Smallest Enclosing Cylinders

Elmar Schömer* Jürgen Sellen† Marek Teichmann‡ Chee Yap§

We address the complexity of computing the smallest-radius infinite cylinder that encloses an input set of n points in 3-space. We further report on experimental work involving an exact and a numerical strategy.

A major topic of geometric optimization is to approximate point sets by simple geometric figures. This includes extensively studied planar problems such as smallest enclosing circles, the minimum width annulus, and the minimum width slab. In higher dimensions, there are few non-trivial complexity results for geometric figures beyond hyperplanes or spheres. We consider the following:

> **Smallest Cylinder Problem (P1):** Let I be a given set of n points in 3-space. Find a line ℓ which minimizes $\max\{ d(\ell, c) : c \in I \}$.

Here, $d(\ell, c)$ denotes the minimum Euclidean distance between c and a point of ℓ. Cylinders constitute an important primitive shape in computer-aided design and manufacturing. For example in the area of *dimensional tolerancing and metrology* (see [SV, Ya]), the task is, given a physical object, to verify its conformance to tolerance specifications by taking probes of its surface.

1 Summary of Theoretical Results

The heart of our approach is to consider restricted versions of the problem, fixing one or more optimization parameters, and using decision scheme for the restricted subproblem. The details and other results can be found in [SSTY]. To obtain efficient decision algorithms, it is often possible to exploit a *linearization* technique. We define an

> **Abstract Decision Problem (D):** Given a set $I \subseteq \mathbf{R}^m$ of n points, decide if there exists a point $c \in \mathbf{R}^\ell$ such that for all $p \in I$, $P(c, p) \leq 0$.

We say $P(\mathbf{x}, \mathbf{y})$ has an *order k linearization* if there exists $2k+1$ polynomials, $X_i = X_i(\mathbf{x})$ $(i = 1, \ldots, k)$ and $Y_i = Y_i(\mathbf{y})$ (for $i = 0, \ldots, k$), such that $P(\mathbf{x}, \mathbf{y}) = Y_0 + \sum_{i=1}^k X_i Y_i$.

Theorem 1
(i) If $P(\mathbf{x}, \mathbf{y})$ has an order k linearization, the decision problem (D) can be solved in $O(n^{\lfloor k/2 \rfloor})$ in the algebraic model.
(ii) In the bit model, if each input coordinate has L bits, the problem (D) can be solved $O(\mu(L) n^{\lfloor k/2 \rfloor})$.

In our application, our focus is the fixed-parameter problem to decide whether there exists an anchored cylinder of given radius r that encloses all input points. We obtain an order 9 linearization for this decision problem. Applying theorem 1, we conclude that the fixed-parameter problem for (P1) can be decided in time $O(n^4)$ in an algebraic model, and in time $O(\mu(L) n^4)$ in a bit model. We then obtain

Theorem 2 *Problem (P1) can be solved in time:*
(i) $O(n^4 \log^{O(1)} n)$ in an algebraic model; and
(ii) $O(L\mu(L) n^4)$ in a bit model.

Here, $\mu(L) = O(L \log L \log\log L)$ denotes the complexity of multiplying two L-bit integers. The algebraic and bit complexity models are described in [SSTY]. We assume that in the algebraic model, each input point c satisfies $||c|| = O(1)$, and that in the bit model, the coordinates of each c are given as homogeneous rational numbers of bit-size L.

The first result follows from an application of the parametric search paradigm (see eg. [Me]), and the use of a parallel convex hull algorithm [AGR]. For the second result, we exploit ideas from the theory of exact computation, and show in the bit model that the combinatorial solution of (P1) i.e. a set of 4 or 5 points defining the smallest enclosing cylinder, can be obtained from an ε-approximate solution for r^* if $\varepsilon = 2^{-O(L)}$. To compute this approximate solution, it suffices to run the decision algorithm for the fixed-parameter problem $O(L)$ times, with radii of bit-size $O(L)$ as input.

We also describe approximation algorithms for the smallest cylinder problem. We obtain complexity trade-offs between n and ε:

Theorem 3 *In an algebraic model of computing, an ε-approximate solution of (P1) can be found in times (resp.):*
$$O(n\varepsilon^{-2} \log \varepsilon^{-1}),\ O(n^3 \varepsilon^{-1} \log \varepsilon^{-1}),\ O(n^4 \log \varepsilon^{-1}).$$

These algorithms are based on *discretizing* various sets of input parameters. For the first bound, the direction of the cylinder axis discretized. We give bounds on the amount of change of the radius of the smallest cylinder for small changes in the orientation of its axis, from which the result follow. The second result requires yet another application of the linearization technique. We find a decision problem with an "intermediate" number of free parameters. The third result is an extension of the previous theorem.

Finally, in an attempt at understanding the combinatorial complexity of the problem, we show that for n given input points, there can be $\Omega(n)$ globally smallest and $\Omega(n^2)$ locally smallest enclosing cylinders.

2 Experimental Results

In this section we describe a simple optimization method and evaluate this method for accuracy by comparing its results against "exact" results that we obtained with MAPLE.

*schoemer@cs.uni-sb.de, Universität des Saarlandes.
†sellen@cs.uni-sb.de, Universität des Saarlandes.
‡teichman@cs.nyu.edu, Courant Institute, NYU.
§yap@cs.nyu.edu, Courant Institute, NYU.

Exact solution: We first treat our solution in MAPLE and its algebraic formulation.

A cylinder C in 3-space is specified by 5 real parameters, its axis line ℓ and its radius r. We first specify the set $C(c_1, \ldots, c_4)$ of cylinders that touch 4 given points $c_1, \ldots, c_4$. We can assume $c_1 = (0,0,0)$. Let $u \in \mathbf{R}^3$ be any direction vector of ℓ. Let E be the plane passing through the origin and orthogonal to u, and let $c_1^*, \ldots, c_4^*$ be the orthogonal projection of the input points $c_1, \ldots, c_4$ onto E. Then the cylinder C passes through $c_1, \ldots, c_4$ if and only if $c_1^*, \ldots, c_4^*$ are cocircular.

A suitable parametrization for the direction vector u is found using barycentric coordinates. Assume u is not parallel to the plane containing c_2, c_3, c_4. Let $u = xc_2 + yc_3 + zc_4$, with $z = 1 - x - y$.

Now, let $R_1(x,y,z)$ (resp. $R_2(x,y,z)$) be the squared radius of the circumcircle of c_1^*, c_2^*, c_3^* (resp. c_1^*, c_3^*, c_4^*) in E. Then $C(c_1, \ldots, c_4)$ can be interpreted as a 2-dimensional surface in 3-space, defined by $R_1(x,y,z) = R_2(x,y,z)$. This condition is equivalent to $P(x,y,z) = 0$, with

$$\begin{aligned} P(x,y,z) = {} & \Delta_{1,2,4}(xz^2 + x^2z) + \Delta_{1,3,4}(yz^2 + y^2z) \\ & + \Delta_{1,2,3}(xy^2 + x^2y) + (\Delta_{1,2,4} + \Delta_{1,3,4} + \Delta_{1,2,3} \\ & - \Delta_{2,3,4})(xyz), \text{ where } \Delta_{i,j,k} = c_i(c_j \times c_k). \end{aligned}$$

In order to compute the cylinders with fixed radius r in the set $C(c_1, \ldots, c_4)$, the additional condition $R_1(x,y,z) = r$ has to be satisfied. This leads to a polynomial equation $Q(x,y) = 0$, with total degree 6.

The set $C_f(c_1, \ldots, c_4, r)$ of all cylinders with radius r that pass through $c_1, \ldots, c_4$ is given by the set of solutions of the system $\{Q(x,y) = P(x,y) = 0\}$, and can be obtained algebraically by computing the roots of the resultants $F_x = Res(P,Q,y)$ and $F_y = Res(P,Q,x)$, each of degree 12. Hence, under certain assumptions, each cylinder is specified uniquely by algebraic numbers of degree at most 12.

Numerical solution: Here we exploit the fact that each axis direction uniquely determines a smallest enclosing cylinder. By this reduction, the optimization problem can be viewed as a search for the minimum on a 2-dimensional surface in 3-space. Each point of this surface can be obtained as the result of a convex optimization problem. Thus, we seek the minimum of a composed function $f \circ g$. Accordingly, we choose an optimization technique which only requires function *evaluations* but not computation of derivatives. We use the standard downhill simplex algorithm [PTVF] which tries to follow the direction of steepest descent. It is applied in two layers, to compute the minimum of f and (recursively) that of g.

For a given start axis, the optimization method converges to some local minimum. To locate a global minimum, one may choose a 2-dimensional grid of start values. But a better choice may be to choose the set of directions of edges in the convex hull of the input points [SSTY].

In the sequel, we shall report on some experimental results with this special set of start values. We first computed smallest enclosing cylinders for randomly generated tetrahedra. In a sequence of 100 tests, at least one of the 6 considered start values (the edge directions of the tetrahedron) led to the optimum. In two additional test sequences, we tested 50 sets of 5 random points, and 10 sets of 8 random points (in convex position.) Similar results were observed.

The most complex examples which we tried consisted of 12 points. The MAPLE implementation ran several days on these sets to find the optimum. The numerical optimization converged within seconds for each starting value. To stimulate further research, we include the data as benchmarks:

Example: The 12 input points are arranged near the 12 vertices of an icosahedron with center at the origin:

	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}	c_{11}	c_{12}
x	−12	−12	9	10	23	10	−10	−24	12	12	−10	−8
y	−19	1	−13	12	−1	−21	21	1	19	−1	−12	12
z	−6	−20	−18	−17	−2	5	−6	2	6	20	17	18

The optimal solution has been computed by MAPLE as the cylinder through the 5 points $c_1, c_3, c_7, c_9, c_{11}$, with radius ≈ 21.0309. The downhill simplex algorithm obtained this solution for the starting values (c_3, c_1), (c_3, c_2), (c_4, c_2), (c_5, c_1), (c_8, c_5) and (c_9, c_8). See [SSTY] for additional examples. To conclude this section, we observe that the proposed downhill algorithm behaves amazingly well, and did not fail for the examples we tried.

3 Final Remarks

As the field of geometric optimization matures, it treats problems of increasingly non-trivial algebraic complexity. The traditional neglect of bit complexity is no longer justified. The smallest cylinder problem is one of these problems. By combining the general linearization technique with parametric search, we developed efficient algorithms in both models. These results seem mainly of theoretical interest.

Our ε-approximation schemes have possibly greater practical applicability. But even here, our numerical experiments suggest that these may not be competitive with some heuristic numerical approaches. A possible reason for the effectiveness of our heuristics may be that the number of local minima is – either generally or in a randomized setting – much smaller than the trivial bound of $O(n^5)$. Indeed, our lower bound of $\Omega(n^2)$ leaves a wide gap for further research.

References

[AGR] N. Amato, M. Goodrich, E. Ramos, "Parallel algorithms for higher-dimensional convex hulls", *IEEE FOCS*, 1995.

[Me] N. Megiddo. "Applying parallel computation algorithms in the design of serial algorithms", *Journal of the ACM*, 30, 1983, pp. 852-865.

[PTVF] W. Press, S. Teukolsky, W. Vetterling, B. Flannery, *Numerical Recipes in C*, Cambridge U. Press, 1988.

[SSTY] E. Schömer, J. Sellen, M. Teichmann, C.K. Yap, "Smallest Enclosing Cylinders", Technical Report, Courant Inst, New York University, 1996, http://www.cs.nyu.edu.

[SV] V. Srinivasan, H. B. Voelcker. *Dimensional Tolerancing and Metrology*, The American Society of Mechanical Engineers, New York, CRTD-Vol. 27, 1993.

[Ya] C. Yap, "Exact Computational Geometry and Tolerancing Metrology", in *Snapshots of Computational and Discrete Geometry, Vol.3*, eds. David Avis and Jit Bose, McGill School of Comp.Sci, Tech.Rep. No.SOCS-94.50, 1994.