

Efficiently Approximating Polygonal Paths in Three and Higher Dimensions*

Gill Barequet[†]

Danny Z. Chen[‡]

Ovidiu Daescu[‡]

Michael T. Goodrich[†]

Jack Snoeyink[§]

Abstract

We present efficient algorithms for solving polygonal-path approximation problems in three and higher dimensions. Given an n -vertex polygonal curve P in \mathbb{R}^d , $d \geq 3$, we approximate P by another polygonal curve P' of $m \leq n$ vertices in \mathbb{R}^d such that the vertex sequence of P' is an ordered subsequence of the vertices of P . The goal is to either minimize the size m of P' for a given error tolerance ε (called the *min-# problem*), or to minimize the deviation error ε between P and P' for a given size m of P' (called the *min- ε problem*). Our techniques enable us to develop efficient near-quadratic-time algorithms in 3-D and sub-cubic-time algorithms in 4-D for solving the min-# and min- ε problems. We discuss extensions of our solutions to d -dimensional space, where $d > 4$.

*Work on this paper by the first and the fourth authors has been supported in part by the U.S. Army Research Office under Grant DAAH04-96-1-0013. Work by the second and third authors has been supported in part by the National Science Foundation under Grant CCR-9623585. Work by the fourth author has been supported also by NSF grant CCR-96-25289. Work by the fifth author has been supported by NSERC, CIES, and by the Geometry Center at Johns Hopkins University.

[†]Center for Geometric Computing, Dept. of Computer Science, Johns Hopkins University, Baltimore, MD 21218. E-mail: [barequet|goodrich]@cs.jhu.edu

[‡]Dept. of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556. E-mail: [chen|odaescu]@cse.nd.edu

[§]Dept. of Computer Science, The University of British Columbia, 201-2366 Main Mall, Vancouver, B.C., Canada V6T 1Z4. E-mail: snoeyink@cs.ubc.ca

1 Introduction

In this paper we consider the problem of approximating an arbitrary n -vertex polygonal path P in 3-dimensional space by another polygonal path P' whose vertices form an ordered subsequence of the vertices in the original path P . We further discuss this problem in the d -dimensional space, where $d \geq 4$.

This problem arises in many applications, such as robotics, image processing, computer graphics, cartography, and data compression. In these applications it is often desirable to approximate a complex graphical or geometric object, which is specified by a polygonal path, by a simpler object that captures the essence of the original object yet achieves a certain degree of data compression [3, 23, 24]. For example, we have been asked to compress 3D path data for optic nerves acquired in medical studies (Kelly Booth, personal communication, 1995) and of river networks in geographic information systems analysis of spawning habitat [20]. Path compression is also important for devices with low bandwidth, from pen plotters to the WWW.

In general, the goal of the path-approximation problem is to replace a (very) high-resolution path P with an approximate path P' that achieves significant time and space improvements while suffering only a small approximation error. Two factors come to play in such an approximation: the number of vertices in the approximation and the error between the approximation and the original path. These two factors give rise to two different versions of the path-approximation problem—the min- ε version and the min-# version. In the min- ε version one is given a parameter $m \leq n$ and asked to find an approximating path P' with m vertices whose *error*, ε , from the original path is minimized.




Figure 1: The error tolerance zone of a single segment $\overline{p_i p_j}$ in 2-D.

This version of the problem is motivated by a desire to obtain the best approximation possible that achieves a certain degree of data compression. In the min-# version of the path-approximation problem one is instead given a parameter $\varepsilon > 0$ and asked to find the smallest-vertex path P' that is within *error* ε of the original path. This version of the problem is motivated by a desire to obtain the smallest-complexity path that maintains a certain level of accuracy. We study both versions of the d -dimensional path-approximation problem, for $d \geq 3$, in this paper.

One issue that immediately arises in the path approximation problem is that one must formalize the “error” between the original path P and the approximation P' . We believe that one of the most natural definitions is the *tolerance zone* error measure [3, 17, 18, 22]. Define the ε -tolerance zone (for $\varepsilon > 0$) of a line segment \overline{pq} to be the region of space (or “zone”) that is the union of all radius- ε balls centered at points along the segment \overline{pq} (see Figure 1). This definition is, of course, parameterized by the metric used to define radius- ε balls. We consider the usual Euclidean metric, L_2 , as well as the L_∞ and L_1 distance metrics. We may then formalize the path-approximation problem as follows:

Polygonal Path Approximation:

Given a polygonal path $P = (p_1, p_2, \dots, p_n)$, find a path $P' = (p_{i_1}, p_{i_2}, \dots, p_{i_m})$, such that, for each $j \in \{1, 2, \dots, m-1\}$:

- $i_j \in \{1, 2, \dots, n\}$, for $j = 1, 2, \dots, m$;
- $i_j < i_{j+1}$; and
- the subpath $(p_{i_j}, p_{i_j+1}, \dots, p_{i_{j+1}})$ is completely contained in the ε -tolerance zone for the segment $\overline{p_{i_j} p_{i_{j+1}}}$.

If the subpath $P_{i_j, i_{j+1}} = (p_{i_j}, p_{i_j+1}, \dots, p_{i_{j+1}})$ is completely contained in the ε -tolerance zone for the segment $\overline{p_{i_j} p_{i_{j+1}}}$, we call $\overline{p_{i_j} p_{i_{j+1}}}$ an *approximating line segment* for $P_{i_j, i_{j+1}}$.

In the min- ε version, we are given $m < n$ and wish to minimize ε ; in the min-# version, we are given $\varepsilon > 0$ and wish to minimize m .

We restrict our attention to the version of this problem that requires that $p_{i_1} = p_1$ and $p_{i_m} = p_n$. In the general version we only require that $i_1 \leq i_m$, and that the subpaths P_{1, i_1-1} and $P_{i_m+1, n}$ are completely contained in the ε -balls centered at P_{i_1} and P_{i_m} , respectively. We explain later in the introduction how to reduce the general problem to the restricted version of it.

1.1 Related Previous Work

Imai and Iri [17, 18] and Melkman and O’Rourke [22] study the 2-D min-# and min- ε polygonal-path approximation problems. Using the tolerance zone measure of error, as defined above, they achieve algorithms whose running times are $O(n^2 \log n)$ and $O(n^2 \log^2 n)$, respectively, and using $O(n^2)$ space, for the two problems. Chan and Chin [4] reduce the time complexity of both results by a logarithmic factor. Chen and Daescu [7] further show that the algorithms of [4] can use only $O(n)$ space without increasing their running times.

Varadarajan [27] studies the min-# and min- ε problems for 2-D polygonal paths that are monotone, i.e., any line parallel to the y axis intersects a path in a point. Given a monotone approximation of a monotone path, Varadarajan measures the error as the maximum distance between the intersection points with some line parallel to the y axis. He gives $O(n^{4/3+\delta})$ time and space algorithms for both problems, where $\delta > 0$ is an arbitrarily small constant. However, those algorithms cannot be extended to the 3-D version of the problem or to the tolerance zone error criterion in 2-D. The recursive simplification heuristic of Douglas and Peucker [8], which is popular in GIS, does extend to curves in 3 dimensions in $O(n^2)$ worst-case time, but solves neither the min-# nor the min- ε problems.

There has been other less-related, but still significant, work done on other metrics for comparing 2-dimensional polygonal paths [18] that do not necessarily use the same set of vertices. Arkin et al. [2], Alt and Godau [1], and Rote [25] describe several metrics for comparing polygonal curves. Guibas et al. [15] study different error measures for the min-# problem, with an eye toward those that could

be implemented in $O(n \log n)$ time by greedy algorithms. Fleischer et al. [11] approximate polygonal shapes by inner and outer polygons and show their applications to the problems of polygon containment and of planar motion planning.

The only work on the general problem, that we are aware of, achieves super-cubic algorithms for the 3-D path-approximation problem. Ihm and Naylor [16] give an $O(n^3 \log m)$ algorithm for the min- ε problem. Eu and Toussaint [10] restrict the polygonal chains to be monotonic with respect to one of the three coordinate axes, and define the approximation in terms of an *infinite beam* measure of error. In this measure of error one requires that each subpath $(p_{i_j}, p_{i_j+1}, \dots, p_{i_{j+1}})$ be completely contained in the ε -tolerance zone of the *infinite line* containing the segment $\overline{p_{i_j} p_{i_{j+1}}}$, called the infinite beam. When infinite beams are defined in terms of the L_1 or L_∞ metrics, Eu and Toussaint show how to achieve running times of $O(n^2)$ and $O(n^3)$ for the min-# and min- ε problems, respectively, for monotone 3-D polygonal paths. Under the Euclidean (L_2) metric, they consider the 3-dimensional version of the problem in [27] and give algorithms that run in $O(n^3)$ and $O(n^3 \log n)$ time, respectively, for the min-# and min- ε problems on monotone 3-D polygonal paths.

1.2 Our Results

In this paper we give the first efficient algorithms for the d -dimensional path approximation problem, where $d \geq 3$, for general (non-monotonic) paths, using the tolerance-zone measure of error for segments. We believe that this notion of error better captures the intuitive concept of the shape of a d -dimensional polygonal path. We illustrate in Figure 2 the difference between approximating a 2-dimensional path with tolerance zones and with infinite beams.

We summarize the running times of our methods in Table 1. Note that all of our bounds are sub-cubic (if not nearly quadratic). Fairly straightforward bounds of $O(n^3)$ for min-# and $O(n^3 \log n)$ for min- ε follow by adapting previous 2-dimensional results [4, 17, 18, 22]. Even so, our approach still follows the general framework of these earlier algorithms. To solve the min-# problem, we first build a graph G on the vertices of the input path




Figure 2: An illustration of the difference between the tolerance zone measure of error and the infinite beam measure of error, for a polygonal path in 2-D, with $\varepsilon = 3\text{mm}$. The path is shown in (a), its approximation using the tolerance zone in (b), and its approximation using infinite beams in (c).

P , where each edge represents a valid “shortcut,” and we then perform a shortest-path computation in G . To solve the min- ε problem we perform a “binary-search-like” computation using the min-# algorithm as a “probe.”

Our algorithms differ from previous approaches in some important ways, however. One of the most fundamental differences is that we develop a d -dimensional generalization of the 2-dimensional approach of Chan and Chin [4] for constructing the “shortcut” graph G . Rather than constructing this graph directly, we instead construct it as the intersection of two other graphs, \overrightarrow{G} and \overleftarrow{G} , which respectively define tolerance zones in terms of “semi-infinite” beams (pointing in opposite directions). The difficult part of this computation is that, fixing an index $1 \leq i < n$, we must determine, for each index j , with $i < j \leq n$, if the ray $\overrightarrow{p_i p_j}$ originating at p_i and passing through p_j intersects each radius- ε ball centered at a vertex p_k for $i < k < j$. We reduce this problem to the following data structuring problem:

Off-Line Ball Inclusion Testing:

Maintain a collection of different-radius balls subject to a given sequence of the following operations:

- **insert(B)**: add the ball B to the collection.
- **inside(p)**: determine if the point p is inside the common intersection of the balls.

Metric	3 Dimensions		$d \geq 4$ Dimensions	
	Min-#	Min- ε	Min-#	Min- ε
L_2	$O(n^2 \log n)$	$O(n^2 \log^3 n)$	$O(n^{3-2/(\lfloor \frac{d}{2} \rfloor + 1)} \text{polylog } n)$	$O(n^{7/3} \text{polylog } n)^\clubsuit$
L_1 & L_∞	$O(n^2)$	$O(n^2 \log n)$	$O(n^2)$	$O(n^2 \log n)$

\clubsuit For $d = 4$ only.

Table 1: Summary of our min-# and min- ε results.

For $d = 3$, we refer to the problem above as the off-line disk-inclusion testing problem (the balls are disks on a 3-D sphere) and provide a data structure for solving this problem in amortized $O(\log n)$ time for inserts and queries, even though each disk insertion can cause $\Theta(n)$ structural changes to the common intersection. We use this data structure to derive an $O(n^2 \log n)$ time algorithm for the min-# path approximation problem. To turn this into an efficient algorithm for the min- ε problem we give an application of the parametric search paradigm that is based upon a non-trivial parallel version of our off-line inclusion testing data structure.

When the ε -tolerance zones are defined using the L_1 or L_∞ metrics, we can do better than this. We present an $O(n^2)$ -time algorithm for the min-# problem and an $O(n^2 \log n)$ -time algorithm for the min- ε problem. For the latter algorithm we avoid using parametric search.

For $d > 3$, we show how to achieve a sub-cubic time for the min-# problem under the L_2 metric. We also give $O(n^2)$ and $O(n^2 \log n)$ -time algorithms (for a fixed $d > 3$) for the min-# and min- ε problems, respectively, under the L_1 and L_∞ metrics. The constants hidden in the big-Oh notation depend on d and on the used metric: they are $O(d^2)$ for L_∞ and $O(d2^d)$ for L_1 .

The general version of the problem can be reduced to the version solved in this paper (where $i_1 = 1$ and $i_m = n$) by adding to G two vertices p_0 and p_{n+1} , and edges p_0p_i (resp., p_ip_{n+1}), for $1 \leq i \leq n$, if all the vertices p_j , for $1 \leq j < i$ (resp., $i < j \leq n$) are contained in an ε -ball centered at p_i . This additional processing takes $O(n^2)$ time, and we then solve the restricted problem for p_0 and p_{n+1} .

We outline the main ideas behind our results in the remaining sections of this extended abstract.

2 Solving the Min-# Problem in 3-D

In this section we present the main ideas behind our $O(n^2 \log n)$ -time algorithm for solving the 3-D min-# problem. The important part of this computation involves a reduction to the off-line inclusion testing problem, which we solve in Section 3.

Suppose then that we are given a 3-D polygonal curve $P = (p_1, p_2, \dots, p_n)$, and an error value $\varepsilon > 0$, and we wish to find the smallest $m \leq n$ such that there is a subpath $P' = (p_{i_1}, p_{i_2}, \dots, p_{i_m})$ that satisfies the ε -tolerance zone criterion (with $i_1 = 1$ and $i_m = n$). As outlined in the introduction, we solve this problem by the following general approach:

1. Construct a graph $G = (V, E)$, where V is the vertex set of P and E consists of directed edges (p_i, p_j) , $1 \leq i < j \leq n$, such that $\overrightarrow{p_i p_j}$ is an approximating line segment of the subchain $(p_i, p_{i+1}, \dots, p_j)$ of P ;
2. Find a shortest path from p_1 to p_n in G (i.e., a path with the fewest number of edges). This path gives us the sought approximating curve P' .

Since the second step is easily solved by a breadth-first search computation, let us concentrate on an efficient technique for constructing G from P .

As mentioned above, we construct G as the intersection of two graphs \overrightarrow{G} and \overleftarrow{G} . Say that edge (p_i, p_j) is in the graph \overrightarrow{G} if the ray $\overrightarrow{p_i p_j}$ originating at p_i and passing through p_j intersects each radius- ε ball $B(p_k, \varepsilon)$, centered at a vertex p_k , for $i < k < j$. Likewise, say that edge (p_i, p_j) is in the graph \overleftarrow{G} if the ray $\overleftarrow{p_i p_j}$ originating at p_j and passing through p_i intersects each $B(p_k, \varepsilon)$, for $i < k < j$. It is an easy observation that an edge (p_i, p_j) is in G if and only if it is in both \overrightarrow{G} and \overleftarrow{G} . Since the constructions of these two graphs are




Figure 3: The projection of the radius- ε ball $B(p_k, \varepsilon)$ centered at p_k onto the sphere S_i .

symmetric computations, let us concentrate further upon the construction of the graph \vec{G} .

In this section we show how to reduce the problem of determining the edges in \vec{G} to the off-line disk-inclusion problem. Let us fix an index i , with $1 \leq i < n$, and focus on computing all the edges of the form (p_i, p_j) in \vec{G} , for $i < j \leq n$. Say that the ray $\vec{p_i p_j}$ properly approximates the polygonal chain (p_i, \dots, p_j) if it intersects every ball $B(p_k, \varepsilon)$ for $i < k < j$. Now imagine a “large” sphere S_i centered at p_i , and project from p_i all the balls $B(p_k, \varepsilon)$ ($i < k < j$) onto this sphere. Each ball $B(p_k, \varepsilon)$ is then projected to a disk D_k on the sphere S_i (see Figure 3).

Lemma 1 *The ray $\vec{p_i p_j}$ is a proper approximation if and only if:*

1. *The intersection of the set of disks $\mathcal{D}(i, j) = \{D_{i+1}, \dots, D_{j-1}\}$ on S_i is nonempty.*
2. *The projection of p_j from p_i onto S_i is found within the intersection of the disks in $\mathcal{D}(i, j)$.*

Proof: Omitted in this extended abstract. \square

This lemma gives us the ability to reduce the construction of all edges of the form (p_i, p_j) in \vec{G} , for our fixed p_i , to the off-line disk-inclusion problem. Specifically, we define the operations:

- **insert(D_j):** Inserts the disk D_j to a data structure, where D_j is the projection of $B(p_j, \varepsilon)$

onto the sphere S_i using p_i as the center of projection.

- **inside(p'_j):** Returns “true” or “false,” depending on whether the point p'_j , the projection of p_j onto S_i using p_i as the center of projection, is contained by all the disks currently in the data structure.

We then construct a sequence Σ_i of **insert** and **inside** operations, as follows:

- 1: Initialize Σ_i to insert a disk containing the entire sphere S_i .
- 2: **for** $j = i + 1$ to n **do**
- 3: Append the operation **inside(p'_j)** to Σ_i .
- 4: Append the operation **insert(D_j)** to Σ_i .
- 5: **end for**

Lemma 2 *The edge (p_i, p_j) is in \vec{G} if and only if the response to the operation **inside(p'_j)** in Σ_i is “true.”*

Proof: Omitted in this extended abstract. \square

Thus, we have reduced the problem of constructing \vec{G} to the problem of solving $n - 1$ independent instances of the off-line disk-inclusion problem. We outline in the next section how we can efficiently solve this problem.

3 The Off-Line Disk Inclusion Problem

Suppose we are given a sequence Σ of n **insert(D_i)** and **inside(p_i)** operations defining an instance of the off-line disk-inclusion problem. In this section we show how to determine the answers to all the **inside(p_i)** queries in $O(n \log n)$ total time. Before we provide our solution we remark that Sharir [26] gives a data structure for maintaining the on-line intersection of a collection of equal-radius disks. This data structure cannot be applied to our disk-inclusion problem, however, as the disks in our problem do not in general have equal radii.

Let us sketch our solution to the off-line disk-inclusion problem. We begin by building a complete binary tree T such that each leaf of T is associated with a different disk given in the sequence Σ . We order the leaves of T from left to right in the order they are given in the sequence Σ . For each internal node v in T , let $I(v)$ denote a representation of the common intersection of the disks

stored at leaf-descendants of v . In the full version of this paper we show how $I(v)$ can be constructed, for all vertices v in T in $O(n \log n)$ total time. The important observation to make at this point is that each leaf of T is associated with a different index in Σ . Thus, if an **inside** query appears in position i in Σ , we can form a search path π_i from the root of T to the leaf corresponding to the first insertion after position i . Say that a node w of T is on the *left fringe* of π_i if w is not on π_i but is a left child of a node on π_i . Noting that the **inside** query is a decomposable search problem, we can answer an **inside**(p) query for position i by determining if p is inside $I(w)$ for each w that is a left fringe node of π_i (answering “true” if and only if it is inside them all). In the full version of this paper we use this observation, together with the fractional cascading technique of Chazelle and Guibas [5, 6], to answer each **inside** query in Σ in $O(\log n)$ time. This gives us the following:

Theorem 3 *Given a sequence Σ of n disk insert updates and point **inside** queries, one can answer all the **inside** queries in Σ in $O(n \log n)$ time.*

Combining this with the discussion from the previous section immediately gives us the following:

Theorem 4 *Given an n -vertex polygonal path P in 3-dimensional space, and a parameter $\varepsilon \geq 0$, one can find the fewest-vertex approximating path P' to P in $O(n^2 \log n)$ time under the ε -tolerance zone-error criterion.*

4 The Min- ε Problem in 3-D

Let us now turn to the min- ε problem. Suppose then that we are given a 3-D polygonal curve $P = (p_1, p_2, \dots, p_n)$, and an integer parameter $m < n$, and we wish to find the smallest $\varepsilon \geq 0$ such that there is a subpath $P' = (p_{i_1}, p_{i_2}, \dots, p_{i_m})$ that satisfies the ε -tolerance zone criterion. For solving the min- ε problem, we parallelize the min-# algorithm and invoke it as a subroutine in a parametric search algorithm, where the sought solution ε^* is the minimum ε for which there exists an approximating curve of no more than m vertices.

In fact, we don’t really have to parallelize the entire min-# algorithm described above. We need

only parallelize the portions of the algorithm that depend upon the parameter ε . In the case of our min-# algorithm, the only place where ε plays a role is in the sizes of the disks that are used in the off-line disk-inclusion algorithm. Thus, the only part of the min-# algorithm that we need to parallelize is the off-line disk-inclusion algorithm itself.

The following problem is a key to our algorithm: In the plane, given n different-radius discs d_1, d_2, \dots, d_n and n points a_1, a_2, \dots, a_n , determine for every $j = 1, 2, \dots, n$, whether $a_j \in \Pi_j = \bigcap_{i=1}^j d_i$. We now give an $O(\log n)$ time, $O(n)$ processor algorithm for solving this problem in the parallel comparison model.

Our algorithm is based on Goodrich’s paradigm [12] for computing the upper envelope of n surfaces in the k -intersecting class. The algorithm is sketched below. We first build an n -leaf complete binary tree T . The i th leaf of T stores d_i and a_i . Every internal node v of T is associated with (1) the intersection $I(v)$ of the discs stored at the leaves of the subtree T_v rooted at v (of course, $I(v)$ needs to be computed); and (2) the point set $A(v)$ consisting of all the points a_i stored at the leaves of T_v . Observe that, for every j , the question of whether $a_j \in \Pi_j = \bigcap_{i=1}^j d_i$ can be answered by checking whether $a_j \in I(v)$ for $O(\log n)$ nodes v of the tree T .

One could use a naive parallel algorithm for computing the $I(v)$ ’s and checking whether $a_j \in I(v)$: Go up the tree T level by level, starting from the leaves; for every node v at each level, first compute $I(v)$ by merging the sorted sequences of the boundary vertices of $I(x)$ and $I(y)$, where x and y are the left and right children of v , and then check whether the points of $A(y)$ belong to $I(x)$. This checking can also be done by merging a sorted sequence of $A(y)$ with the boundary vertices of $I(x)$. At the root of T , we have the answers for all questions of whether $a_j \in \Pi_j$. This naive algorithm can be easily implemented in $O(\log n \log \log n)$ time and $O(n / \log \log n)$ processors in the CREW PRAM model.

To obtain a faster algorithm in the parallel comparison model, we pipe-line the merge operations in computing $I(v)$ from $I(x)$ and $I(y)$ and in checking whether $a_j \in I(v)$, as in Goodrich’s paradigm. A logarithmic number of stages are used for this

pipe-lined procedure. Every $I(v)$ is maintained as a linked list. At the end of each stage t , a list $I_t(v)$ is stored at v . We say v is *full* when $I_t(v) = I(v)$. The merge at the children of v is pipe-lined to v by maintaining a list which is an *approximately-uniform* subsequence of $I_t(v)$, and $I_{t+1}(v)$ is defined as $I_t(x) \cup I_t(y)$. The information for the points in $A(v)$ can be maintained and pipe-lined in a similar fashion as for $I(v)$. These structures at every node v can be maintained in $O(1)$ time per stage. By using the techniques of [12] we can show that the problem of checking whether $a_j \in \Pi_j$ for every j is solvable in $O(\log n)$ time using n processors in the parallel comparison model. More details our found in the full version of the paper.

We then use this parallel algorithm to drive a parametric-search algorithm [21]. We simulate each step of this parallel algorithm sequentially, noting that some of the computations in this step depend upon whether the optimal value ε^* falls in specific intervals $[\varepsilon_1, \varepsilon_2]$. We find the median such ε_i in linear time and use the min-# algorithm as a subroutine, running in $O(n^2 \log n)$ time, to determine if ε^* is above or below this median. The answer to this question resolves half of the comparisons for this parallel step. We can therefore repeat this process recursively and in $O(\log n)$ calls to the min-# algorithm we will have resolved all the comparisons for this parallel step. We then repeat this process for the next parallel step. Since our parallel algorithm takes time $O(\log n)$, this implies that the total running time for our min- ε algorithm is $O(n^2 \log^3 n)$. Thus, we have the following:

Theorem 5 *Given an n -vertex polygonal path P in 3-dimensional space, and an integer parameter $m < n$, one can find the best m -vertex approximating path P' to P in $O(n^2 \log^3 n)$ time under the ε -tolerance zone-error criterion.*

The above discussion assumed that the underlying metric for defining radius- ε balls was the L_2 metric. In a later section we explore improvements and simplifications that can be achieved if we use the L_1 or L_∞ metrics for this purpose.

5 Min-# and Min- ε Algorithm for the L_2 Metric in d -D ($d \geq 4$)

As in three dimensions, the key to solving the the min-# problem in d dimensions is to reduce it to the off-line ball-inclusion problem in a lower dimensional space and apply decomposability to the ball-inclusion problem. We sketch below the main ideas of this technique.

In order to approximate a d -D curve, we solve n off-line ball inclusion problems, each performing $O(n)$ inclusion tests of points in $(d-1)$ -balls. By a standard lifting map, testing if a point is in a set of $(d-1)$ -balls is equivalent to asking whether a point in d -D is below the lower envelope of a collection of d -D hyperplanes: The query point is lifted to a paraboloid, and the balls are mapped to hyperplanes; the intersections of the d -D hyperplanes with the paraboloid, projected into $d-1$ dimensions, are the original balls.

Thus, to solve one instance of off-line ball inclusion, we build a static data structure for answering ray-shooting queries (see [9, 19]), using a parameter s to control a tradeoff between query time and space/preprocessing time. In particular, we build a ray-shooting data structure in $O(s \text{ polylog } n)$ space and preprocessing time, that can answer each ray-shooting query in $O(n \log n / s^{1/\lfloor \frac{d}{2} \rfloor})$ time. In the full paper we give a more complete discussion of space/time tradeoffs for the off-line ball inclusion problem.

Now, to solve a min-# problem, we have to build n data structures and ask $O(n^2)$ queries. If we balance the time required for the preprocessing and for performing the queries, we find out that $s = O(n^{2-2/(\lfloor \frac{d}{2} \rfloor + 1)} \text{ polylog } n)$ and consequently obtain an $O(n^{3-2/(\lfloor \frac{d}{2} \rfloor + 1)} \text{ polylog } n)$ -time algorithm. For $d = 4$ or 5 , for example, this gives us an algorithm for the min-# problem (under the L_2 metric) that runs in $O(n^{7/3} \text{ polylog } n)$ time. Only for $d > 19$ does the running time of the algorithm exceed $O(n^{2.8} \text{ polylog } n)$.

Since each ray-shooting data structure has sub-quadratic space, and we only need one structure at a time, the amount of space required by this algorithm is still dominated by the maximum size of the graph G , which is $\Theta(n^2)$. Thus, for d -dimensions we obtain:

Theorem 6 Given an n -vertex polygonal path P in d -dimensional space, and an integer $m < n$, one can find the best m -vertex approximating path P' to P in $O(n^{3-2/(\lfloor \frac{d}{2} \rfloor + 1)} \text{polylog } n)$ time by using $O(n^2)$ space.

For the min- ε problem in 4-D, we need only observe (as for the 3-D version of the problem) that we have to parallelize only the steps in the sequential min-# algorithm that depend on ε . For this purpose we use the ray-shooting data structure of [13, 14]. With some preprocessing, the construction of this data structure does not depend on the value of ε . In this data structure the points are sorted according to their respective coordinates in parallel. As in the 3-D case, we perform the ray-shooting queries in parallel, which costs us an $O(\log^2 n)$ factor in the running time. Over all, we obtain an $O(n^{7/3} \text{polylog } n)$ -time algorithm for the min- ε 4-D problem under the L_2 metric. The details are found in the full version of the paper. Note that this method does not extend to dimensions higher than 4.

6 Min-# and Min- ε Algorithms for L_1 and L_∞ Metrics ($d \geq 3$)

In this section we present our algorithms for the min-# and min- ε problems under the L_∞ and L_1 metrics. Decomposability of on-line ball inclusion, and the fact that the unit balls in these metrics are polytopes of constant complexity for constant d allow us to give efficient general solutions in d dimensions. We describe our algorithm for the ε -tolerance zone measure of error.

Recall that to compute the graph \vec{G} we had to determine, for a fixed index $1 \leq i < n$ and each $i < j \leq n$, whether the ray $\vec{p_i p_j}$ intersects each radius- ε ball centered at a vertex p_k for $i < k < j$. As Figure 3 illustrates, the rays from p_i that intersect the radius- ε ball around p_k form a *cone*. If the ball at p_k contains p_i then the cone is the entire space.

Under L_∞ , L_1 , or any metric whose ball is a polytope whose complexity depends only on d , we can reduce the problem of constructing the graph \vec{G} to a constant number of instances of the 1-D version of Ball Inclusion, which is the 1-D version of Off-Line Interval Inclusion Testing.




Figure 4: (a) Three different views of a cube from a point. (b) Three corresponding cross-sections for a cone.

Lemma 7 In d dimensions, the question of whether a ray $\vec{p_i p_j}$ intersects each ball $B(p_k, \varepsilon)$, for $i < k < j$, reduces to $O(d^2)$ instances of Off-line Interval Inclusion Testing for the L_∞ metric or $O(d^{2d})$ instances for the L_1 metric.

Proof: Construct the cones of the rays from p_i through radius- ε balls centered at each p_k . The ray $\vec{p_i p_j}$ intersects all balls if and only if it is in the intersection of these cones. The key is to observe that halfspaces that define the boundary of these cones can be separated into a constant number of families, where each family projects to a halfplane in 2-D.

Since the metric is defined by a convex polytope, each cone is convex and polyhedral—it is the intersection of hyperplanes through p_i and the $(d - 2)$ -dimensional faces, called *ridges*, that are on the silhouette of the ball at p_k . For example, ridges in 3-D are one-dimensional edges of the 3-cube. Different views of a 3-cube give cones of different cross-sections, as shown in Figure 4, but all boundary facets of the corresponding cone are planes through p_i and a ridge of the 3-cube.

To determine whether a point p_j is in the intersection of cones defined by balls at p_k for $i < k < j$, we form an instance of Off-Line Interval Inclusion Testing for each ridge: A translated version of the ridge r generates a halfspace whenever it appears on the silhouette when viewed from p_i ; since all halfspaces generated by translates of r are bounded by hyperplanes parallel to r that contain p_i , we

can project each translate onto the circle at infinity that is contained in the 2-flat through p_i and orthogonal to ridge r . (Knowing the projection, the ridge, and p_i we can reconstruct the hyperplane that forms one of the cone's boundary facets.)

By decomposability, point p_j is in the intersection of cones if and only if it is in the intersection of their halfspaces, which happens if and only if the projection of p_j is included in each instance of Interval Inclusion Testing. Thus, \vec{G} can be constructed by solving an instance of Interval Inclusion Testing for each ridge and intersecting the solutions.

We complete the proof by counting ridges in L_∞ and L_1 polytopes. In the d -cube, every pair of non-parallel planes generates a ridge, giving $2d(d-1)$ ridges. In the d -crosspolytope (a d -D L_1 ball), a ridge is obtained by choosing positive or negative vertices for $d-1$ of the d coordinate directions, giving $d2^{d-1}$ ridges. Not every ridge is on the silhouette—for the cube, at most $2d$ will be—ridges not on the silhouette and ridges whose cubes contain p_i do not contribute constraints. \square

It is now straightforward to solve the min-# problem in $O(n^2)$ time for any constant dimension d with a constant-sized polytope for the metric. Note that the reduction to a number of 2-D problems allows us to use only $O(n)$ space, as in [7].

Theorem 8 *Given an n -vertex polygonal path P in d -D, and a parameter $\varepsilon \geq 0$, one can find the fewest-vertex approximating path P' to P in $O(n^2)$ time and $O(n)$ space under the ε -tolerance zone-error criterion using the L_∞ or L_1 error metric.*

We solve the d -D min- ε problem in $O(n^2 \log n)$ time without using parametric search. Instead, we can explicitly construct the set, $ERR(P)$, of all candidate optimal ε values and then perform a simple binary search in this set using the min-# algorithm as the “probe.” We use the fact that the optimal ε value must be determined by some edge (p_i, p_j) in G such that the farthest point p_k , with $i < k < j$, from the segment $\overline{p_i p_j}$ is at distance exactly ε from $\overline{p_i p_j}$. In the full version of the paper we show how, fixing the starting point p_i , we can determine the farthest such point p_k for each edge $\overline{p_i p_j}$ in $O(n \log n)$ time. Thus, we can compute the set $ERR(P)$ in $O(n^2 \log n)$ time, and then

perform a binary search in this quadratic-size set in $O(n^2 \log n)$ time using our quadratic-time min-# algorithm as the “probe.”

We can actually achieve an $O(n)$ space bound, but to do so we cannot afford to maintain the error set $ERR(P)$ explicitly (since $|ERR(P)| = O(n^2)$). Instead, we make use of the sampling technique of [7], which allows us to store only $O(n)$ errors of $ERR(P)$, while still being able to perform binary search on the $O(n^2)$ errors. Therefore, we can achieve the following:

Theorem 9 *Given an n -vertex polygonal path P in d -D, and an integer parameter $m \leq n$, one can find the best m -vertex approximating path P' to P in $O(n^2 \log n)$ time under the ε -tolerance zone-error criterion using the L_∞ or L_1 metric.*

References

- [1] H. ALT AND M. GODAU, Measuring the resemblance of polygonal curves, *Proc. 8th Ann. ACM Symp. on Computational Geometry*, Berlin, Germany, 102–109, 1992.
- [2] E.M. ARKIN, L.P. CHEW, D.P. HUTTENLOCHER, K. KEDEM, AND J.S.B. MITCHELL, An efficiently computable metric for comparing polygonal shapes, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(3):209–216, 1991.
- [3] B. BUTTENFIELD, Treatment of the cartographic line, *Cartographica*, 22:1–26, 1985.
- [4] W.S. CHAN AND F. CHIN, Approximation of polygonal curves with minimum number of line segments or minimum error, *Int. J. of Computational Geometry and Applications*, 6(1):59–77, 1996.
- [5] B. CHAZELLE AND L.J. GUIBAS, Fractional cascading: I. A data structuring technique, *Algorithmica*, 1:133–162, 1986.
- [6] B. CHAZELLE AND L.J. GUIBAS, Fractional cascading: II. Applications, *Algorithmica*, 1:163–191, 1986.
- [7] D.Z. CHEN AND O. DAESCU, Space-efficient algorithms for approximating polygonal curves in two-dimensional space. Manuscript, 1997.
- [8] D.H. DOUGLAS AND T.K. PEUCKER, Algorithms for the reduction of the number of points required to represent a digitized line or its caricature, *Canadian Cartographer*, 10(2):112–122, 1973.
- [9] J. ERICKSON, Space-time tradeoffs for emptiness queries, *Proc. 13th Ann. ACM Symp. on Computational Geometry*, Nice, France, 304–313, 1997.
- [10] D. EU AND G.T. TOUSSAINT, On approximation polygonal curves in two and three dimensions, *CVGIP: Graphical Models and Image Processing*, 56(3):231–246, 1994.
- [11] R. FLEISCHER, K. MEHLHORN, G. ROTE, E. WELZL, AND C.-K. YAP, Simultaneous inner and outer approximation of shapes, *Algorithmica*, 8:365–389, 1992.

- [12] M.T. GOODRICH, Using approximation algorithms to design parallel algorithms that may ignore processor allocation. *Proc. 32nd Ann. IEEE Symp. on Foundations of Computer Science*, San Juan, Puerto Rico, 711–722, 1991.
- [13] M.T. GOODRICH AND R. TAMASSIA, Dynamic ray shooting and shortest paths via balanced geodesic triangulations, *Proc. 9th Ann. ACM Symp. on Computational Geometry*, San Francisco, CA, 318–327, 1993.
- [14] M.T. GOODRICH AND R. TAMASSIA, Dynamic trees and dynamic point location, *SIAM J. Computing*, to appear.
- [15] L.J. GUIBAS, J.E. HERSHBERGER, J.S.B. MITCHELL, AND J.S. SNOEYINK, Approximating polygons and subdivisions with minimum link paths, *Int. J. of Computational Geometry and Applications*, 3(4):383–415, 1993.
- [16] I. IHM AND B. NAYLOR, Piecewise linear approximations of digitized space curves with applications, in N. M. Patrikalakis (ed.), *Scientific Visualization of Physical Phenomina*, Springer-Verlag, Tokyo, 545–569, 1991.
- [17] H. IMAI AND M. IRI, Computational-geometric methods for polygonal approximations of a curve, *Computer Vision, Graphics, and Image Processing*, 36:31–41, 1986.
- [18] H. IMAI AND M. IRI, Polygonal approximations of a curve-formulations and algorithms, *Computational Morphology*, 71–86, North-Holland, Amsterdam, 1988.
- [19] J. MATOUŠEK AND O. SCHWARZKOPF, On ray shooting in convex polytopes, *Discrete & Computational Geometry*, 10:215–232, 1993.
- [20] M. MCALLISTER AND J. SNOEYINK, Medial axis generalisation of hydrology networks, *AutoCarto 13: ACSM/ASPRS Ann. Convention Technical Papers*, Seattle, WA, 164–173, 1997.
- [21] N. MEGIDDO, Applying parallel computation algorithms in the design of serial algorithms, *J. of the ACM*, 30:852–865, 1983.
- [22] A. MELKMAN AND J. O’ROURKE, On polygonal chain approximation. *Computational Morphology*, 87–95, North-Holland, Amsterdam, 1988.
- [23] B.K. NATARAJAN, On comparing and compressing piecewise linear curves, Technical Report, Hewlett Packard, 1991.
- [24] B.K. NATARAJAN AND J. RUPPERT, On sparse approximations of curves and functions, *Proc. 4th Canadian Conf. on Computational Geometry*, 250–256, 1992.
- [25] G. ROTE, A new metric between polygons, and how to compute it, *Proc. 19th Int. Colloq. Automata Lang. Program., Lecture Notes in Computer Science*, 623:404–415, Springer-Verlag, 1992.
- [26] M. SHARIR, A near-linear algorithm for the planar 2-center problem, *Proc. 12th Ann. ACM Symp. on Computational Geometry*, Philadelphia, PA, 106–112, 1996.
- [27] K.R. VARADARAJAN, Approximating monotone polygonal curves using the uniform metric, *Proc. 12th Ann. ACM Symp. on Computational Geometry*, Philadelphia, PA, 106–112, 1996.