Topological Neighborhoods for Spline Curves: Practice & Theory

Lance Edward Miller, Edward L.F. Moore, Thomas J. Peters, and Alexander Russell

Department of Computer Science & Engineering,
University of Connecticut,
Storrs, CT 06269-2155
tpeters@cse.uconn.edu

Abstract. The unresolved subtleties of floating point computations in geometric modeling become considerably more difficult in animations and scientific visualizations. Some emerging solutions based upon topological considerations for curves will be presented. A novel geometric seeding algorithm for Newton's method was used in experiments to determine feasible support for these visualization applications.

1 Computing the Pipe Surface Radius

Parametric curves have been shown to have a particular neighborhood whose boundary is non-self-intersecting [9]. It has also been shown that specified movements of the curve within this neighborhood preserve the topology of the curve [12, 13], as is desired in visualization. This neighborhood is defined by a single value, which is the radius of a pipe surface, where that radius depends on curvature and the minimum length over those line segments which are normal to the curve at both endpoints of the line segment [9]. Since computation of curvature is a well-treated problem, the focus of this paper is efficient and accurate floating point techniques to compute the other dependency for that radius.

Definition 1. *For a non-self-intersecting, parametric curve c, where*

$$c : [0, 1] \to \mathbb{R}^3,$$

and for distinct[1] *values* $s, t \in [0, 1]$*, then the line segment* $[c(s), c(t)]$ *is* doubly normal *if it is normal to c at both of the points* $c(s)$ *and* $c(t)$*.*

To avoid unnecessary complications with computing derivatives, only curves with regular parameterization [7] are considered.

Definition 2. *The* global separation *is the minimum over the lengths of all doubly normal segments. (For compact curves, this minimum has been shown in be positive [10].)*

[1] If the curve is closed, the s and t should be distinct values in $[0, 1)$.

P. Hertling et al. (Eds.): Real Number Algorithms, LNCS 5045, pp. 149–161, 2008.

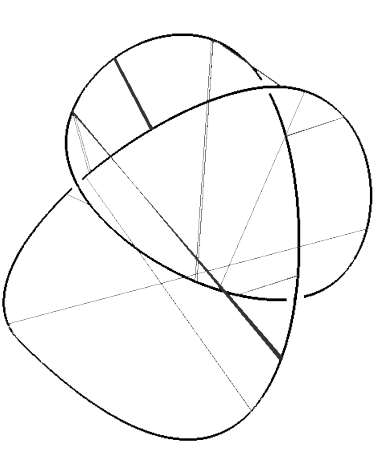

Fig. 1. Many doubly normal segments exist on this curve

An example cubic B-splines curve is given in Figure 1, with

1. control points: (0.0 0.0 0.0) (-1.0 1.0 0.0) (4.5 5.5 2.0) (5.0 -1.0 8.5) (-1.5 2.5 -4.5) (4.5 6.0 8.5) (3.5 -3.5 0.0) (0.0 0.0 0.0), and
2. knot vector: {0 0 0 0 0.2 0.4 0.6 0.8 1 1 1 1}

For this curve, there exist many doubly normal segments, as shown in Figure 1. The problem is how to efficiently find all these doubly normal segments, and then find the pair which represents the global separation distance, denoted as σ. A pair of distinct points at parametric values s and t on a curve will be endpoints of a doubly normal segment if they satisfy the two equations [9] for $s, t \in [0, 1]$:

$$[c(s) - c(t)] \cdot c'(s) = 0 \tag{1}$$

$$[c(s) - c(t)] \cdot c'(t) = 0. \tag{2}$$

In principle, the system given by Equations 1 and 2 could be solved algebraically by writing them in their power basis form, but this approach results in well-known algorithmic difficulties [14]. Hence, alternative techniques will be presented. For the software infrastructure available to these authors, it was convenient to convert the B-spline curve into a composite Bézier curve by the usual technique of increasing the multiplicity of each interior knot to 3. This produces 5 subcurves (depicted in Figure 2 by differing line fonts), with control points:

- (0 0 0) (-1 1 0) (1.75 3.25 1) (3.21 3.29 2.58),
- (3.20 3.29 2.58) (4.67 3.33 4.17) (4.83 1.17 6.33) (3.83 0.67 5.25),
- (3.83 0.67 5.25) (2.83 0.17 4.17) (0.67 1.33 -0.17) (0.58 2.5 -0.17),
- (0.58 2.5 -0.17) (0.5 3.67 -0.17) (2.5 4.83 4.17) (3.25 3.04 4.21),
- (3.25 3.04 4.21) (4 1.25 4.25) (3.5 -3.5 0) (0 0 0).

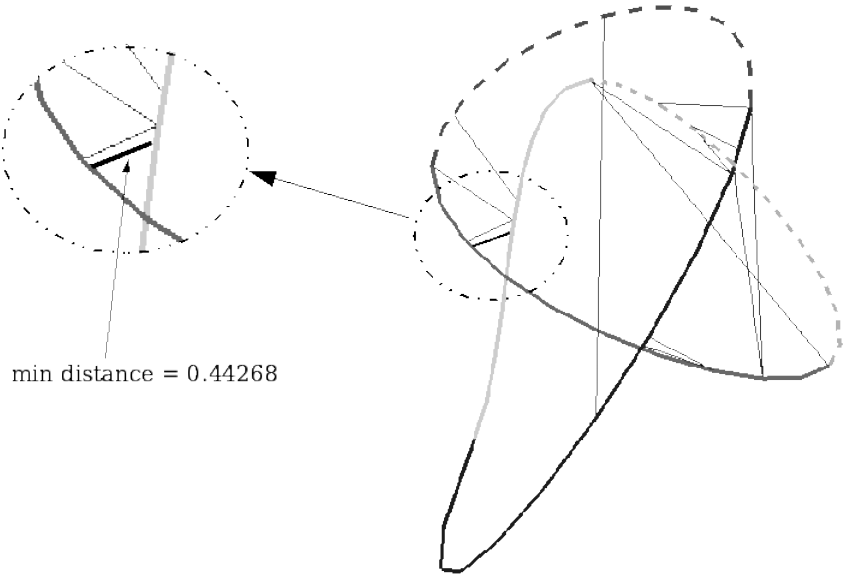

Fig. 2. Newton's method

Newton's method for two variables [11] was applied to Equations 1 and 2. The numerical experiments reported on prototype code suggest that this approach could be sufficiently rapid to support scientific visualization. These experiments were performed on a 64-bit AMD processor with Red Hat Linux Fedora Core 2 and OpenGL with double buffering. As always, the integration with a specific graphics subsystem is highly dependent upon the underlying architecture, and incorporation of this code on any platform would require further development and experimentation.

As is typical, the 'art' required for the successful use of Newton's method is highly dependent upon the determination of reasonable initial estimates, within the following standard formulation

$$\begin{bmatrix} s_{n+1} \\ t_{n+1} \end{bmatrix} = \begin{bmatrix} s_n \\ t_n \end{bmatrix} - J^{-1}(s_n, t_n) \begin{bmatrix} f(s_n, t_n) \\ g(s_n, t_n) \end{bmatrix}, n = 0, 1, ... \tag{3}$$

until $|J^{-1}(s_n, t_n)[f(s_n, t_n)\ g(s_n, t_n)]^T|$ is less than some $\epsilon > 0$, where $J^{-1}(s_n, t_n)$ is the inverse Jacobian matrix.

A viable approach to this art is presented and verified on an illustrative example. The general idea is to take finitely many points on each subcurve and consider all line segments between each pair of points as a candidate for being doubly normal. Many of these segments can be excluded from further consideration by an easy culling technique based upon a lack of normality at one end point or the other.

Let $\langle c(s), c(t) \rangle$ denote the vector of unit length, formed by taking the vector between $c(s)$ and $c(t)$ and dividing that vector by its norm. Let $c'(s)$ and $c'(t)$

denote the unit tangent vectors at points $c(s)$ and $c(t)$, respectively. Let ϵ_1 and ϵ_2 be positive. The following modifications of Equations 1 and 2 are used

$$\langle c(s), c(t) \rangle \cdot \ c'(s) < \epsilon_1, \tag{4}$$

and

$$\langle c(s), c(t) \rangle \cdot \ c'(t) < \epsilon_2. \tag{5}$$

If the result of the preceding comparisons[2] are false, then this segment is rejected. Otherwise, it is sufficiently close to being doubly normal to serve as an initial estimate for Newton's method.

These candidate double normal points are shown graphically in Figure 2 with line segments connecting the pairs of candidate points from Bézier segments. When Bézier segments are shown with no connecting line segment, that means that no candidate doubly normal points were found. When only one connecting line segment is depicted, that indicates that Newton's method did not converge for those particular points. When two pairs of connecting line segments are shown, that indicates that Newton's method did converge, and the resulting pair of minimum double normal points is one of the two line segments from each pair. Typically, convergence with $\epsilon = 0.0001$ occurred after 3 or 4 iterations, where similar behavior was corroborated in independently implemented code [2]. Note that Figure 2 depicts the same curve as in Figure 1, but now the curve is rotated about the y-axis to get a better view of doubly normal points, with σ illustrated in the zoomed-in section of Figure 2.

Table 1 summarizes experimental work completed. Tests 1 - 3 report on a naive, direct approach. This relies purely upon the limiting notion that sufficiently many approximation pairs will produce a list that contains a reasonable estimate for σ. This produces the reliable estimates shown in both Tests 1 and 2, but at prohibitively slow performance for visualization applications. Furthermore, Test 3 shows that further coarsening on the partition results in both poor estimates for σ *and* unacceptable performance. Alternatively, Test 4 shows that Newton's method produces a reliable estimate of σ with acceptable performance over a very coarse partition [3]. It should also be noted that the timing for the Newton's implementation is a very rough estimate and that the prototype code is not fully optimized, so further efficiencies could be expected. Even with these disclaimers, the time reported is encouraging for scientific visualization purposes.

2 Guaranteeing a Lower Bound

The estimate of σ produced by Newton's Method can be done quickly, but it could easily be an overestimate of σ. In this section we show how to efficiently

[2] The possibility of choosing different values for ϵ_1 and ϵ_2 is left as a user-option and is fully permissible within the theory presented. In practice, these values may often be chosen to be the same.

[3] As a verification of the Newton's code implemented, the value of σ for this experimental curve was corroborated by an independently created code [2].

Table 1. Estimating σ

Test #	Method	Partition Size, n	Time(s)	σ
1	Direct	10,000	85	0.44268
2	Direct	2,000	6	0.44268
3	Direct	1,000	2	0.91921
4	Newton	10	.0004	0.44268

determine a *guaranteed approximation* to the length of the shortest ϵ-nearly doubly normal line segment, a quantity we call $\sigma(\epsilon)$. (This is defined precisely below.) Note that $\sigma(\epsilon) \leq \sigma(0) = \sigma$, and that in order for this to be a guaranteed approximation of σ, one would have to establish a relationship between σ and $\sigma(\epsilon)$. However, if s is a good multiplicative approximation to $\sigma(\epsilon)$ in the sense that $\alpha^{-1} \leq s/\sigma(\epsilon) \leq \alpha$ (for some small $\alpha > 1$), then certainly $\alpha^{-1}s$ is a guaranteed lower bound on σ.

2.1 Partitioning by Taylor's Theorem

Given $\epsilon > 0$, the algorithm presented in this section depends on a subroutine PIPE(δ) that returns a PL approximation to c. Specifically, PIPE is called with a parameter δ and computes a PL approximation of c for which

- the Hausdorff distance between c and the PL approximation is bounded above by $\delta/2$, and
- the PL edges "ϵ-approximate" the tangents of c associated with this edge.

To do this for the curve c, the subroutine PIPE will determine a uniform partition of the parametric interval $[0, 1]$ by the increasing sequence of points

$$0 = s_0, s_1, \ldots, s_\ell = 1.$$

Then a PL approximation to the c is created by connecting the interpolant points

$$c(s_0), c(s_1), \ldots c(s_\ell).$$

Both conditions can be met by invoking Taylor's Theorem [6]. Taylor's Theorem is stated as follows. For $f : \mathbb{R} \to \mathbb{R}$ and $n > 0$, suppose that $f^{(n+1)}$ exists for each x in a non-empty open interval $I \subset \mathbb{R}$ containing a. For each $x \neq a$ in I, there exists t_x strictly between a and x such that

$$f(x) = f(a) + f'(a)(x-a) + \ldots + \frac{f^n(a)}{n!}(x-a)^n + r_n(x),$$

where

$$r_n(x) = \frac{f^{(n+1)}(t_x)}{(n+1)!}(x-a)^{n+1}.$$

Note that this statement of Taylor's Theorem is for the univariate case into $\mathbb{R}$, whereas the present application is to the map $c : [0, 1] \to \mathbb{R}^3$, a univariate function into $\mathbb{R}^3$. However, the x, y and z components can be treated independently as functions into $\mathbb{R}$.

Definition 3. *For any compact set* $K \subset [0,1]$ *and any continuous function* $f : K \to \mathbb{R}^3$*, and any* $t \in K$*, denote the components of* f *as* $f_x(t), f_y(t)$ *and* $f_z(t)$*. Then the* max *norm of* $f(t)$ *is denoted as* $\|f(t)\|_{max}$*, with*

$$\|f(t)\|_{max} = max\{f_x(t), f_y(t), f_z(t)\}.$$

Condition 1. PL Approximation within $\delta/2$**:** This part discusses the creation of a PL approximant of c that is within $\delta/2$ of c.

Since only C^2 functions defined on the compact set $[0,1]$ are considered, there is a maximum positive value for $\|c'(t)\|_{max}$, denoted as M_0. Recall that $c'(t)$ is non-zero. Then for any $t \in [t_0, t_1]$, (when $|t_1 - t_0|$ is sufficiently small), a straightforward application of Taylor's Theorem to the x component of $c(t)$, denoted as $c_x(t)$ would give,

$$c_x(t) = c_x(t_0) + E_x(t^*)$$

for some $t^* \in [t_0, t]$, where

$$E_x(t^*) = (t - t_0)c'_x(t^*),$$

with $E_x(t^*)$ playing the role of $r_1(x)$ above. Clearly, this can be done in each component. Then, since the final intent is to use the Euclidean norm on the vector-valued c, denoted as $\|c(t) - c(t_0)\|$, an elementary algebraic argument shows that the component-wise inequalities can be combined to yield

$$\|c(t) - c(t_0)\| \leq (t_1 - t_0)\sqrt{3}M_0.$$

Observe then that if $|s_{i+1} - s_i| \leq \delta/(2\sqrt{3}M_0)$ for each i, the curve c and this PL approximation are nowhere more than $\delta/2$ apart, as desired.

Note that this analysis only applies to a single curve, and recall that a curve c can be composed of many Bézier sub-curves. Suppose there are j many sub-curves. Then, the Taylor's theorem analysis must be applied to each of the j-many sub-curves.

Condition 2. Guaranteeing Good Local Tangent Approximations: This is analogous to the preceding argument. Suppose the curvature is positive somewhere. If not, the curve is the trivial case of a straight line. Let M_1 denote the maximum value of $\|c''(t)\|_{max}$, and let μ_0 denote the minimum value of $\|c'(t)\|_{max}$. A similar application of Taylor's Theorem yields,

$$\|c'(t) - c'(t_0)\| \leq |t_1 - t_0|\|c''(t^*)\| \leq (t_1 - t_0)\sqrt{3}M_1.$$

Let θ_t denote the angle between $c'(t_0)$ and $c'(t)$. Then,

$$|sin(\theta_t)| \leq \frac{\|c'(t) - c'(t_0)\|}{\|c'(t)\|}.$$

For a sufficiently small value of ϵ chosen to be greater than 0, the arcsine function is monotonically increasing on $[-\epsilon/4, \epsilon/4]$. Therefore, to show that $sin(\theta_t) < sin(\epsilon/4)$ over that interval, it is sufficient to have $|\theta_t| < \epsilon/4$, yielding

$$|sin(\theta_t)| \leq \frac{\|c'(t) - c'(t_0)\|}{\|c'(t)\|} \leq (t_1 - t_0)\frac{M_1}{\mu_0}.$$

Observe then that if $|s_{i+1} - s_i| \leq \sin(\epsilon/4)\mu_0/M_1$ for each i, the angular deviation along the curve will be bounded as desired.

The subroutine PIPE(δ), then, returns the PL approximation obtained by uniformly dividing the interval so that each

$$|s_{i+1} - s_i| \leq \min\left(\frac{\sin(\epsilon/4)\mu_0}{M_1}, \frac{\delta}{2\sqrt{3}M_0}\right) \quad .$$

2.2 Lower Bound for $\sigma(\epsilon)$

The introduction, here, of the terminology "ϵ-nearly doubly normal" is similar to the conditions previously set forth for the seeds for Newton's Method, as expressed in Equations 1 and 2 in Section 1.

Let $c(s_\sigma)$ and $c(t_\sigma)$ be two distinct points of c such that $d(c(s_\sigma), c(t_\sigma)) = \sigma$. Consider those circumstances, where for sufficiently small positive ϵ there exist $\tilde{s}_\sigma, \tilde{t}_\sigma \in [0, 1]$ such that the the normal planes P_1 and P_2 at $c(\tilde{s}_\sigma)$ and $c(\tilde{t}_\sigma)$, respectively, are distinct and intersect in a line near to the segment connecting $c(s_\sigma)$ and $c(t_\sigma)$ such that ν is a point on $P_1 \cap P_2$ which minimizes the sum $d(c(\tilde{s}_\sigma), \nu) + d(c(\tilde{t}_\sigma), \nu)$ and such that the angle ϕ formed between the segments connecting $c(\tilde{s}_\sigma)$ to ν and ν to $c(\tilde{t}_\sigma)$ is between $\pi - \epsilon$ and π. An illustration is shown in Figure 3, where $a = d(c(\tilde{s}_\sigma), \nu)$ and $b = d(c(\tilde{t}_\sigma), \nu)$ denote the lengths along the indicated line segments.

Any two points $c(s)$ and $c(t)$ are said to be ϵ-*nearly doubly normal* if

$$(c(s) - c(t)) \cdot c'(s) = 0 \quad \& \quad (c(s) - c(t)) \cdot c'(t) = 0,$$

or

$$\pi - \epsilon < \phi < \pi.$$

The triangle inequality gives $d(c(\tilde{s}_\sigma), c(\tilde{t}_\sigma)) \leq a + b$, and that $a + b \leq \sigma$. The algorithm described will estimate the global separation distance using approximations of $d(c(\tilde{s}_\sigma), c(\tilde{t}_\sigma))$. Since $d(c(\tilde{s}_\sigma), c(\tilde{t}_\sigma)) \leq \sigma$, the estimate produced, denoted as $\sigma(\epsilon)$ (defined immediately, below) will also be shown be no more than σ. The value $\sigma(\epsilon)$ (See Figure 4) is defined over any two ϵ-nearly normal points $c(t), c(s)$ with $t \neq s$,

$$\sigma(\epsilon) = min_{\{c(t), c(s)\}}\{d(c(t), c(s)).\}$$

The transition to providing an estimate of the more conservative value $\sigma(\epsilon)$ rather than trying to directly approximate σ is motivated by the following example. Let α be a planar C^∞ curve containing an arc of the unit circle with

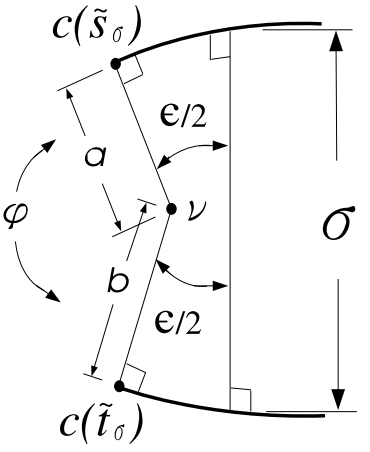

Fig. 3. The points $c(\tilde{s}_\sigma)$ and $c(\tilde{t}_\sigma)$ are ϵ-nearly doubly normal

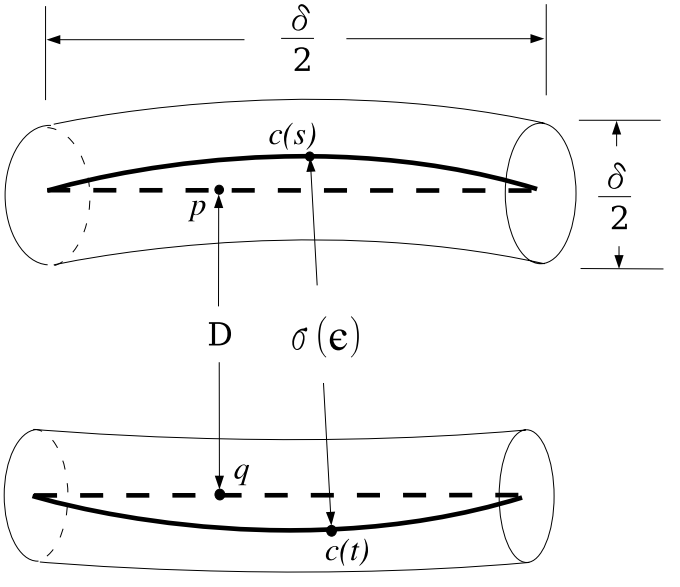

Fig. 4. The points $c(s)$ and $c(t)$ on the curve segments inside each cylinder are ϵ-nearly doubly normal, and D is the distance between the PL segments that approximates the curve segments

arc-length strictly less than π, but where α has its minimum separation distance being much greater than 2 and found elsewhere on the curve. For any algorithm that attempts to approximate σ by focusing upon pairs of points that were nearly normal within some fixed tolerance, there would always be

Global Separation Distance Estimate Algorithm

Input: A spline curve c & ϵ.
0. Initialize $\delta = 1$ and ω for upper precision bound.
1. Initialize $A(\epsilon) = 0$.
2. Apply PIPE(δ) to create a PL approximation.
3. Find pw-distances, $d(e_i, e_j)$
with points p, q that realize $d(e_i, e_j)$.
4. If p and q are ϵ-nearly doubly normal
retain $d(e_i, e_j)$ for further consideration,
Else discard.
Let $D = min(d(e_i, e_j))$ over the remaining points.
5. If $D \geq 4\delta$
$A(\epsilon) = D - 2\delta$
Else $\delta = \delta/2$, and Go to Step **1**.
Output: $A(\epsilon)$ = estimate for global separation distance.

Fig. 5. General algorithm for estimating the global separation distance

some input curve like α which would return some value near 2, since this arc-length can be made arbitrarily close to π.

The value $\sigma(\epsilon)$ is now accepted as a good estimate of σ, and the focus shifts to approximating $\sigma(\epsilon)$, recalling that $\sigma(\epsilon) \leq \sigma$. Then, the algorithm below in Figure 5 will return an approximation $A(\epsilon)$ of $\sigma(\epsilon)$, with the following two guarantees:

- $A(\epsilon) \leq \sigma(\epsilon) \leq \sigma$, and
- $A(\epsilon) > (\sigma(\epsilon))/2$.

Recall that the previous Taylor analyses guarantees that the result of PIPE(δ) satisfies the following three conditions:

- the length of each cylinder is strictly less than $\delta/2$,
- the radius of each cylinder is strictly less than $\delta/2$, and
- the angular deviation between tangents on the curve segments in each cylinder is strictly less than $\epsilon/4$.

The value for δ is initialized to 1. (It can be assumed that the curve has been normalized so that it lies in a sphere of radius 1 (Note that this also makes $\sigma(\epsilon) < 1$ for all $\epsilon > 0$, which is invoked later).) The resultant estimate $A(\epsilon)$ is then tested for validity (See algorithm in Figure 5), and failure results in halving the value of δ, repeating the iterations until a valid value is obtained. In this way, the overall algorithm is logarithmic in $1/\sigma(\epsilon)$.

2.3 Termination and Satisfactory Value

A termination, let $\check{D}$ be defined as the distance between two PL segments that approximate the curve segments in which $\sigma(\epsilon)$ is actually realized. Note that $D \leq \check{D} \leq \sigma(\epsilon) + \delta$, since the radius of the cylinders shown in Figure 4 is at most $\delta/2$, as given previously by Taylor's analysis.

The algorithm will terminate when $2\delta < \sigma(\epsilon)$. Several applications of the triangle inequality in Figure 4 show that $D \geq \sigma(\epsilon) - 2\delta$, or equivalently $D + 2\delta \geq \sigma(\epsilon)$, yielding

$$\frac{D}{\sigma(\epsilon)} \geq \frac{D}{D + 2\delta} = \frac{1}{1 + \dfrac{2\delta}{D}} \geq \frac{1}{1 + \dfrac{2\delta}{4\delta}} = 2/3.$$

Hence, $D \geq (2/3)\sigma(\epsilon)$ and $D \leq \sigma(\epsilon) \leq \sigma$, as desired.

The global separation distance algorithm in Figure 5 assumes the existence of a geometric distance predicate $d(e_i, e_j)$ between two line segments, e_i and e_j, which returns:

- the distance $d(e_i, e_j)$ between the two line segments, and
- the points p and q on e_i and e_j, respectively, where that distance is realized.

2.4 Asymptotic Time Bound

The time taken to approximate the global separation distance by this algorithm is quadratic in the bounds derived earlier for the Taylor's analysis. The final bound $\sigma(\epsilon)$ is computed within an *a priori* upper bound on the total number of subdivisions required as is standard practice [8]. As the algorithm is guaranteed to terminate when $\delta < \sigma(\epsilon)/2$ and δ is halved during each iteration, no more than $O(\log \sigma(\epsilon)^{-1})$ calls to PIPE are invoked. In the worst case, checking validity for a given PL approximation produced by PIPE takes quadratic time in the number of edges. The total time is thus no more than

$$O\left(\log(1/\sigma(\epsilon)) \max\{\ (\frac{M_0}{\delta})^2, \quad (\frac{M_1}{\mu_0})^2\ \}\right).$$

2.5 Example Analysis

For the curve already used, the values for the indicated parameters, above, were computed using the Maple computer algebra system as

- $M_0 = 14.9$,
- $\mu_0 = 3.4$,
- $M_1 = 21.9$.

Then an easy analysis shows that the number of subintervals generated for each sub-curve is 2048, which is consistent with the empirical findings in Table 1, where approximately 2000 sampled points per sub-curve produced an acceptable approximation. However, this algorithm for $\sigma(\epsilon)$ provides the additional

information that $\sigma(\epsilon)$ is a lower estimate and is truly close in the precisely defined sense given in Section 2.3. Of course, the input numerical parameters between the two algorithms would cause slight variances, but the agreement within this order of magnitude comparison is of interest.

2.6 Open Issues for Future Work

Within the Taylor's analysis performed, a well-defined lower estimate is established. Neither of the two algorithms presented (based upon Newton's Method or Taylor's Theorem) can effectively preclude output of a value that might be generated due purely to local properties of the curve. In practice, though, this is not quite as problematic as it may first appear. Recall that the purpose in estimating σ was to find the *global* factor that contributed to determining the radius of the neighborhood around the curve, while the *local* factor was in terms of curvature. A minimization is taken over those two factors to determine the radius. So, if either of the algorithms presented here returns a minimum value that is reflective of local properties, then this may suggest that the curvature is the determining factor for the radius. In those cases where curvature is the determining factor, then one need not even estimate σ, but these authors know of no *a priori* way to discriminate these cases, in order to avoid unnecessary computations. Resolving this issue remains beyond the scope of the present article but it is of interest for future investigation, both

- *experimentally*, with the algorithms discussed on more examples, and
- *theoretically*, by examination of adaptive skeletal structures, such as the medial axis [1], but also inclusive of more recent alternatives [3, 4, 5].

Extensions to higher dimensional geometric elements appear to be possible, but remain the subject of future work.

3 Experimental Observations

The curves here were assumed to be C^2. While this is sufficient for Newton's method, it is clearly *not* necessary to have the curve be C^2 *globally* . Clearly, Newton's method is local, so it it will be sufficient to have the C^2 condition *locally* within neighborhoods of the seeds. This is shown in Figure 6. This composite cubic Bézier curve has a point of non-differentiability at the top, where the three segments are shown in differing line fonts. Yet Newton's method easily and quickly estimates σ as 1.52, using only 10 partitioning points per sub-curve. This value of 1.52 was verified by the direct method discussed in Section 1. In Figure 6, two line segments are shown, with the thinner font indicating the seed and the thicker font denoting the converged value from Newton's method.

This composite Bézier curve has three segments and its control points are

- (0, 0.5, 0), (0.75, -1, 0), (0.83, -1.67, 0), (0.72, -2.11, 0),
- (0.72, -2.11, 0), (0.5, -3, 0), (-0.5, -3, 0),
- (-0.72, -2.11, 0), (0.83, -1.67, 0), (0.75, -1, 0), (0, 0.5, 0).

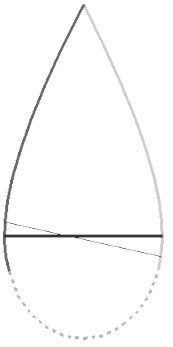

Fig. 6. A composite Bézier curve with a non-differentiable point

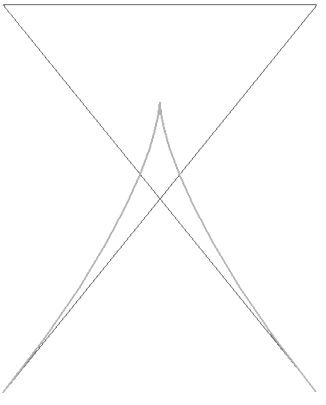

Fig. 7. A Bézier curve with a cusp

When the bound on the angular deviation of **Condition 2** of Subsection 2.1 is $\epsilon = 0.1$ (as done Subsection 2.5), the Taylor's analysis yields a value of $\ell = 10$, indicating $2^\ell = 1024$ partition points, far in excess of the 10 used here for Newton's method.

Of course, care must still be exercised in using Newton's method, as shown in Figure 7. Here there is a cusp at the top and the control polygon is shown. Using a very fine sampling relative to Inequalities 4 and 5, results in accepting a seed that is far into the cusp. Under Newton's method such a seed converges to an estimate of zero for σ. For the particular curve σ does equal zero, but the curve shown could be merely a subset of a much larger closed curve having a non-zero value for σ, meaning that this zero estimate would be inappropriate. Note that the algorithm for $\sigma(\epsilon)$ would specifically detect this difficulty by its check on the magnitude of the derivatives, thereby identifying this unbounded derivative and terminating the algorithm. Similar checks should also be incorporated into any practical code for Newton's method in this application. This

example provides further motivation for studying the trade-offs regarding local and global properties, as mentioned in Subsection 2.6.

4 Conclusion

Newton's method in two variables, when implemented with some novel geometric seeding techniques, provides an approach that is promising for preservation of topological characteristics during scientific visualization. Experiments and an alternative theoretical analysis, based upon Taylor's Theorem, are presented.

Acknowledgements. The authors were partially supported by NSF grants DMS-9985802, DMS-0138098, CCR-022654, CCR-0429477 and/or by an IBM Faculty Award. All statements here are the responsibility of the authors, *not* of the National Science Foundation *nor* of IBM. The authors thank the Dagstuhl Seminar organizers and the Dagstuhl staff for providing the intellectually stimulating environment for refinement of these ideas, which were initially based upon the dissertation of E. L. F. Moore.

References

1. Amenta, N., Peters, T.J., Russell, A.C.: Computational topology: ambient isotopic approximation of 2-manifolds. Theoretical Computer Science 305, 3–15 (2003)
2. Bisceglio, J.: Personal communication. justin.bisceglio@gmail.com (October 2005)
3. Damon, J.: On the smoothness and geometry of boundaries associated to skeletal structures, i: sufficient conditions for smoothness. Annales Inst. Fourier 53, 1941–1985 (2003)
4. J. Damon.: Determining the geometry of boundaries of objects from medial data (pre-print, 2004)
5. J. Damon.: Smoothness and geometry of boundaries associated to skeletal structures, II: geometry in the Blum case (pre-print, 2004)
6. Ellis, R., Gullick, D.: Calculus with Analytic Geometry, 3rd edn., Harcourt Brace Jovanovich (1986)
7. Farin, G.: Curves and Surfaces for Computer Aided Geometric Design: A Practicle Guide, 2nd edn. Academic Press, San Diego (1990)
8. Lutterkort, D., Peters, J.: Linear envelopes for uniform B–spline curves. In: Curves and Surfaces, St Malo, pp. 239–246 (2000)
9. Maekawa, T., Patrikalakis, N.M.: Shape Interrogation for Computer Aided Design and Manufacturing. Springer, New York (2002)
10. Maekawa, T., Patrikalakis, N.M., Sakkalis, T., Yu, G.: Analysis and applications of pipe surfaces. Computer Aided Geometric Design 15, 437–458 (1998)
11. Mathews, J.H.: Numerical Methods for Computer Science, Engineering and Mathematics. Prentice-Hall, Inc., Englewood Cliffs (1987)
12. Moore, E.L.F.: Computational Topology of Spline Curves for Geometric and Molecular Approximations. PhD thesis, The University of Connecticut (2006)
13. Moore, E.L.F., Peters, T.J., Roulier, J.A.: Preserving computatational topology by subdivision of quadratic and cubic Bézier curves (to appear)
14. Piegl, L., Tiller, W.: The NURBS Book, 2nd edn. Springer, Heidelberg (1997)