N. Lev-Tov and D. Peleg. Polynomial time approximation schemes for base station coverage with minimum total radii. *Computer Networks*, 47:489–501, 2005.

[8] 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.*, 28:1298–1309, 1999.

[9] K. Pahlavan and A. H. Levesque. *Wireless information networks*. Wiley-Interscience, New York, NY, USA, 1995.

[10] J. B. Reid. The carrot calculator: a decision support tool for carrot crop management. Acta Hort. (ISHS) 670:131-136, 2005. `http://www.actahort.org/books/670/670_14.htm`