Short Cuts in Higher Dimensional Space

Gautam Das * Giri Narasimhan †

1 Introduction

In this paper we generalize to higher dimensions a result by Kenyon and Kenyon [2] on the problem of adding *short cuts* to a polygonal chain on the plane, so that any two points on the chain are linked by a path composed of the chain and the short cuts, whose length is at most a constant times their Euclidean distance. More precisely, assume that we are given 2-dimensional polygonal chain, which can be represented as an Euclidean graph $G = (V, E)$. Let $P(E)$ be the (infinite) set of points that compose the edges in E. Let $d(u, v)$ denote the Euclidean distance between two points u and v. Let $wt(E)$ denote the sum of the Euclidean lengths of the edges in E. We restate below a simpler version of a theorem proved in [2].

Theorem 1 *Given a 2-dimensional polygonal chain $G = (V, E)$, a (possibly infinite) number of new edges E' can be added between points in $P(E)$ such that, (1) in the new structure there is a path between any points u and v in $P(E)$ of length at most $t_1 \cdot d(u, v)$, and (2) $wt(E') \leq c_1 \cdot wt(E)$. Here $t_1 > 1$ and $c_1 > 0$ are two absolute constants.*

We refer to t_1 as the *stretch factor* and the new edges as the *short cuts*. There are several interesting aspects to the above mentioned result. Firstly, though the number of segments added may be infinite, the algorithm only outputs a *finite length description* of the set of short cuts. The second aspect concerns the concept of *finite precision*. Suppose we only need to provide short cuts between pairs of points that are at least some $\epsilon > 0$ distance apart. In this case, their algorithm can be easily modified to only add a finite number of short cuts. The third aspect concerns a concept stronger than finite precision, which we call *finite extendibility*. In this, it should be possible to obtain a structure with smaller finite precision from one with a larger finite precision by simply *adding* more segments, without having to remove the previous ones. The total weight of the resulting structure should still be at most $c_1 \cdot wt(E)$. The algorithm in [2] can easily be modified to achieve finite extendibility.

Kenyon and Kenyon generalized their results to *rectifiable curves* on the plane, and also claim that their scheme could be generalized to deal with the case of planar graphs. However, they left open the problem of constructing short cuts for chains and Euclidean graphs in higher dimensional space.

In this paper we prove the following theorem and thus solve the problem of constructing short cuts for Euclidean graphs in higher dimensional space.

*Mathematical Sciences Dept., The University of Memphis; *e-mail*: dasg@next1.msci.memphis.edu; Supported in part by NSF Grant CCR-930-6822.

†Mathematical Sciences Dept., The University of Memphis; *e-mail*: giri@next1.msci.memphis.edu; Supported in part by NSF Grants INT-911-5870 and CCR-940-9752.

Theorem 2 *Given a k-dimensional connected Euclidean graph $G = (V, E)$ and any $t > 1$, a (possibly infinite) number of short cuts E' can be added between points in $P(E)$ such that, (1) in the new structure there is a path between any points u and v in $P(E)$ of length at most $t \cdot d(u, v)$, and (2) $wt(E') \leq c(t, k) \cdot wt(E)$, where $c(t, k) > 0$ is a constant dependent on t and k.*

We summarize the important aspects of our results.

1. Theorem 2 is a generalization of Theorem 1 to higher dimensions.

2. While Theorem 1 proves the existence of a positive stretch factor t_1, Theorem 2 is true for any given stretch factor $t > 1$. Of course, the smaller the t, the larger the constant factor in the total length bound of our short cuts.

3. In Section 2, we present an algorithm which adds short cuts such that the resulting structure has a small total weight, but only provides short paths between pairs of points that are not "too close" (finite precision). Unfortunately, it seems unlikely that this method can be modified to achieve either the finite extendibility or the finite description requirements. Later, in Section 3 we describe a more robust algorithm for constructing a structure which achieves finite precision. Moreover, this method can be easily modified to achieve either finite extendibility or finite description.

4. Theorem 2 can be easily generalized to smooth curves in higher dimensional space.

5. The methods in [2] are heavily dependent on planar geometry, and considerably different ideas are required to generalize the results to k-dimensional space. One of the main tools in our results is the k-dimensional spanner algorithm developed in [1, 3]. Given a set V of n points in k-dimensional space, and any $t' > 1$, an Euclidean graph $G = (V, E)$ is a *t'-spanner* if between every u and v in V, the shortest path in G has length at most $t' \cdot d(u, v)$. The spanner algorithm constructs a t'-spanner whose total edge length is at most $c'(t', k) \cdot wt(SMT(V))$, where $SMT(V)$ is a Steiner minimum tree of V and $c'(t', k) > 0$ is a constant dependent on t' and k.

As with the scheme in [2], we do not analyze the running time of our algorithm, other than showing that it runs in finite time. This is because the running time is very sensitive to the way the edges of G are configured in space, and in the case of finite precision, also on the required tolerance. For example, a pair of long, parallel edges will generate a lot of short cuts.

2 Achieving Finite Precision: The Grid Method

In this section we give an algorithm to achieve finite precision. It is quite straightforward; however it is unlikely that it can be modified to achieve either finite extendibility or finite descriptions.

We use a procedure called *grid*. Given a line segment (u, v) and any $\delta > 0$, this procedure introduces Steiner points $w_1, w_2, \ldots, w_p$ along the segment (u, v), such that $d(u, w_1) = d(w_1, w_2) = \cdots = d(w_{p-1} w_p) = \delta$, and the last fragment $d(w_p, v) \leq \delta$.

Our algorithm works as follows. Suppose $\epsilon > 0$ is the required tolerance. Select $\delta < \epsilon \cdot \left(\frac{t-\sqrt{t}}{1+\sqrt{t}}\right)$ and grid all the edges of the input graph $G = (V, E)$. Let V' be the union of the original vertices and the new Steiner points. Run the spanner algorithm on V' with parameter $t' = \sqrt{t}$. Let the edges of the spanner, E', be the short cuts to be added to G.

We first show that a stretch factor of t is indeed achieved.

Lemma 1 *Let u and v be two points on $P(E)$ such that $d(u,v) > \epsilon$. Then there is a path between them in the structure of length at most $t \cdot d(u,v)$.*

Proof: Consider any pair of points u and v in $P(E)$ such that $d(u,v) > \epsilon > \delta \cdot \left(\frac{1+\sqrt{t}}{t-\sqrt{t}}\right)$. We shall show that in the new structure, there is a short path between u and v. Note that from u (resp. v) one can get to a vertex in V', say u_1 (resp. v_1), by traveling no more than $\delta/2$ distance along the structure. But because of the spanner, there is a path from u_1 to v_1 of length $\leq t' \cdot d(u_1, v_1) = \sqrt{t} \cdot d(u_1, v_1) \leq \sqrt{t} \cdot (d(u,v) + \delta)$. Thus, there a path between u and v of length $\leq \sqrt{t} \cdot (d(u,v) + \delta) + \delta$. Let this path have length x. Thus $\frac{x}{d(u,v)} \leq \frac{\sqrt{t} \cdot (d(u,v)+\delta)+\delta}{d(u,v)}$. That is, $\frac{x}{d(u,v)} \leq \sqrt{t} + \frac{(\sqrt{t}+1)\cdot\delta}{d(u,v)}$. But we have assumed that $d(u,v) > \delta \cdot \left(\frac{1+\sqrt{t}}{t-\sqrt{t}}\right)$. Substituting in the right hand side of the above and simplifying, we get $\frac{x}{d(u,v)} < t$. □

We next estimate the total weight of the short cuts.

Lemma 2 *The total weight of the short cuts is at most $c'(\sqrt{t}, k) \cdot wt(E)$.*

Proof: $wt(E') \leq c'(t', k) \cdot wt(SMT(V'))$, by the properties of the spanner. But $wt(SMT(V')) \leq wt(E)$, since G connects all points in V'. Thus, for any δ, $wt(E') \leq c'(\sqrt{t}, k) \cdot wt(E)$. That is, $c(t,k) = c'(\sqrt{t}, k)$, which proves the lemma. □

This algorithm cannot be easily modified to achieve finite extendibility. If a smaller precision is later required, we cannot simply recompute from scratch and add new short cuts, because if we do not remove the previous short cuts, the structure will eventually get too heavy. Secondly, if no tolerance is allowed, it is not clear how the algorithm can be easily modified to output a finite length description of the possibly infinite set of short cuts. A more complex algorithm which addresses all these issues is presented in the next section.

3 Achieving Finite Precision, Extendibility, and Description

We present here another algorithm that achieves finite precision. Later we show how it can be easily modified to achieve finite extendibility. We also show how to modify it to run without any tolerance requirements, yet produce a finite description of the (possibly infinite) set of short cuts.

The algorithm takes as input a graph $G = (V, E)$, a stretch factor $t > 1$, and a tolerance $\epsilon > 0$. The algorithm constructs short cuts to satisfy this tolerance.

Step 1: Preliminaries

Let l be the length of the shortest edge in G. As a preliminary step, we insert (at most) two Steiner points u_1 and v_1 on each edge (u,v) of G, such that $d(u, u_1) = d(v, v_1) = l/2$. Let this set of Steiner points be S. We may now imagine the graph to be composed of the vertices $V \cup S$, where each original edge has been split into (at most) three new edges.

Step 2: Pairs of Points on Non-Adjacent Edges

Consider the graph obtained in Step 1. Two edges are *non-adjacent* if they do not share a common end point. Define the *distance* between a pair of non-adjacent edges as the length of the shortest line connecting a point on each. Let γ be the shortest such distance over all pairs of non-adjacent edges. Select δ such that $\delta < \gamma \cdot \left(\frac{t-\sqrt{t}}{1+\sqrt{t}}\right)$. Grid the edges of the graph using this δ, and construct a $\sqrt{t}$-spanner for the vertices in $V \cup S$ and the new Steiner points formed by the gridding process. Add the edges generated by the spanner as short cuts to the graph.

Lemma 3 *Let u and v be two points on non-adjacent edges. Then there is a path in the structure of length at most $t \cdot d(u,v)$.*

Proof: The proof is similar to the proof of Lemma 1 presented in the previous section. □

We next estimate the total weights of the short cuts added so far.

Lemma 4 *The total weight of the short cuts is at most $c'(\sqrt{t}, k) \cdot wt(E)$.*

Proof: The weight of a Steiner minimum tree of the vertices in $V \cup S$ and the Steiner points generated in Step 2 is at most $wt(E)$. So by spanner properties, the total weight of the short cuts is at most $c'(\sqrt{t}, k) \cdot wt(E)$. □

At this stage, the short cuts we have added provide short paths between all pairs of points belonging to non-adjacent edges. We now add short cuts to provide short paths between all of pairs of points belonging to adjacent edges.

Step 3: Pairs of Points on Adjacent Edges

Observe that we can immediately ignore those pairs of edges that share a common point in S. This is because there is already a short path (a straight line) between pairs of points on such edges.

We next consider for each vertex v of V, the edges incident to it. Note that each edge has length $l/2$. Select a small proper fraction α such that $\alpha \leq \frac{t-\sqrt{t}}{t+1}$. Consider the decreasing geometric series $(l/2), (l/2)\alpha, (l/2)\alpha^2, \ldots, (l/2)\alpha^i, \cdots$. Imagine an infinite series of nested concentric spheres centered at v with these radii. For every edge e incident to v, let $v_0, v_1, \ldots, v_i, \ldots$ correspond to the points of intersection of each sphere's surface with the edge. Let e_i represent the subedge (v_i, v_{i+1}), and E_i represent the set of these subedges (i.e., the fragments of all edges incident to v that lie between the consecutive concentric spheres).

The algorithm proceeds by performing the following operations for each vertex v of V. (Intuitively, for each vertex v we will compute a series of "concentric and overlapping" spanners, up to a certain depth determined by the tolerance).

Let m be the smallest integer such that $(l/2)\alpha^m + (l/2)\alpha^{m-1} < \epsilon$. For each edge incident to v, for $i = 0, 1, \ldots, m$, compute each intersection point v_i. For $i = 0, 1, \ldots m-1$, compute each set E_i. For $i = 0, 1, \ldots, m-2$, let the set F_i be defined as $F_i = E_i \cup E_{i+1}$. Let γ_i be the shortest distance between all pairs of non-adjacent edges in F_i. Select δ_i such that $\delta_i < \gamma_i \cdot \left(\frac{\sqrt{t}-\sqrt[4]{t}}{1+\sqrt[4]{t}}\right)$. Grid the edges of each F_i using this δ_i, and construct a $\sqrt[4]{t}$-spanner for the end vertices and the new Steiner points of F_i formed by the gridding process. Add the edges generated by the spanner as short cuts to the graph.

Note: A crucial observation is that these spanners will be *scaled versions* of one another; thus we only need compute F_0, compute its spanner, and scale it in size as well as position to describe the others. This has an important part to play when we address finite description concerns later.

Lemma 5 *Let u and w be two points on adjacent edges incident to v, such that $d(u,w) > \epsilon$. Then there is a path between them in the structure of length at most $t \cdot d(u,w)$.*

Proof: The proof has several cases.

Case 1: (*u and w both belong to the same F_i*). In a manner similar to the proof of Lemma 1, we can show that there is a path from u to v in the structure of length at most $\sqrt{t} \cdot d(u,w)$ (the only difference is that we will be using $\sqrt{t}$ instead of t, $\sqrt[4]{t}$ instead of $\sqrt{t}$, and δ_i instead of δ).

Since $\sqrt{t} < t$, Lemma 5 is proved for this case.

Case 2: (*u belongs to some E_i, w belongs to some E_j, $j > i+1$*). Let w belong to e, where e is an edge of the graph after Step 2. Consider v_{i+2}, the intersection point of this edge with the sphere of radius $(l/2)\alpha^{i+2}$. Clearly v_{i+2} belongs to F_i, *but does not belong to E_i*. We construct a path from u to w as follows: use the spanner edges of F_i to go from u to v_{i+2}, then go from v_{i+2} to w along e. Let the length of this path be x. Clearly $x \leq d(u, v_{i+2}) \cdot \sqrt{t} + d(v_{i+2}, w)$. That is, $x \leq (d(u,w) + d(w, v_{i+2})) \cdot \sqrt{t} + d(v_{i+2}, w)$. Dividing both sides by $d(u,w)$ we get $\frac{x}{d(u,w)} \leq \left(1 + \frac{d(w,v_{i+2})}{d(u,w)}\right) \cdot \sqrt{t} + \frac{d(w,v_{i+2})}{d(u,w)}$. But it is also easy to verify that $\frac{d(w,v_{i+2})}{d(u,w)} \leq \frac{\alpha}{1-\alpha}$. Substituting in the previous inequality, we get $\frac{x}{d(u,w)} \leq \left(1 + \frac{\alpha}{1-\alpha}\right) \cdot \sqrt{t} + \frac{\alpha}{1-\alpha}$. If we substitute the value of α selected earlier and simply, we get $\frac{x}{d(u,w)} \leq t$, which proves Lemma 5 for this case.

Case 3: (*w is within the sphere of radius $(l/2)\alpha^m$*). We have assumed that $d(u,w) > \epsilon$. From the way m has been selected earlier, it is clear that u has to be outside the sphere of radius $(l/2)\alpha^{m-1}$. Assume u belongs to an edge of E_i, $i < m-1$. Let v_{i+2} be the intersection of the original edge on which w lies and the sphere of radius $(l/2)\alpha^{i+2}$. Clearly u and v_{i+2} both belong to F_i. We construct a path from u to w as follows: use the spanner edges of F_i to go from u to v_{i+2}, then go from v_{i+2} to w along e. Let the length of this path be x. In a manner similar to Case 2, it can be shown that $\frac{x}{d(u,w)} \leq t$, which proves Lemma 5 for this final case. □

We next estimate the weight of the short cuts created in Step 3.

Lemma 6 *The total weight of the short cuts created in Step 3 is at most $c''(t,k) \cdot wt(E)$, where $c''(t,k)$ is a function of t and k.*

Proof: Let us consider any vertex v of the original graph. Let H be the set of edges adjacent to it after Step 1 (each such edge has length $l/2$). Let H_i be the portions of these edges that are within the sphere of radius $(l/2)\alpha^i$. Clearly $wt(H_i) = \alpha^i \cdot wt(H)$.

Consider any F_i. It is clear that $wt(F_i) < wt(H_i)$. In Step 3, each F_i is gridded with Steiner points, and then a spanner is created for these points. But H_i is itself a Steiner tree for these points. Thus the weight of the spanner for F_i is $< c'(\sqrt[4]{t}, k) \cdot wt(H_i)$. That is, $< \alpha^i \cdot c'(\sqrt[4]{t}, k) \cdot wt(H)$. Summing over $i = 0, 1, \ldots, m-2$, we see that the weight for all the short cuts added when considering vertex v is a geometric series whose sum is $< \left(\frac{1}{1-\alpha}\right) \cdot c'(\sqrt[4]{t}, k) \cdot wt(H)$. Finally, if we sum this over all the vertices, we see that the total weight of all short cuts added in Step 3 is no more than $c''(t,k) \cdot wt(E)$, where $c'''(t,k)$ is some function of t and k. $\square$

We now investigate how to achieve finite extendibility. Assume that the above algorithm has been originally run with a tolerance of $\epsilon > 0$. If later an even smaller tolerance is required, say $\epsilon_1 < \epsilon$, we only need to *add* more short cuts. This is done as follows.

We recompute a new and larger m_1, defined to be the smallest integer such that $(l/2)\alpha^{m_1} + (l/2)\alpha^{m_1 - 1} < \epsilon_1$. For every vertex v, we do the following. For $i = m-2, \ldots, m_1 - 2$, compute new sets F_i, and for each such set, compute a spanner as described earlier in Step 3. The edges of all these new spanners are then added as new short cuts to the structure.

It is easy to see that the new structure will still have a stretch factor of t (but with the additional capability of handling smaller tolerances). Moreover, the weight will still remain within a constant multiple of $wt(E)$, since it will be upper bounded by the geometric series summation described earlier in Lemma 6.

Let us now address the problem of finite description of the output. In this case, there is no ϵ in the input, hence no integer m is computed by the algorithm to bound the terms in the geometric series. Thus it may seem that the algorithm will have to compute an infinite number of spanners in Step 3, one for each F_i in the infinite series. But as we have observed earlier, these spanners are simply scaled copies of one another. Thus a finite description of the set of short cuts can be achieved by computing the spanner for F_0, and specifying the geometric series (by simply specifying α, the ratio of successive terms), which determines the positions and sizes of the succeedingly smaller spanners.

Incidently, the above observations formally prove Theorem 2.

4 Generalization to Smooth Curves

Our problem for Euclidean graphs can be easily generalized to smooth curves in k-dimensional space. As in [2], given a *smooth curve* S, define $\beta > 0$ so that any disk of diameter less than β intersects S in a connected set, and in any disk of diameter β centered on a point of S, the curve S does not deviate much from a straight line. We omit the details, but essentially since the curve looks locally flat, we can grid it with a parameter δ which will be dependent on t and β (similar to what was done in Step 2 of the previous section). A spanner of the grid points, then, defines the short cuts.

References

[1] Arya S.; Das G.; Mount D.; Salowe J.; Smid M., Euclidean Spanners: Short, Thin and Lanky, to be presented at the *27th ACM STOC*, 1995.

[2] C. Kenyon; R. Kenyon, How to Take Short Cuts, *Discrete and Computational Geometry*, 8:251-264 (1992).

[3] Das G.; Narasimhan G.; Salowe J., A New Way to Weigh Malnourished Euclidean Graphs, *5th ACM-SIAM SODA*, (1995), 215-222.