On the Ball Spanned by Balls

Nimrod Megiddo

IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120-6099, USA, and
School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

Abstract. The procedure for linear programming in linear time in fixed dimension is extended to solve in linear time certain nonlinear problems. Examples are the problem of finding the smallest ball enclosing n given balls, and the weighted-center problem in fixed dimension.

1. Introduction

The problem of finding the smallest circle enclosing n given points in the plane has a long history (see [M1] for references). A linear-time algorithm was given in [M1] and it follows from [M3] that in any fixed dimension the smallest ball enclosing n points can be found in $O(n)$ time since the problem can be formulated as a convex quadratic optimization problem subject to linear constraints. The problem of the smallest enclosing circle is also known as the 1-center problem, where one seeks to locate a point (a "facility") so as to minimize the largest distance between the facility and any of the n given ("demand") points. A known generalization of the latter is the *weighted* 1-center problem, where each of the demand points has a multiplicative weight associated with it and one seeks to minimize the largest weighted distance between the facility and any demand point. Different algorithms for the weighted 1-center problem were proposed in [M2], [MZ], and [Co] until Dyer [D1] proposed an embarrassingly simple trick that solves the problem in linear time in any fixed dimension using the search method proposed in [M3]. In the same paper Dyer also improves the search technique itself obtaining better constants in the linear functions. An analogous improvement was proposed by Clarkson [C1].

Dyer's trick was to reformulate the problem as a linear optimization problem with linear inequalities plus only one quadratic inequality. In this note we consider the problem of the smallest ball enclosing n given balls, where it seems that such a trick does not work. This problem is in a sense a weighted-center problem with

an "additive" weight. We propose a different approach where the constraints remain nonlinear but the search technique still works. This approach yields linear-time algorithms for a class of problems including the (multiplicative) weighted-center problem.

2. The Smallest Ball Containing a Set of Balls

For any point $p \in R^d$ and any nonnegative real number r, let $B(p, r)$ denote the ball with radius r centered at p, that is,

$$B(p, r) = \{y \in R^d : \|p - y\| \le r\}.$$

The following problem arose in an algorithm by Pollack and Sharir [PS] for the geodesic center of a simple polygon.

Problem. Given n balls $B(p^i, r_i) \subset R^d$ $(i = 1, \ldots, n)$, find a ball $B(x, \rho) \subset R^d$ with minimal radius ρ, such that $B(p^i, r_i) \subset B(x, \rho)$ $(i = 1, \ldots, n)$.

It is easy to verify the following fact:

Fact. *A ball $B(p, r)$ is contained in another ball $B(x, \rho)$ if and only if*

$$\|x - p\| + r \le \rho.$$

Thus, the underlying optimization problem is the following:

(P)
$$\begin{array}{ll} \text{Minimize} & \rho \\ \text{subject to} & \|x - p^i\| + r_i \le \rho \qquad (i = 1, \ldots, n). \end{array}$$

Notice that the feasible domain of (P) is a convex subset of R^{d+1} since the function of the distance from any point is convex. The problem (P) can be interpreted more generally. In the solution procedure below we do not need the numbers r_i to be nonnegative. Assume the points p^i are demand points so that the point p^i is currently served by a facility at a distance r_i. The problem (P) is to find a location x for a facility so as to minimize the largest loss to any demand point, where the loss is measured as the difference between the new service distance, $\|x - p^i\|$, and the previous one r_i.

To simplify the form of the problem, the inequalities

$$\|x - p^i\| + r_i \le \rho$$

can be written as

$$\|x - p^i\|^2 \le (\rho - r_i)^2,$$

where ρ is constrained to be greater than or equal to the largest r_i. Finally, the inequalities can be written in the form

$$\|x\|^2 - 2(p^i)^T x + \|p^i\|^2 - \rho^2 + 2r_i\rho - r_i^2 \le 0.$$

Let us denote the left-hand side of the latter by $f_i(x, \rho)$,

$$f_i(x, \rho) = \|x\|^2 - 2(p^i)^T x + \|p^i\|^2 - \rho^2 + 2r_i\rho - r_i^2.$$

$$\text{Minimize} \quad \rho$$

(P) $$\text{subject to} \quad f_i(x, \rho) \le 0 \qquad (i = 1, \dots, n),$$

$$\rho \ge R,$$

where $R = \max\{r_i\}$.

Although the functions $f_i(x, \rho)$ are not linear, they all have the same nonlinear terms, namely, $\|x\|^2$ and $-\rho^2$. This implies that, for any $i \ne j$, the equation

$$f_i(x, \rho) = f_j(x, \rho)$$

defines a hyperplane

$$H_{ij} = \{(x, \rho): 2(p^i - p^j)^T x - 2(r_i - r_j)\rho = \|p^i\|^2 - \|p^j\|^2 - r_i^2 + r_j^2\}.$$

This hyperplane has the following useful property. If the point (x, ρ) is known to lie in the open halfspace delimited by H_{ij} where $f_i(x, \rho) < f_j(x, \rho)$ then the constraint $f_i(x, \rho) \le 0$ becomes redundant. Similarly, if the point (x, ρ) is known to lie in the complementary open halfspace $f_i(x, \rho) > f_j(x, \rho)$ then the constraint $f_j(x, \rho) \le 0$ becomes redundant. Here we can apply the multidimensional search technique proposed in [M3] as we outline below. For more detail the reader is referred to [M3] and [D2].

The technique can be loosely stated as follows. Assume there exists a certain unknown point $x^* \in R^d$, and an "oracle" is available which can tell for any hyperplane $H \subset R^d$ whether the point x^* lies on H, or else which of the halfspaces bounded by H contains x^* in its interior. Given n hyperplanes $H_1, \dots, H_n \subset R^d$, we would like to decide the position of the unknown point x^* relative to "many" of the hyperplanes H_i $(i = 1, \dots, n)$, consulting the oracle only a "small" number of times. The interpretation of "many" and "small" depends on the particular version used. In our case the unknown point is the optimal solution (x, ρ) to our problem, which is unique by strict convexity. An oracle is an algorithm that decides the position of this point relative to any given hyperplane H. As pointed out in [M3], when the objective function is convex, we can use as an oracle an algorithm that solves the optimization problem with an additional requirement that the point lie on H; then, another optimization problem over the neighborhood of the constrained optimum reveals the "correct" side of H.

In our case the oracle would work as follows. Consider a constrained problem:

$$\text{Minimize} \quad \rho$$

(P_c) $$\text{subject to} \quad \|x - p^i\| + r_i \le \rho \qquad (i = 1, \dots, n),$$

$$a^T x + \alpha\rho = \beta,$$

where $a \in R^d$ and α and β are scalars. To solve the problem (P_c), we recurse on the dimension of the domain. Suppose (x', ρ') is the optimal solution of (P_c). There is at most one side of the hyperplane $H = \{x: a^T x + \alpha \rho = \beta\}$ where the objective function of (P) may attain a value smaller than ρ'. Such a side may be found as follows. First, let I denote the set of indices i of "active" constraints at (x', ρ'), that is, for $i \in I$,

$$\|x' - p^i\| + r_i = \rho'.$$

We are interested in a direction h of movement from x' along which all the functions

$$d_i(x) = \|x - p^i\| + r_i$$

are decreasing. Obviously, h is such a direction if and only if, for every $i \in I$,

$$(x - p^i)^T h < 0.$$

If no such h exists then the optimum of the constrained problem coincides with that of the unconstrained one; otherwise, a solution h indicates a side of the hyperplane H that contains the optimal solution of (P) as follows. Consider points x of the form

$$x = x(\varepsilon) = x' + \varepsilon h$$

and let

$$\rho = \rho(\varepsilon) = \max\{\|x(\varepsilon) - p^i\| + r_i\},$$

where $\varepsilon \geq 0$. At least the first linear piece of $\rho(\varepsilon)$ indicates a direction of improvement and hence the "correct" side of the hyperplane H. A direction h as above can be found by solving a linear-programming problem (in fixed dimension).

Notice that the recursion gives rise to slightly more general problems, where in addition to the initial constraints we also have several *linear* equality constraints through which the dimension is reduced. This does not represent a difficulty at all. Linear equality constraints can be used to eliminate variables. We then end up with a generalized form of the problem (P):

$$(\tilde{P}) \qquad \begin{aligned} &\text{Minimize} \quad \rho \\ &\text{subject to} \quad \|Ax - p^i\| + r_i \leq \rho \qquad (i = 1, \ldots, n), \end{aligned}$$

where A is a certain matrix. The latter can be solved with essentially the same methods.

As in the linear programming algorithm in fixed dimension, the search technique lets us discard in linear time a fixed fraction of the constraints and hence the entire procedure runs in linear time in any fixed dimension. Using the improvements proposed by Dyer [D2] and Clarkson [C1], the coefficient of the linear function varies with the dimension like $O(3^{d^2})$.

3. The Weighted-Center Problem

The weighted-center problem can be solved using an idea similar to the one presented in Section 2. The optimization problem is the following:

$$(W) \qquad \begin{aligned} &\text{Minimize} \quad \lambda \\ &\text{subject to} \quad w_i\|x - p^i\| \le \lambda \qquad (i = 1, \ldots, n), \end{aligned}$$

where w_i $(i = 1, \ldots, n)$ are given positive constants. This is equivalent to

$$\begin{aligned} &\text{Minimize} \quad \lambda \\ &\text{subject to} \quad \|x - p^i\| \le \lambda / w_i \qquad (i = 1, \ldots, n), \end{aligned}$$

or

$$\begin{aligned} &\text{Minimize} \quad \mu \\ &\text{subject to} \quad \|x\|^2 - 2(p^i)^T x + \|p^i\|^2 \le \mu / w_i^2 \qquad (i = 1, \ldots, n). \end{aligned}$$

Following the idea of Section 2, it becomes useful to consider the hyperplanes

$$H_{ij} = \left\{ (x, \mu): 2(p^i - p^j)^T x + \left(\frac{1}{w_i^2} - \frac{1}{w_j^2} \right) \mu = \|p^i\|^2 - \|p^j\|^2 \right\}.$$

The rest of the algorithm is very similar to the algorithm developed in Section 2.

Acknowledgment

References

[Ch] R. Chandrasekaran, The weighted Euclidean 1-center problem, *Operations Research Letters* **1** (1982), 111–112.

[Cl] K. Clarkson, Linear programming in $O(n \times 3^{d^2})$ time, *Information Processing Letters* **22** (1986), 21–24.

[Co] R. Cole, Slowing down sorting networks to obtain faster sorting algorithms, in: *Proceedings of the 25th Annual IEEE Symposium on Foundations of Computer Science* (1984), IEEE Computer Society Press, Los Angeles, 1984, pp. 255–260.

[D1] M. E. Dyer, Linear-time algorithms for two- and three-variable linear programs, *SIAM Journal on Computing* **13** (1984), 31–45.

[D2] M. E. Dyer, On a multidimensional search technique and its application to the Euclidean one-center problem, *SIAM Journal on Computing* **15** (1986), 725–738.

[M1] N. Megiddo, Linear-time algorithms for linear programming in R^3 and related problems, SIAM *Journal on Computing* **12** (1983), 759–776.

[M2] N. Megiddo, The weighted Euclidean 1-center problem, *Mathematics of Operations Research* **8** (1983), 498–504.

[M3] N. Megiddo, Linear programming in linear time when the dimension is fixed, *Journal of the Association for Computing Machinery* **31** (1984), 114–127.

[MZ] N. Megiddo and E. Zemel, A randomizing $O(n \log n)$ algorithm for the weighted Euclidean 1-center problem, *Journal of Algorithms* **7** (1986), 358–368.

[OKM] J. O'Rourke, S. R. Kosaraju, and N. Megiddo, Computing circular separability, *Discrete & Computational Geometry* **1** (1986), 105–113.

[PS] R. Pollack and M. Sharir, Computing the geodesic center of a simple polygon, Technical Report No. 231, Courant Institute of Mathematical Sciences, New York University, July 1986.