

A Fast Octree-Based Algorithm for Computing Ropelength

Ted Ashton* and Jason Cantarella†

Department of Mathematics, University of Georgia, Athens, GA 30602

(Dated: December 31, 2003; Revised: September 16, 2004)

The *ropelength* of a space curve is usually defined as the quotient of its length by its *thickness*: the radius of the largest embedded tube around the knot. This idea was extended to space polygons by Eric Rawdon, who gave a definition of ropelength in terms of doubly-critical self-distances (local minima of the distance function on pairs of points on the polygon) and a function of the exterior angles of the polygon.

A naive algorithm for finding the doubly-critical self-distances of an n -edge polygon involves comparing each pair of edges, and so takes $O(n^2)$ time. In this paper, we describe an improved algorithm, based on the notion of *octrees*, which runs in $O(n \log n)$ time. The speed of the ropelength computation controls the performance of ropelength-minimizing programs such as Rawdon and Piatek’s TOROS.

An implementation of our algorithm is freely available under the Gnu Public License.

Keywords: ropelength, ideal knot, tight knot, minrad, polygonal injectivity radius, octree, TOROS, RIDGERUNNER

1. INTRODUCTION

For a smooth curve in 3-space, *ropelength* is the quotient of the length of the curve by its *thickness*: the diameter of the largest embedded tube around the curve. Minimizing ropelength is the same as fixing the diameter of the tube and minimizing its length—if the tube is knotted, we are pulling the knot tight, and so the minimum ropelength curves in any knot type are often called *tight* knots. Since the problem is such a natural one, the definition of thickness has been discovered and rediscovered by several authors [1, 3, 12], with the earliest results known (to these authors) on the problem credited to Krötenheerdt and Veit in 1976 [11].

In the past decade, there has been a great deal of interest in exploring the geometry of tight knots; the definition of thickness has been refined and fully understood [9], it has been shown that $C^{1,1}$ minimizers exist in each knot type [5, 7, 8], some minimizing links have been found [5], and a theory of ropelength criticality has started to emerge [4, 19]. The development of this theory has been fueled by a steady stream of numerical data on ropelength minimizers, from Pieranski’s original SONO algorithm [13] and Rawdon’s TOROS [14], to second-generation efforts such as Smutny and Maddocks’ biarc computations [17] and the RIDGERUNNER project of Cantarella, Piatek, and Rawdon.

All of these algorithms have in their innermost loops a computation of the ropelength of a curve in 3-space. Intuitively, the thickness of a tube is controlled locally by the curvature of the core curve, and globally by the approach of “distant” sections of the tube. Rawdon, in his thesis [15], defined a radius of curvature for a corner of a polygon. A given corner has two circles which are tangent to both incident edges and tangent to one of the edges at its center. He proved that we can define a sensible polygonal radius of curvature as the radius of the smaller of those two circles.

More precisely:

Definition 1. If P_n is a polygonal curve in \mathbb{R}^3 with edges e_1, \dots, e_n , and α_i is the exterior angle of the polygon made by edges e_i and e_{i+1} , then let

$$\text{minRad}(P_n) = \min_{i \in 1, \dots, n} \left\{ \frac{|e_i|}{2 \tan\left(\frac{\alpha_i}{2}\right)}, \frac{|e_{i+1}|}{2 \tan\left(\frac{\alpha_i}{2}\right)} \right\}. \quad (1)$$

where we take $e_{n+1} = e_1$ if the polygon is a closed curve, and take $i \in 1, \dots, n-1$ otherwise.

Then we define the distance function $D : P_n \times P_n \rightarrow \mathbb{R}$ by $D(x, y) = |x - y|$, and say that a chord xy of P_n is a *pair of closest approach* of P_n if it is a local minimum of the distance function. The length of the shortest such chord is denoted $\text{POCA}(P_n)$. We then define *thickness* by

$$\text{SThi}_\lambda(P_n) = \min \left\{ \frac{2 \text{minRad}(P_n)}{\lambda}, \text{POCA}(P_n) \right\}. \quad (2)$$

We note that this is *not* Rawdon’s original definition of polygonal thickness [15], which used the minimum distance between all “doubly-turning” pairs of the polygon and not just those pairs which were local minima of self-distance. The difference is made clear when one considers a regular $2n+1$ -gon in the plane, which has no pairs of closest approach. In this case, minRad compensates for the absence of POCA lengths and the results of the definitions are identical. In fact, we conjecture that the definitions are equivalent in general, and intend to address the subject in a future publication.

The factor λ in our definition requires some explanation. For a smooth curve P , fixing $\text{SThi}_\lambda(P) = 1$ models a tube of unit diameter whose diameter of curvature is bounded below by λ . When $\lambda = 1$, the tube is perfectly flexible: it can make a “U”-shaped bend with a semicircle of unit diameter and two parallel sections of tube in contact along a straight line. When $\lambda > 1$, we model a more realistic “stiff” rope whose minimum diameter of curvature is somewhat higher.

As computing the radius of curvature at a given corner only involves the edges incident to that corner, computing $\text{minRad}(P_n)$ requires only $O(n)$ time. On the other hand, all previous efforts to compute thickness have used some variant of Algorithm 1 on page 2 for computing $\text{POCA}(P_n)$. This algorithm is clearly $O(n^2)$. So we have focused our attention on improving the $\text{POCA}(P_n)$ calculation.

*Email: ashted@uga.edu

†Email: cantarel@math.uga.edu

```

for  $i = 1$  to  $n$  do
  for  $j = i + 1$  to  $n$  do
    check  $e_i$  and  $e_j$  for local min chords;
    compare to previous shortest local min chord;
  end
end

```

Algorithm 1: Standard Algorithm for $\text{POCA}(P_n)$.

Our algorithm concentrates on reducing the total number of edge-edge checks performed by grouping the edges according to their positions in space into a data structure known in computer graphics as an *octree*. We will use the octree to optimize the inner loop of Algorithm 1, and show that we can isolate a constant-size set of candidate e_j 's for any given e_i in time $O(\log n)$. The new algorithm will then perform $O(n \log n)$ edge-edge checks, and one octree construction (which will also require time $O(n \log n)$), for a significant improvement in speed.

At this point, the reader might be concerned that our (significantly more complicated) algorithm loses its theoretical advantage over Algorithm 1 for most n of practical interest due to large constants on the $O(n \log n)$ terms. In fact, this is not so (and we invite the reader to consult the performance charts of Section 6 for an empirical verification of this claim).

2. EDGE-EDGE CHECKS

The quantity $\text{POCA}(P_n)$ is defined to be the smallest local minimum of the distance function $D(x, y)$ on pairs of points on the polygon P_n . To understand it, we first make an observation about the nature of these local minima.

Lemma 2.1. *If we orient the curve P_n and let $T^-(x)$, $T^+(x)$ denote the inward and outward tangent vectors of P_n at x (they are different if and only if x is a vertex with nonzero exterior angle). Every pair xy which locally minimizes $D : P_n \times P_n \rightarrow \mathbb{R}$ has*

$$T^-(x) \cdot (y - x) \geq 0 \quad T^+(x) \cdot (y - x) \leq 0 \quad (3)$$

$$T^-(y) \cdot (x - y) \geq 0 \quad T^+(y) \cdot (x - y) \leq 0. \quad (4)$$

We note that if x is in the interior of an edge, then the above relations force $T^\pm(x) \cdot (x - y) = 0$.

Proof. There are three cases: either both x and y are on the interior of an edge, one is an edge point and one a vertex, or both are vertices, as shown in Figure 1. At x , the distance from y must not decrease to first order as one moves away from x in either direction along the curve: a computation verifies that this is equivalent to the first line of the statement of the Lemma. A similar argument at y completes the proof. \square

We now make a definition:

Definition 2. The i^{th} *ramp*, R_i of a polygonal curve P_n is the union of the planes through edge $e_i = v_i v_{i+1}$ with normal vector $v_{i+1} - v_i$, together with the wedge of vectors w defined by the inequalities

$$(w - v_i) \cdot T^-(v_i) \geq 0 \quad (w - v_i) \cdot T^+(v_i) \leq 0. \quad (5)$$




FIG. 1: We see the three cases in the proof of Lemma 2.1, from left to right an *edge-edge* pair, a *vertex-edge* pair, and a *vertex-vertex* pair. In the center and right figures we see $T^-(x)$ and $T^+(x)$, and in the righthand figure we also see $T^-(y)$. We have not drawn $T^+(y)$, but it would be colinear with $v_j v_{j+1}$ as with $T^+(x)$.




FIG. 2: The shaded area represents the region of space in which the second point y of a locally minimal pair xy can lie when x is on the edge e_i or is the vertex v_i . This region consists of the infinite slab of parallel planes normal to e_i which pass through e_i , together with the wedge extending from vertex v_i in the outward direction from the vertex.

See Figure 2.

This leads naturally to the Lemma:

Lemma 2.2. *If xy is a pair of points on P_n which locally minimizes D , and x is on the half-open edge $e_i - \{v_{i+1}\}$, then y is in the ramp R_i .*

Proof. If x is in the interior of the edge, Lemma 2.1 implies that xy must be perpendicular to e_i , and hence that y is in the union of normal planes through e_i . If x is at v_i , the inequalities above are those of the statement of Lemma 2.1. \square

It may seem like we have only rephrased Lemma 2.1. In fact, we have gained an important geometric insight about the problem— for any edge e_i , the ramp R_i will probably be very close to a thin slab which only intersects the remainder of P_n in a few places (see Figure 3 on page 3). If we can isolate these intersections quickly, we can complete the task of finding $\text{POCA}(P_n)$ by a more detailed comparison of these candidates to e_i .

3. THE OCTREE DATA STRUCTURE

With the discussion in Section 2, we have reduced the problem of identifying edges e_j which may form locally minimal pairs with points on edge e_i to the problem of finding which




FIG. 3: A typical ramp in a trefoil of about 90 edges consists of a very thin slab which only intersects the remainder of the knot in a few places. If we could isolate ramp-knot intersections quickly, it would reduce the number of edge-edge checks required to find the shortest POCA.

edges of P_n intersect e_i 's ramp. To do so efficiently, we will need a new data structure for P_n : the octree [10].

The octree representation of a collection of points in space is a tree where each node represents the bounding box of a subset of that collection. The eight daughter nodes of a parent represent the bounding boxes of subsets of the points in the parent box created by dividing that point set in two in each of the coordinate directions. The most detailed octree representation of a point set has leaf nodes which each contain a single point but it is common to stop subdividing when the point sets are smaller than some fixed number. Figure 4 illustrates three levels of this process for a set of points in the plane, while Figure 5 shows the resulting tree.

From the description above, one can observe that it is easy to build an octree using the recursive procedure of Algorithm 2.

Data : A node of the tree, a corresponding box B , and a list of points.

Result : An octree representation of this point list.

if the list of edges is shorter than some fixed N **then**

- assign these points to this node;
- make this node a leaf of the tree;
- return;

else

- partition B into 8 child boxes;
- for** each child box **do**
 - create a sublist of points intersecting that box;
 - recurse if this sublist is nonempty;
- end**

end

Algorithm 2: One way to Build an Octree

For an n -point dataset, if one chooses each box partition so that no child box contains more than half the total number of points in the parent box, the number of levels in this tree is less than $\log_2 n$, and one expects this algorithm to run in $O(n \log_2 n)$ time. However, this algorithm involves a non-trivial amount of overhead in keeping track of lists of points, and making procedure calls. Since we are very concerned

with the final performance of our implementation, we now present a more insightful octree construction algorithm which has the same asymptotic time bound of $O(n \log_2 n)$, but is much faster than Algorithm 2 in practice.

To describe the new algorithm, we start with a numbering scheme. As we mentioned before, it is conventional to denote the upper right quadrant of the plane by the number 1, and proceed counterclockwise to the fourth quadrant on the lower right. A more natural numbering scheme assigns each quadrant a 2-digit binary number, $d_x d_y$, where $d_x = 0$ has lower values of x (the left hand side) and $d_x = 1$ has higher values of x (the right hand side), while $d_y = 0$ denotes lower values of y (the bottom half), while $d_y = 1$ denotes higher values of y (the top half). For octants in 3-space, we could assign three digit binary numbers $d_x d_y d_z$ similarly.

Now consider the process of quadtree construction again. At the first subdivision, we divide the point set in two parts by x -coordinate and by y -coordinate. This gives us four groups of points, which we can number as above by the 2-digit binary numbers $d_x d_y$. These groups are the members of the 4 boxes in the next level of the tree, as we saw above.

But there is something else to notice here: If the collection of points is sorted by x and by y , the digits d_x and d_y for any particular point are the most significant binary digits of that point's *position* in the sorted array. Further, if we continue to subdivide the points into fourths by x and y , the next pair of binary digits associated to each point, $d_{x_1} d_{y_1}$ will be the next pair of binary digits in that point's position in the x and y arrays as well. Again, for octrees the situation is similar, but we sort by z as well, and create a sequence of 3-digit binary (or 1-digit octal) numbers.

Continuing this process, we see that each point in the collection has a unique *octal tag* generated by interleaving the digits of its position in the sorted x , y , and z arrays. This tag specifies its position in the octree. Further, if we made a least-first traversal of the octree (descending to octants in the order of their octal labels), the order in which we would encounter the points would be by increasing octal tags. These observations give rise to a new octree-building algorithm:

Data : A list of points in \mathbb{R}^3 .

Result : An octree representation of point list.

- Sort the points by x , y , and z coordinates;
- Shuffle binary digits of array positions to create octal tags;
- Sort again by octal tags;
- Build tree from this traversal-ordered list;

Algorithm 3: A faster Octree-building algorithm

The problem of building a tree from a traversal-ordered list of its contents is a standard one in computer science. Our particular solution is discussed in some detail in Section 5 below. We note that building the octree from the list also has time complexity $O(n \log n)$, since every node in the octree must be visited, but that this algorithm is still much faster than the previous method of octree construction (Algorithm 2). We are among many rediscoverers of this method of octree construction, which traces its roots to the "linear quadtree" construction of Gargantini [6].




FIG. 4: From left to right, these pictures show three stages in the construction of a quadtree representation (the planar version of an octree representation) of a set of points. On the left, the bounding box of the entire point set is computed. This is the root of the tree. In the center, we see the points divided in two by x and y coordinates, and then grouped by quadrant into four subcollections, with bounding boxes as shown. On the right, we again divide the subcollections and group into subquadrants. The resulting 3-level tree is shown in Figure 5.




FIG. 5: This picture shows the quadtree constructed in Figure 4 in a more familiar form. The boxes are labelled according to the usual numbering convention for quadrants of the plane, where the first quadrant is on the top right and numbering proceeds counter-clockwise. The final numbers show the number of points in each leaf box, and should be compared to the boxes shown in Figure 4 in the right-hand image.

4. THE CORE OF THE ALGORITHM

We can now describe our algorithm. Given any e_i , we must identify all edges e_j which might be part of a shortest POCA with e_i . Such edges must obey two conditions: they must intersect e_i 's ramp, and they must be closer to e_i than the shortest POCA found so far. Since both conditions can be checked for sub-boxes of the octree, we can use them to eliminate groups of e_j from consideration before performing edge-edge checks.

In pseudo-code, this is a collection of n calls to the (recursive) Algorithm 4 (one for each e_i). We refer to the entire algorithm (minRad computation, octree construction by Algorithm 3, and calls to Algorithm 4 for each edge) as **Octrope**.

Each call to this algorithm might require it to traverse the entire depth of the octree before reaching leaf nodes and performing the edge-edge checks. Yet this depth is bounded above by $\log_2 n$, so the expected running time for the algorithm is $O(\log n)$. In pathological cases, many or all of the boxes may intersect the ramp. If all the boxes intersect all the ramps, this algorithm may be asymptotically slower than the naive one: we are forced to visit $O(n \log n)$ tree nodes against each of n edges, for a total time complexity of $O(n^2 \log n)$.

Data : An octree node, the current minimum strut length l and a ramp from e_i .

Result : All struts between e_i and edges in this octree.

```

if this box is within  $l$  of  $e_i$  then
  if this box intersects the ramp from  $e_i$  then
    if this box is a leaf then
      check the (at most  $N$ ) edges against  $e_i$ ;
      if struts are found then
        update  $l$ ;
        return struts;
      end
    else
      for each nonempty child node do
        recurse on the child node and the ramp;
      end
    end
  end
end

```

Algorithm 4: Recursively identifying candidate e_j 's.

We have seen Algorithm 1 outperform Algorithm 4 only for a particularly bad class of examples: knots formed by connecting vertices chosen at random inside a fixed volume. (See Section 6 for details.)

5. IMPLEMENTATION ISSUES

While being able to replace an $O(n^2)$ algorithm with one which is $O(n \log n)$ will certainly save time for large enough values of n , there is no guarantee that this will help with problems of practical size. Indeed, Algorithm 4 threatens to consume a fair amount of overhead, while Algorithm 1 involves only edge-edge checks, which could be coded very efficiently. So in this section we turn our attention from the “ $n \log n$ ” to its multiplier — from mathematics to program design. In this discussion, we’ll refer to function names and prototypes from our publically available library version of **Octrope**, which is called `liboctrope`.




FIG. 6: Our example is a polygonal Hopf-link approximation composed of two regular pentagons. The lighter lines are the 9 struts for which the **Octrope** algorithm is searching. They extend from the midpoints of one side of each interlocked polygon to all the midpoints of the edges of the other.

Component	
1	2
$v_{00} = (14.5, 20, 0)$	$v_{10} = (0, 0, 14.5)$
$v_{01} = (23.5, -7.6, 0)$	$v_{11} = (0, 27.6, 23.5)$
$v_{02} = (0, -24.7, 0)$	$v_{12} = (0, 44.7, 0)$
$v_{03} = (-23.5, -7.6, 0)$	$v_{13} = (0, 27.6, -23.5)$
$v_{04} = (-14.5, 20, 0)$	$v_{14} = (0, 0, -14.5)$

TABLE I: The approximate vertices of the pentagonal Hopf link shown in Figure 6 are given here. We have rounded the numbers to the nearest tenth to simplify the table. This does not affect the octree-building procedure under discussion, but would change the strut picture shown in the figure above.

The depth of the octree. Since searching the octree involves some overhead, it is to be expected that we will not get the best performance from the deepest octree. Rather, we expect it to be more efficient to group some number of edges in each box and do simple checks between the current edge and the edges in an implicated leaf box.

We implement this by bounding the maximum number of levels in the tree by some ℓ . This can be set by the user, using the `octree_set_levels()` call or it will default to $\ell = \lceil \frac{3}{4} \log_2 n \rceil$ (where $\lceil r \rceil$ is the least integer greater than or equal to r), a formula at which we arrived empirically. It also uses m , the maximum number of edges in any leaf box, which is given by $m = \lceil \frac{n}{2^{\ell-1}} \rceil$.

A concrete example. We will now trace through our implementation of the fast octree construction procedure of Algorithm 3 for a particular example: a Hopf link where each edge is given by a regular pentagon (see Figure 6 and Table I). In this example, ℓ would default to 3 and m would then also equal 3.

Sorting edges. The algorithm begins by gathering all of the edges into a single n -element array which we call `by_oct`. To avoid double-checking edge pairs later on, we need each edge to “belong” to only one of the leaf boxes in our tree. So we identify each edge by its midpoint and store that, as well

as the edge’s length and its starting vertex in `by_oct`.

As we create `by_oct`, we also build `by_x`, `by_y` and `by_z`, three n -element arrays of pointers to the elements of `by_oct`. We then sort these by x , y , and z order. The result is shown in Table II.

We divide `by_x`, `by_z`, and `by_z` into sections of m points each (shown in Table II by spacing) and walk through them, labeling the edges with the binary numbers of the sections in which they lie in the following unusual fashion: if $x = x_1 x_2 \cdots x_{\ell-1}$, $y = y_1 \cdots y_{\ell-1}$ and $z = z_1 \cdots z_{\ell-1}$ are the respective box numbers and their binary representations, we interleave those bits to produce a single octal number, $z_1 y_1 x_1 z_2 y_2 \cdots y_{\ell-1} x_{\ell-1}$. This is the *octal tag* of Section 3 above¹. We then sort `by_oct` by that octal tag, as shown in Table III on page 6.

As we discussed in Section 3, the sorted `by_oct` array is in the same order as that of a traversal of the full octree.

Building the tree. The actual building of the octree can be approached in various directions. We could simply use the `by_oct` array with no further indexing, traversing it with binary searches (an approach which saves space at the expense of time). On the other hand, if we are to index it, we can build our index in a top-down fashion, establishing the root node and building out to the leaf nodes. We can build in a bottom-up fashion, partitioning off parts of `by_oct` as leaf nodes and collecting them together in groups until we reach a single top node. Or we can (and do) use a “sideways” or “limb-by-limb” approach. We take an array of ℓ box pointers and on them build the “left-hand limb” of the tree, all the way from the smallest numbered leaf box down to the root box. Each of the boxes knows its first edge in `by_oct` and how many edges it has (which are grouped together thanks to the octal sort).

Then we walk once through `by_oct`, watching the octal tags. As long as the tag is the same as the one before it, we simply increment the count of edges in that box. When it changes, we do a binary XOR with the previous tag to see how much they differ (that is, which of the octal digits changed). That tells how many of the boxes in this “limb” are complete. After some cleanup (which may include pruning the “limb”) we leave those boxes and create the new ones necessary to hold this edge. Figure 7 on page 7 shows our example tree after this process is complete.

Searching the tree. We have now created the tree and can move into using it. We can now check each edge and its ramp against the tree, looking for leaf boxes on which to run edge-edge checks. Since struts are symmetric, we do not ever want to compare the same pair of edges twice. To avoid this we do the edge-edge check only if the edge in question precedes our

¹ To construct octal tags, we take a single pass simultaneously through `by_x`, `by_y`, and `by_z`, starting with the second box, which has binary tag $00 \cdots 01$. As we walk the arrays, we spread the bits of the box number apart (e.g. $1101 \rightarrow 100100001$) using a lookup table similar to that of Shaffer [16], shift them left 1 bit for y or 2 bits for z , and OR them with the tag constructed so far. The tags are thus built up over time and guaranteed to be correct only when we reach the end of the pass.

by_x	by_y	by_z
$e_{03} : (-19, 6.2, 0)$	$e_{02} : (-11.75, -16.15, 0)$	$e_{13} : (0, 13.8, -19)$
$e_{02} : (-11.75, -16.15, 0)$	$e_{01} : (11.75, -16.15, 0)$	$e_{12} : (0, 36.15, -11.75)$
$e_{04} : (0, 20, 0)$	$e_{14} : (0, 0, 0)$	$e_{00} : (19, 6.2, 0)$
$e_{10} : (0, 13.8, 19)$	$e_{03} : (-19, 6.2, 0)$	$e_{01} : (11.75, -16.15, 0)$
$e_{11} : (0, 36.15, 11.75)$	$e_{00} : (19, 6.2, 0)$	$e_{02} : (-11.75, -16.15, 0)$
$e_{12} : (0, 36.15, -11.75)$	$e_{10} : (0, 13.8, 19)$	$e_{03} : (-19, 6.2, 0)$
$e_{13} : (0, 13.8, -19)$	$e_{13} : (0, 13.8, -19)$	$e_{04} : (0, 20, 0)$
$e_{14} : (0, 0, 0)$	$e_{04} : (0, 20, 0)$	$e_{14} : (0, 0, 0)$
$e_{01} : (11.75, -16.15, 0)$	$e_{11} : (0, 36.15, 11.75)$	$e_{11} : (0, 36.15, 11.75)$
$e_{00} : (19, 6.2, 0)$	$e_{12} : (0, 36.15, -11.75)$	$e_{10} : (0, 13.8, 19)$

TABLE II: In this table, we see the 10 edges of the pentagons in Table I sorted by x , y , and z . The edges are sorted by their midpoints, and numbered by the index of their first vertices. The spacing reminds us that since $m = 3$, we are grouping the midpoints by threes when constructing boxes.

edge	x -box	y -box	z -box	bits	octal	decimal
e_{02}	0	0	1	000100	04 ₈	4
e_{03}	0	1	1	000110	06 ₈	6
e_{00}	3	1	0	001011	13 ₈	11
e_{01}	2	0	1	001100	14 ₈	12
e_{12}	1	3	0	010011	23 ₈	19
e_{13}	2	2	0	011000	30 ₈	24
e_{10}	1	1	3	100111	47 ₈	39
e_{14}	2	0	2	101000	50 ₈	40
e_{04}	0	2	2	110000	60 ₈	48
e_{11}	1	2	2	110001	61 ₈	49

TABLE III: This table contains the edges with their box numbers in the x , y , and z directions, the binary numbers generated by interleaving the bits of these box numbers, and the corresponding octal tags (in octal and decimal). The data is sorted by octal tag, and so appears in the same order in which it appears in the `by_oct` array. In principle, as many as m edges can share an octal tag, which means that they occupy the same leaf node of the resulting octree, but this does not happen in our example.

chosen edge in `by_oct`. By so doing, we can eliminate entire boxes because their lowest edges is not in range. The savings from this technique has been significant.

6. PERFORMANCE

We tested our algorithm using a 1.25 Ghz *G4* Macintosh computer on high-resolution discretizations of simple knots and on random walks. To make the tests, we compared the run times between `liboctrope` with the tree depth set to 1 (forcing $\frac{n(n-3)}{2}$ edge-edge checks) and with the default tree depth of $\ell = \lceil \frac{3}{4} \log_2 n \rceil$. We compiled our code with `gcc 3.3` and used the `-O3` option.

For both of these classes of knots, `liboctrope` was much faster than our reference implementation of Algorithm 1. Figure 8 shows the relative performance of the two algorithms on trefoil knots. It is also worthwhile to compare the performance of our code for a 2499-edge random walk at different

	Algorithm		
	Standard	Octrope	Max depth
Octree levels	1	9	13
Edge-edge checks	3, 121, 251	32, 033	8189
Box/ramp checks	0	51, 131	93, 187
Time	1.9 sec	0.16 sec	0.22 sec

TABLE IV: This table compares the performance of the `liboctrope` library on a 2500-edge random walk at three levels of tree depth: 1 (the standard $O(n^2)$ algorithm), $\lceil \frac{3}{4} \log_2 n \rceil$ (**Octrope**), and 13 (the maximum resolution). We see a trade-off between edge-edge checks and box/ramp checks as the octree resolution increases. Increasing the number of levels in the octree from 9 to 13 cuts the number of final edge-edge checks performed by a factor of 4, but doubles the number of box/ramp checks. Since the box/ramp checks are more computationally expensive, this is not a favorable ratio, and the overall execution time increases.

The data shows that we have been very effective at reducing the number of edge-edge checks. On average, **Octrope** compares each edge to less than 13 carefully chosen candidates when searching for struts.

levels of octree resolution, as in Table IV.

It is worth noting that **Octrope** does not always outperform the standard algorithm. For instance, for random knots constructed by choosing vertices inside a fixed volume, neither ramp-checking nor distance checking eliminates a significant number of pairs from consideration. The edges are simply too long, and pass too close to one another to decide in advance which pairs are likely to control thickness. But even in this case, `liboctrope` was only a few times slower than the standard algorithm.

Currently, minimizing the ropelength of 850-edge knots by simulated annealing is a relatively taxing task, requiring a few weeks of computer time on a standard desktop machine. Our timings above indicate that a simulated annealer built around `liboctrope` could do the job 5 times faster, or, alternatively, minimize the ropelength of a 2000-edge knot in the same time.

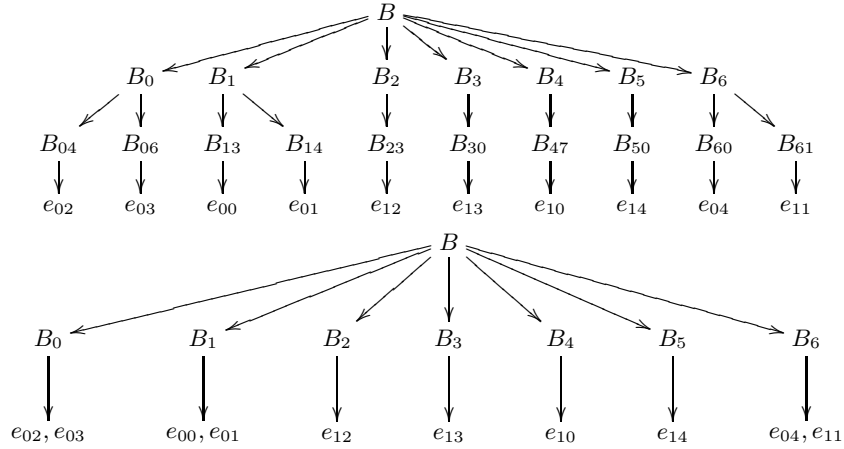


FIG. 7: These two trees show the octree as initially constructed (top) and after our pruning procedure (bottom). This has grouped some edges together in single nodes (such as e_{02} and e_{03}) and deleted some nodes with only one child (such as B_{23}). It is desirable to eliminate extra nodes of this kind, since even though we keep the octal tree as compact in memory as possible, each jump from node to node runs the risk of straying outside the memory cache of the processor and incurring a delay as more information is loaded from main memory.




FIG. 8: The two plots show the time in seconds required to find ropelength for trefoils (solid lines) and random walks (dashed lines) as a function of edgelenh. The timings were computed on a 1.25 Ghz Macintosh G4 computer, and represent averages over 10 runs (for times above one second), or 100 runs (for times below one second). Some variation is still evident in the right-hand graph, which is probably due to inaccuracy in the system's `time()` function.

The top lines in each graph show the time required by the standard $O(n^2)$ algorithm, while the bottom lines show that of the `liboctrope` library. Notice that the right-hand graph shows that `liboctrope` is faster for trefoils of more than about 100 edges, and essentially all of the random walks tested.

On the right-hand graph, the standard algorithm's performance fits to $O(n^2)$, as expected. **Octrope's** performance, on the other hand, appears to fit $O(n^{4/3})$. Our clustering algorithm appears to be significantly more effective for random walks than trefoils. We conjecture that this is due to the fact that the walks occupy more volume than the trefoils, but we do not have a theoretical model for this effect.

7. CONCLUSIONS AND FUTURE DIRECTIONS

We have given an outline of an improved algorithm for computing the ropelength of polygonal space curves in time $O(n \log n)$ and contrasted it to the previous standard algorithm which required time $O(n^2)$. We have implemented the algorithm efficiently in ANSI C, and given timings which show that our algorithm is also much faster in practice than previous methods used in the field.

The increase in speed from using our method should enable researchers to consider significantly more complicated knots, and to get much higher-resolution data for simpler knots. Both of these are valuable goals. It has always been a goal of the geometric knot theory community to apply our results to large biomolecules such as DNA and proteins. Since these curves may involve thousands of vertices, they have been out of the reach of tools based on Algorithm 1.

However, our methods do not entirely settle the problem of fast ropelength computation. As mentioned in the introduction, Cantarella, Fu, Kusner, Sullivan, and Wrinkle [4] have discovered tiny straight segments in a ropelength-critical simple clasp. These segments are a few one-thousands of one unit in length out of a total clasp length of about 6 units (a similar clasp has been constructed by Starostin [18], though he does not prove it is critical). To resolve these very small scale phenomena numerically will require ropelength-minimized configurations with tens of thousands of edges.

At about 3 seconds per ropelength computation on a standard desktop machine, it would simply be untenable to minimize the ropelength of a 20,000-edge knot using our library and simulated annealing on a desktop machine. However, Octrope parallelizes well, so one could bring supercomputing cluster machines to bear on the problem, reducing the time to evaluate a configuration to tenths or hundredths of a second. This might allow for a long enough cooling schedule to resolve some small-scale phenomena, but there is no guarantee.

We are hence considering two further approaches to the problem: the **Multiresolution Ropelength Algorithm** and the use of Edelsbrunner’s “segment trees”([20]). The first algorithm is based on the idea that very high resolution knots can be well-approximated by subsampling the vertex set. If one keeps track of the distance between the subsampled knot and the original, one can again eliminate groups of edges from edge-edge checking. The potential advantages of this scheme are twofold: first, the construction of the corresponding tree is

linear in time and second, the subsampled knots can be handled by **Octrope** itself. The disadvantage of the multiresolution algorithm is that it will not help with random walks or other very complicated knots, such as large protein backbones, as their subsamples will not be close to the original curve.

We would also like to observe that while the discussion above is phrased in terms of polygons, the general octree/ramp method is equally applicable to other discretization schemes for curves, such as biarcs. In that case, the relative speed advantage of this algorithm should be greater, since the “edge-edge check” for a pair of arcs or spline segments is much slower than the edge-edge check for polygonal edges described above.

In conclusion, we hope that our algorithm and implementation will become a standard software component in numerical investigations of the ropelength problem. If others can improve our code, we hope that they will do so, and invite them to contact us. We also hope that our public release of the library (the first that we know of in geometric knot theory since Brakke’s Evolver [2]) will inspire others in the field to contribute from their personal and laboratory collections of code to the public domain.

Those interested in obtaining the code should browse to <http://ada.math.uga.edu/research/software/octrope/> for further information.

8. ACKNOWLEDGEMENTS

The authors are grateful to many colleagues for discussions about ropelength and algorithms, including Herbert Edelsbrunner, Mark Peletier, and John Sullivan. The 2002-2003 VIGRE group in Geometric Knot Theory (in particular Xander Faber, Chad Mullikin, and Nancy Wrinkle) contributed to our understanding of the computational issues surrounding ropelength. And Monica Shaw and Allison Diana, members of the 2003 Summer Undergraduate Research Experience, worked on a prototype implementation of the octrope algorithm. Michael Piatek and Eric Rawdon served as the liboctrope beta-test team, as well as contributing insight on efficient code and library design. The authors would also like to acknowledge the support of the National Science Foundation through the University of Georgia VIGRE grant (DMS-00-89927), DMS-99-02397 (to Cantarella), and DMS-02-04826 (to Cantarella and Fu).

-
- [1] Leonard M. Blumenthal and Karl Menger. *Studies in geometry*. W. H. Freeman and Co., San Francisco, Calif., 1970. (The three-point curvature appears on page 320.).
- [2] Kenneth A. Brakke. The surface evolver. *Experiment. Math.*, 1(2):141–165, 1992.
- [3] Gregory Buck and Jeremy Orloff. A simple energy function for knots. *Topology Appl.*, 61(3):205–214, 1995.
- [4] Jason Cantarella, Joseph H.G. Fu, Robert B. Kusner, John M. Sullivan, and Nancy Wrinkle. Criticality for the Gehring link problem. arXiv:math.DG/0402212, 2004.
- [5] Jason Cantarella, Robert B. Kusner, and John M. Sullivan. On the minimum ropelength of knots and links. *Invent. Math.*, 150(2):257–286, 2002.
- [6] Irene Gargantini. An effective way to represent Quadrees. *Commun. ACM*, 25(12):905–910, 1982.
- [7] O. Gonzalez and R. de la Llave. Existence of ideal knots. *J. Knot Theory Ramifications*, 12(1):123–133, 2003.
- [8] O. Gonzalez, J. H. Maddocks, F. Schuricht, and H. von der

- Mosel. Global curvature and self-contact of nonlinearly elastic curves and rods. *Calc. Var. Partial Differential Equations*, 14(1):29–68, 2002.
- [9] Oscar Gonzalez and John H. Maddocks. Global curvature, thickness, and the ideal shapes of knots. *Proc. Natl. Acad. Sci. USA*, 96(9):4769–4773 (electronic), 1999.
- [10] Chris L. Jackins and Steven L. Tanimoto. Oct-trees and their use in representing three-dimensional objects. *Comp. Graphics and Image Proc.*, 14:249–270, 1980.
- [11] Otto Krötenheerdt and Sigrid Veit. Zur Theorie massiver Knoten. *Wiss. Beitr. Martin-Luther-Univ. Halle-Wittenberg Reihe M Math.*, 7:61–74, 1976.
- [12] Alexander Nabutovsky. Non-recursive functions, knots “with thick ropes”, and self-clenching “thick” hyperspheres. *Comm. Pure Appl. Math.*, 48(4):381–428, 1995.
- [13] Piotr Pierański. In search of ideal knots. In *Ideal knots*, volume 19 of *Ser. Knots Everything*, pages 20–41. World Sci. Publishing, River Edge, NJ, 1998.
- [14] Eric Rawdon. TOROS: Thickness or Ropelength Optimizing System. Personal Communication.
- [15] Eric Rawdon. *The Thickness of Polygonal Knots*. PhD thesis, The University of Iowa, 1997.
- [16] Clifford A. Shaffer. Bit interleaving for Quad- or Octrees. In Andrew S. Glassner, editor, *Graphics Gems*, pages 443–447. Morgan Kaufman, 1990.
- [17] Jana Smutny and John Maddocks. Approximation of space curves with biarcs. 2004. Preprint.
- [18] Eugene L. Starostin. A constructive approach to modelling the tight shapes of some linked structures. *Proc. Appl. Math. Mech.*, 3:479–480, 2003.
- [19] Heiko von der Mosel and Friedemann Schuricht. Characterization of ideal knots. *Calculus of Variations*, 2003. Online. DOI:10.1007/s00526-003-0216-y.
- [20] A. Zomorodian and H. Edelsbrunner. Fast algorithms for box intersections. *Internat. J. Comput. Geom. Appl.*, 12:143–172, 2002.