Improved Thin-Tube Models for Slender Vortex Simulations[1]

Omar M. Knio* and Rupert Klein†

*Department of Mechanical Engineering, The Johns Hopkins University, 3400 N. Charles Street, Baltimore, Maryland 21218-2686; and †FB Mathematik & Informatik, Freie Universität Berlin, Arnimallee 2-6, D-14195 Berlin, Germany

E-mail: *knio@jhu.edu, †rupert.klein@zib.de

This paper explores three numerical schemes for efficient simulation of slender vortex filaments. The schemes defeat the spatial and temporal stiffness of the equations of motion by requiring only adequate resolution of the filament centerline and allowing large integration time steps. In order to correctly capture the self-induced filament velocity, the first scheme uses an explicit velocity correction method, the second scheme relies on a logarithmic extrapolation of two velocity predictions, and the third scheme employs a local refinement algorithm. The performances of the three schemes are contrasted in light of unsteady computations of a perturbed vortex ring with small core to radius ratio.   © 2000 Academic Press

Key Words: slender vortices; asymptotic techniques; stiffness.

1. INTRODUCTION

Vortex element schemes are designed for application to flows with highly concentrated vorticity. These schemes are typically based on discretization of the vorticity field into spherically smoothed Lagrangian elements of overlapping cores and transport of these elements along particle trajectories. The advantages of this approach stem from the Lagrangian discretization, which naturally concentrates computational elements into regions of high vorticity, and from the Lagrangian transport procedure, which minimizes numerical diffusion.

One ad hoc ansatz for the application of vortex element schemes to the simulation of slender filaments is to represent the filament centerline using a "chain" of regularized vortex elements [1, 2]. When the cores of the vortex elements are overlapping, the smoothing procedure leads to a well-defined numerical vortex core structure, which is "typically"

---

[1] Dedicated to Professor Lu Ting on the occasion of his 75th birthday.

identified as the physical core vorticity distribution. However, Klein and Knio [3] show that in the resulting thin-tube model O(1) errors appear in the prediction of the vortex filament centerline motion. A detailed asymptotic analysis of the numerical and the physical vorticity structures in the vortex core reveals the origin of the discrepancy and naturally leads to a correction of the scheme. These ideas are extended in [4] (see also [5]) to include nontrivial vortex stretching as well as the effects of viscosity.

Application of thin-tube vortex element models to very thin vortices is difficult for two reasons. First, when the ratio δ/R of the core size to the characteristic radius of curvature of the filament centerline is small, high spatial resolution requirements arise. This is the case because, in the original approaches suggested in [3], the numerical core radius is essentially of the same order as the physical core radius. In addition, vortex elements are required to strongly overlap to ensure that the numerical core structure is well defined. Thus, when the slenderness ratio δ/R is on the order of 10^{-2} or smaller, the number of elements needed to satisfy overlap with cores of order δ may exceed by several orders of magnitude the number of elements needed for adequate resolution of the filament centerline. This would, consequently, necessitate prohibitively large CPU resources.

A second (related) difficulty is that the small spacing between elements may impose severe restriction on the integration step. These stiffness problems are well known in various Lagrangian calculations in which computational elements may tend to cluster. The temporal stiffness compounds the spatial resolution problem and substantially increases the demand for computational resources.

In this paper, we explore several means to overcome the spatial and temporal stiffness of slender vortex simulations. We focus on the corrected thin-tube models proposed in [3] and [4] and implement several approaches to enhance their performance. In order to describe these approaches, we start in the following section with a brief outline of the corrected thin-tube models. Section 3 then summarizes the various approaches used to address the stiffness issues. Section 4 provides a brief discussion regarding the selection of model parameters, based on static velocity predictions. Implementation of the optimized schemes is illustrated in Section 5 in light of unsteady 3D computations of a perturbed slender ring. We conclude in Section 6 with a brief summary.

2. CORRECTED THIN-TUBE MODELS

2.1. *Governing Equations for Filament Motion*

Corrected thin-tube models are Lagrangian vortex methods that "simulate" the asymptotic filament evolution equation [6]

$$\dot{\mathbf{X}}(s) = \frac{\Gamma}{4\pi} \left[\ln\left(\frac{2}{\delta}\right) + C \right] \kappa(s)\mathbf{b}(s) + \mathbf{Q}^f(\mathbf{X}(s)), \tag{1}$$

where Γ is the circulation of the filament, s is the arc length parameter along the time-dependent filament centerline $\mathcal{L}(t)$, κ and $\mathbf{b}$ are the curvature and unit binormal at $\mathbf{X}$, respectively, and $C(t)$ is the time-dependent core structure coefficient. $C(t)$ represents the contributions of the local swirling and axial velocities to the leading-order velocity of the filament; it is expressed as

$$C(t) = C_v(t) + C_w(t), \tag{2}$$

where

$$C_v(t) = \lim_{\bar{r}\to\infty} \left\{ \frac{4\pi^2}{\Gamma^2} \int_0^{\bar{r}} \bar{r} v^{(0)^2} \, d\bar{r} - \ln\bar{r} \right\} - \frac{1}{2} \tag{3}$$

$$C_w(t) = -\frac{1}{2}\left(\frac{4\pi}{\Gamma}\right)^2 \int_0^{\infty} \bar{r} w^{(0)^2} \, d\bar{r}. \tag{4}$$

$v^{(0)}$ and $w^{(0)}$ are the leading-order swirling and axial velocities within the filament core, respectively, and $\bar{r} \equiv r/\delta$. Meanwhile, $\mathbf{Q}^f$ is the so-called finite part of the line Biot–Savart integral [3, 6]; it represents the nonlocal self-induction of the filament.

The core structure coefficients evolve in time according to the evolution of the leading-order axial vorticity and axial velocity distributions. As shown by Callegari and Ting [6], the leading-order axial vorticity and velocity within the core obey inhomogeneous heat equations with a source term that depends on the stretching of the filament centerline. In the inviscid limit, simple closed-form expressions for the evolution of C_v and C_w have been obtained in [7],

$$C_v(t) = C_v(0) + \ln\sqrt{\frac{S(t)}{S(0)}} \tag{5}$$

$$C_w(t) = \left[\frac{S(0)}{S(t)}\right]^3 C_w(0), \tag{6}$$

where $S(t)$ is the instantaneous total arc length of the filament. When viscous effects are present, the expressions describing the evolution of the core structure coefficient are more involved [4, 5, 7]; for brevity, they are omitted.

2.2. Numerical Simulation

Construction of thin-tube models is based on discretization of the centerline $\mathcal{L}$ into a finite number of regularized vortex elements with spherical overlapping cores. The vortex elements are described in terms of their Lagrangian position vectors, $\chi_i(t)$, $i, \ldots, N$, which are indexed consecutively such that the resulting "chain" describes the filament centerline [2, 3]. Based on the Lagrangian variables, a smooth representation of the *regularized* filament vorticity is obtained using the expression [8]

$$\omega(x, t) = \sum_{i=1}^{N} \Gamma \delta\chi_i(t) f_\sigma(x - \chi_i(t)), \tag{7}$$

where $f_\sigma(x) \equiv \sigma^{-3} f(|x|/\sigma)$, f is a rapidly decaying spherical core function of unit mass, $\delta\chi_i(t)$ is the arc length increment associated with the ith element, and σ is the *numerical* core radius. When inserted into the three-dimensional Biot–Savart integral, the above representation yields the desingularized velocity field

$$\mathbf{v}^{\text{ttm}}(x, t) = -\frac{\Gamma}{4\pi} \sum_{i=1}^{N} \frac{(x - \chi_i(t)) \times \delta\chi_i(t)}{|x - \chi_i(t)|^3} \kappa_\sigma(x - \chi_i(t)), \tag{8}$$

where $\kappa_\sigma(x) \equiv \kappa(|x|/\sigma)$, and $\kappa(r) = 4\pi \int_0^r \xi^2 f(\xi)\,d\xi$ is the velocity smoothing kernel corresponding to f [8]. In the computations, the arc length increments $\delta\chi_i(t)$ are related to the distribution of particle positions using the procedure in [3]. It is based on a Lagrangian spectral collocation interpolation of the filament geometry onto the particle positions and approximating the arc length based on spectral collocation derivatives of the interpolated filament centerline.

The above discretization and regularization procedure endows the thin tube with a numerical core structure that depends on the choice of the core size and the smoothing function. Unfortunately, if the numerical core vorticity distribution is identified with the physical core vorticity distribution, and O(1) velocity error in the predicted centerline velocity occurs.

To obtain a velocity prediction that is consistent with Eq. (1) and the definitions of C (Eq. (2)) and $\mathbf{Q}^f$ [6], Klein and Knio [3] propose three correction procedures for the above thin-tube model. We restrict our attention to two of these approaches, which will be used later in the optimization of the numerical scheme.

In the first approach, the numerical core radius σ is taken to be equal to the physical core size δ, and the velocity correction

$$\delta\mathbf{v} \equiv \frac{\Gamma}{4\pi}(C - C^{\text{ttm}})\kappa\mathbf{b} \tag{9}$$

is added to the thin-tube prediction $\mathbf{v}^{\text{ttm}}$ to absorb the discrepancy in the centerline velocity. Here, C^{ttm} is the numerical core constant which corresponds to the choice of core smoothing function [3].

In the second approach, the numerical core radius σ is selected according to

$$\sigma = \delta \exp(C^{\text{ttm}} - C). \tag{10}$$

This ensures that the corresponding regularized velocity field $\mathbf{v}^{\text{ttm}}$ (8) coincides with the theoretical prediction in (1) at the particle positions. An advantage of the present approach is that explicit evaluation of the curvature and binormal at the filament centerline is avoided.

Once the corrected centerline velocity, $\mathbf{v}^{\text{cor}}$, is determined it is used to update the positions of the vortex elements. This is achieved by integrating

$$\frac{\partial\chi_i}{\partial t} = \mathbf{v}^{\text{cor}}(\chi_i(t), t). \tag{11}$$

In the computations, a third-order Adams–Bashforth scheme is used to numerically integrate the above system.

3. OPTIMIZATION TECHNIQUES

In this section, we outline three methods for optimizing the corrected thin-tube model outlined in the previous section. In all cases, our approach is motivated by the observation that (1) in the previous construction the number of elements required to achieve overlap is much larger than that needed to adequately describe the centerline geometry and (2) consequently, both the temporal stiffness and high spatial resolution problems can be effectively addressed if a coarser (or generally more efficient) discretization can be adopted. This observation suggests a number of different optimization approaches, of which we discuss the following three.

3.1. *Method 1: Extrapolation Technique*

This approach is based on selecting a discretization level that is fine enough for the representation of the filament centerline irrespective of the filament core size δ. As previously discussed, selection of numerical core size levels using either of the approaches of the previous section becomes problematic since overlap among neighboring elements may not be satisfied. In fact, for small slenderness ratios, overlap is likely to be everywhere violated.

To avoid this issue, a large value of the numerical core size is used. Let $\tilde{\sigma}$ denote the core size parameter predicted by Eq. (10), let

$$\sigma_o(t) = \max_{i=1,\ldots,N} |\delta\chi_i(t)| \tag{12}$$

be the maximum inter-element separation distance, let $\sigma_1 = K\sigma_o$, and let $\sigma_2 = \alpha\sigma_1$. Here K and α are real constants. Note that for any choice $K > 1$ and $\alpha > 1$ implementation of a thin-tube model using either σ_1 or σ_2 will automatically satisfy the overlap condition.

Next, we recall that for any σ the velocity field predicted by the uncorrected thin-tube model corresponds to [3]

$$\mathbf{v}^{\text{ttm}} = \frac{\Gamma}{4\pi}\left[\log\left(\frac{2}{\sigma}\right) + C^{\text{ttm}}\right]\kappa\mathbf{b} + \mathbf{Q}^f \tag{13}$$

as long as strong overlap is satisfied. Consequently, if we let $\mathbf{v}_1$ and $\mathbf{v}_2$ denote the thin-tube velocity prediction based on σ_1 and σ_2, respectively, then the *corrected* velocity field prediction can be estimated based on the following logarithmic extrapolation procedure:

$$\mathbf{v}^{\text{cor}} = \mathbf{v}_1 + (\mathbf{v}_1 - \mathbf{v}_2)\frac{\log\left(\frac{\sigma_1}{\tilde{\sigma}}\right)}{\log\alpha}. \tag{14}$$

In other words, the corrected centerline velocity can be obtained at the cost of two "coarse mesh-fat core" evaluations. Note that the present procedure does not require estimates of the filament curvature and binormal.

3.2. *Method 2: Explicit Velocity Correction*

The velocity correction technique is based on a single velocity evaluation using large σ and implementing an explicit correction to the corresponding velocity prediction. This approach has been first implemented in Zhou's study of Kelvin waves on a slender vortex [9]. Here, we implement a variant of the velocity correction approach which is summarized by

$$\mathbf{v}^{\text{cor}} = \mathbf{v}_1 + \frac{\Gamma}{4\pi}\log\left(\frac{\sigma_1}{\tilde{\sigma}}\right)\kappa\mathbf{b}, \tag{15}$$

where σ_1 is dynamically determined during the simulation using the same definitions discussed for method 1. In the computations, the curvature and binormal are obtained following the procedure given in [3].

3.3. *Method 3: Local Mesh Refinement*

This approach is motivated by the observation that the numerical core structure along a particular location on the axis of the thin tube is actually induced by neighboring elements only. Since the core smoothing function decays rapidly, the numerical *vorticity* structure at

a given point is determined by the fields of elements lying within a few numerical core radii of that point. This suggests the following approach for estimating the corrected thin-tube velocity.

As in the previous two methods, we rely on a discretization level that is fine enough to adequately represent the filament geometry. However, unlike the previous two approaches, the rescaled numerical core size predicted by Eq. (10) is used, so that overlap is not globally satisfied. Based on this choice of parameters, a modified velocity evaluation procedure is implemented. To evaluate the velocity field at a given point, the elements are divided into two disjoint groups: a group of neighboring elements and a group of well-separated elements. A separation distance of three core radii is used as basis for this segregation procedure. The contribution of elements belonging to the second group is computed directly on the coarse grid. Meanwhile, the contribution of neighboring elements is accounted for by locally remeshing the corresponding segments using a fine-enough grid for core overlap to be locally satisfied.

In the present computations, a simplified version of the local remeshing procedure is implemented. It consists of determining, at every time step and in a global fashion, a Lagrangian grid that is fine enough for overlap to be everywhere satisfied. The fine grid is determined by Fourier transformation of the filament geometry, extending the resulting spectrum by array padding, and then inverse transforming the extended spectrum onto physical space.

Remark. The improved schemes outlined above address the stiffness of the slender filament equations using different strategies. From a practical standpoint, method 1 (M1) is very attractive as it does not require any modification of a corrected thin-tube code [3, 4]. Specifically, all the code elements that determine the element velocity based on Eq. (8) remain unchanged, and the optimization operates on their output only. The attractive feature of method 2 (M2) is that it isolates the singular part of the element evolution equations and, as outlined in [9], enables the optimization of the numerical integration using operator-splitting approaches. However, a disadvantage of M2 is that it requires explicit evaluation of higher order derivatives of the filament centerline. This has adverse effects on the numerical stability of the scheme, as observed in [3] and discussed in Section 5.2 below. Method 3 (M3) would be quite attractive in the context of fast vortex element methods (e.g., [10, 11]), which use similar clustering ideas to defeat the $O(N^2)$ cost of direct vortex interactions. The considerable bookkeeping and storage management associated with these methods compared to the direct methods would then lose its importance. A key disadvantage of M3, however, concerns the stiffness issue. As the scheme operates with the physical core size and evaluates velocities by directly accounting for neighboring (overlapping) computational elements, there will be strong cancellation of large kernels as the core size diminishes. The phenomenon is due to the fact that the local induced velocities are by a factor $O(1/\delta)$ larger than the net filament velocity [3, 6]; it leads to amplification of roundoff errors and compounds the stiffness problem.

It follows from the present discussion that, from a practical standpoint, M1 appears as the most attractive alternative. This claim is further supported in the computational tests below.

4. PARAMETER SELECTION

The optimization methods introduced above involve various numerical parameters that are related to the discretization of the slender filament and to the choice of numerical core

size(s). Specifically, for M1 and M2 one needs to specify the number of elements, N, and the overlap parameter, K. For M1, one must in addition specify a second parameter, α. For M3, core coarsening parameters are not needed but fine and coarse discretization levels are used instead. In this section, we explore general guidelines for the selection of these parameters.

To this end, we start by emphasizing that the constructions of the previous two sections are based on the key assumption that the predictions of the thin-tube model agrees with the asymptotic expression. For this to hold, two conditions must be satisfied [3]: (a) the selected (coarsened) numerical core size must still be small enough for the corresponding prediction to fall within the range of validity of the asymptotic theory, and (b) the numerical Biot–Savart integral must be accurately evaluated. As discussed below, one can use these two conditions to develop guidelines for the selection of optimization parameters.

To illustrate these guidelines, we consider the case of a circular vortex ring. Variables are normalized so that the dimensionless ring radius $R = 1$, and the dimensionless circulation $\Gamma = 1$. The core smoothing function $f(r) = \mathrm{sech}^2(r^3)$ is used. The corresponding velocity kernel is $\kappa(r) = \tanh(r^3)$ [8], and the numerical core structure coefficient $C^{\mathrm{ttm}} = -0.4202$ [3]. For brevity, we focus exclusively on static velocity predictions; unsteady computations are presented in the following section.

We start by addressing condition (b) above, and so we examine the role of the parameter K that enters in M1 and M2. As discussed earlier, K measures the ratio between the maximum element length and the numerical core size, and consequently it measures the degree of overlap among neighboring cores. Thus, K is referred to as an "overlap" parameter. This aspect has been extensively discussed in theoretical convergence studies of particle methods (e.g., [8]) as well as computational studies (e.g., [2]). Based on previous experiences, one would expect that the numerical evaluation of the Biot–Savart integral becomes valid when neighboring cores overlap (i.e., in the range $K > 1$). In Fig. 1, we examine the effect of K by applying M2 to compute the self-induced velocity of a slender vortex ring with $\tilde{\sigma} = 0.001$. The results indicate that for $K > 1.5$, the predictions become essentially independent of the selected value of K. In all the tests below, we conservatively choose overlap parameters in the range $K \geq 2$.


FIG. 1. Effect of overlap parameter K on the velocity prediction. Results are obtained using M2 with $N = 512$. The core radius $\tilde{\sigma} = 0.001$.

FIG. 2. Dependence of the discrete estimate of the numerical core constant, C^{ttm}, from Eq. (16) on the ring slenderness ratio $\epsilon = \tilde{\sigma}/R$. The ring is discretized using $N = 2048$ elements.

We now turn our attention to the first condition, which concerns the slenderness ratio of the filament. Note that both M1 and M2 rely on numerical Biot–Savart integrals with enlarged cores, and one must consequently ensure that the corresponding slenderness ratio $\epsilon \equiv \kappa_{\max}\sigma$ is small enough to fall within the regime of validity of the theory. To explore the regime of validity of the constructions, we rely on the computed self-induced velocity of the circular ring to generate a discrete estimate of the numerical core constant; we use

$$C^{\text{ttm}} \simeq \frac{4\pi V}{R} - \log\left(\frac{8R}{\tilde{\sigma}}\right), \tag{16}$$

where V is the computed self-induced velocity and $\tilde{\sigma}$ is the numerical core radius. In Fig. 2, we plot the values of C^{ttm} computed using Eq. (16) against the ring slenderness ratio $\epsilon = \tilde{\sigma}/R$. As can be observed in Fig. 2, for $\epsilon < 0.1$ the computed core constant C^{ttm} is weakly dependent on the slenderness ratio. It is also interesting to note that as ϵ decreases the predicted values tend toward the theoretical value [3].

The above exercise can be used, in the application of the optimization techniques, to guide the selection of appropriate grid coarsening levels, or alternatively of the number of elements. For instance, for a filament with characteristic peak curvature $\kappa_{\max}$ the condition $\epsilon < 0.1$ can be used to estimate a maximum allowable element length, $\delta\ell_{\max}$. A reasonable estimate for M2 is $\delta\ell_{\max} \simeq 0.1/(K\kappa_{\max})$; for M1, we set $\delta\ell_{\max} \simeq 0.1/(K\alpha\kappa_{\max})$. Assuming that the overall filament arc length is S and that the lengths of vortex elements are fairly uniform, then the number of elements may be estimated using $N \simeq S/\delta\ell_{\max}$. In addition to relying on this initial estimate, the computations of the following section continuously monitor ϵ in order to verify that it remains within acceptable levels. Our experiences indicate that the results of the optimized methods are valid when ϵ remains small; in contrast, an example is provided in Section 5.2 which shows that the predictions deteriorate for $\epsilon > 0.1$.

We conclude this section with a brief comment regarding the parameter α used in M1. Following the discussion above, when K and N have been optimized, α must naturally be restricted to the range $\alpha > 1$. In order not to penalize the discretization level, it would obviously be advantageous to choose α as close to unity as possible. At the same time, one should guard against the possible amplification of extrapolation errors. To examine this

FIG. 3. Effect of α on the computed self-induced velocity of a circular vortex ring with $\tilde{\sigma} = 0.001$. Results are obtained using M1 with $K = 2$ and $N = 2048$.

issue, we use M1 to compute the self-induced velocity of a slender circular vortex ring; the results are plotted in Fig. 3 against the corresponding values of α. As shown in the figure, for $\alpha \geq 1.25$ the predictions become essentially insensitive to the selected value of α. This restriction is imposed in all the calculations below.

5. RESULTS AND DISCUSSION

Implementation of the optimized schemes outlined above is described in light of inviscid 3D simulations of a slender vortex ring. We use the same normalization convention and the same core smoothing function as in the previous section.

In order to observe a nontrivial slender vortex evolution the ring centerline is perturbed using plane, azimuthal bending waves. Initially, the perturbed ring radius is given by

$$\rho(\theta) = R[1 + a_1 \sin(k_1\theta) + a_2 \sin(k_2\theta)], \tag{17}$$

where θ is the azimuthal angle, k_1 and k_2 are integer wavenumbers, and a_1 and a_2 are the corresponding amplitudes.

5.1. *Performance Measures*

The evolution of the slender ring described above is computed using the rescaled numerical core radius approach [3], which is referred to as original scheme, as well as the optimized schemes based on methods 1, 2, and 3 above, which are labeled M1, M2, and M3, respectively. Results are obtained for a slender vortex ring with $\tilde{\sigma} = 0.02$, $k_1 = 2$, $k_2 = 3$, and $a_1 = a_2 = 0.02$. The parameters used in the runs are summarized in Table I. Note that core overlap using $\tilde{\sigma}$ would not be satisfied in M1, M2, and M3 at the selected coarse-grid resolutions.

The predictions of the original and optimized schemes are first illustrated in Fig. 4, which shows the evolution of the peak centerline curvature and velocity. The results show that the

TABLE I
Numerical Parameters and Performance Measures

	Original	M1	M2	M3
σ	$\tilde{\sigma}$	(σ_1, σ_2)	σ_1	$\tilde{\sigma}$
N	1024	256	256	256/1024
K	—	2	3	—
α	—	1.5	—	—
CPU time[a] (s)	708.9	68.19	45.98	71.37
Performance gain	1	10.39	15.42	9.93

[a] SGI R10000 195-MHz processor; $\Delta t = 0.002$; 1500 iterations.

peak curvature of the filament undergoes large-amplitude periodic oscillations and that the behavior of the peak velocity follows closely that of the peak curvature. The figures show that the predictions of the optimized schemes are in close agreement with each other and with the predictions of the original scheme.

The agreement between the curves in Fig. 4 is quantified by computing the relative deviation between the predictions of the improved schemes and the results of the original method. For instance, the relative deviation in peak curvature for method M_i is defined as

$$D_i = \frac{\max_{[0,t]} \left| \kappa_{\max}^{\mathrm{orig}}(t) - \kappa_{\max}^{M_i(t)} \right|}{\max_{[0,t]} \kappa_{\max}^{\mathrm{orig}}(t)}. \tag{18}$$

A similar definition based on the peak velocity is used. Based on these definitions, the computed relative deviation in peak curvature is 0.159% for M1, 0.118% for M2, and 0.194% for M3. Meanwhile, the relative deviations in peak velocity are 0.109% for M1, 0.064% for M2, and 0.234% for M3.

Very close agreement between the predictions of the original and optimized schemes is also evident in Fig. 5, which depicts the spatial distribution of centerline curvature along the circumference of the ring. As observed earlier, the results of the optimized schemes are nearly identical to those obtained using the original model.


FIG. 4. Evolution of the maximum curvature (left) and maximum velocity (right) for a perturbed vortex ring with $\tilde{\sigma} = 0.02$. Plotted are results obtained using the original method and schemes M1, M2, and M3. Parameters are indicated in Table I.

FIG. 5. Spatial distribution of the curvature along the filament axis at $t = 3$. Curves are generated using the results of the original method and schemes M1, M2, and M3. Parameters are indicated in Table I.

The performance of the original and optimized schemes is quantified in Table I. The table provides the total CPU time spent and the performance gain which is defined as the ratio of the current CPU time to the CPU time spent in the original scheme. The results show that the optimized schemes produce an order-of-magnitude enhancement in the performance of the computations. This is a substantial performance gain, especially since a relatively large core-to-radius ratio is selected.

5.2. *Temporal Stiffness*

The numerical construction outlined in the previous section suggests that by only requiring adequate representation of the filament centerline the improved schemes may also be effective in defeating the temporal stiffness of the slender filament equations of motion. In the present section, we briefly explore this possibility in light of a simplified heuristic analysis of the critical time step, Δt_{crit}. Specifically, the analysis is based on performing unsteady calculations with different values of the time step, Δt, and monitoring the value at which the simulations become numerically unstable. To interpret the observations, simulations are performed with different values of the core size; results are obtained with $\tilde{\sigma} = 0.02$, 0.002, and 0.0002. Again, we restrict our attention to the initial centerline geometry of the previous section and to improved schemes M1 and M2 with $N = 256$.

In the examples of Section 5.1, the optimization parameters were selected following the guidelines developed in Section 4. In particular, the criterion $\epsilon < 0.1$ was satisfied throughout the computations. In the present implementations, the parameters for M1 are deliberately chosen so that this criterion is violated; we use $K = 3$ and $\alpha = 2$. This enables us to observe the effects of a poor choice in optimization parameters. For M2, we still use $K = 3$.

The results of the analysis are summarized in Table II, which provides the observed value of Δt_{crit} for different values of $\tilde{\sigma}$, and in Fig. 6, which depicts the evolution of the peak centerline curvature for $\tilde{\sigma} = 0.002$ and 0.0002. Table II also shows the dominant period of the oscillations in the filament centerline, which is deduced from the curvature evolution curves in Figs. 4 and 6.

TABLE II
Critical Time Step and Period

$\tilde{\sigma}$	Δt_{crit}: M1	Δt_{crit}: M2	Period
0.02	4.43×10^{-2}	1.35×10^{-2}	1.420
0.002	1.80×10^{-2}	4.17×10^{-3}	0.836
0.0002	1.10×10^{-2}	2.56×10^{-3}	0.586

The results show that for both M1 and M2 Δt_{crit} decreases slowly as $\tilde{\sigma}$ is reduced by two orders of magnitude. The behavior in Δt_{crit} appears to be consistent with the reduction in the observed period of the centerline oscillations and with the logarithmic nature of the filament self-induction law. This indicates that at a given *filament* resolution level the critical time step is restricted by the *centerline* evolution rate. Since the optimized schemes require adequate resolution of the filament centerline only, the present tests show that they are also effective in overcoming the stiffness of the equations of motion.

It is also interesting to note that Δt_{crit} for scheme M1 is consistently larger than that of M2. For the conditions of the present simulations, the results in Table II show that at the critical time step there are about 30–50 iterations per period for scheme M1 and about 100–200 iterations per period for scheme M2. This behavior can be attributed to the fact that scheme M1 avoids the evaluation of higher order derivatives [3, 12] which enter in the expressions of curvature and binormal. Combined with the performance measures of the previous section, the results in Table II indicate that if the differences in Δt_{crit} were accounted for then scheme M1 would outperform M2.

One should also note that for $\tilde{\sigma} = 0.0002$, the agreement between the predictions for M1 and M2 is evidently degraded. The deviation between the two curves in Fig. 6 is 3.24%, and the corresponding deviation based on peak velocity is 2.37%. In order to examine the origin of the deviation, we note that for M2 with $N = 256$ and $K = 3$, the peak slenderness ratio based on σ_1 is $\epsilon_1 \equiv \max_{[0,t]} \kappa_{\max}(t)\sigma_1(t) = 0.0926$. For M1 with $N = 256$, $K = 3$, and $\alpha = 2$, the peak slenderness ratio based on σ_1 is $\epsilon_1 = 0.0926$; based on σ_2, we have $\epsilon_2 = 0.1853$. As discussed in Section 4, slenderness ratios above 0.1 are beyond the range of validity of the slender vortex theory and the predictions of the optimized schemes may


FIG. 6. Evolution of the maximum curvature for a perturbed vortex ring with $\tilde{\sigma} = 0.002$ (left) and $\tilde{\sigma} = 0.0002$ (right). Plotted are results obtained using schemes M1 and M2.

consequently suffer from significant modeling errors. Thus, the deviations observed in Fig. 6 at $\tilde{\sigma} = 0.0002$ appear to be due to the poor choice of K and α in M1. Note that at $\tilde{\sigma} = 0.002$ the deviations between M1 and M2 are harder to detect, which also suggests that large-ϵ effects become more pronounced as one extrapolates across a larger core size range.

To verify the above claim, computations were repeated using M1 with $K = 2$ and $\alpha = 1.5$. For these values of K and α, the computed values of ϵ_1 and ϵ_2 are 0.0617 and 0.0926, respectively. Consistent with the experiences of Section 4, the corresponding results (not shown) are again in close agreement with those of M2, with relative deviations of 0.978% for peak curvature and 0.912% for peak velocity. This enables us to emphasize that the predictions of the improved schemes and original method are consistent with each other and independent of the choice of numerical parameters, as long as these parameters obey the restrictions of the numerical construction and of the underlying slender vortex theory.

5.3. *Tests with Large Deformation*

In the numerical examples of Sections 5.1 and 5.2, the perturbation initially imposed had a small amplitude and the deformations of the filament centerline remained moderate. Specifically, spatial variations of the filament curvature were within $\pm 20\%$ of the mean, approximately.

This section provides a numerical example which illustrates the performance of the optimized schemes in a regime with $O(1)$ changes in the filament curvature. To this end, we consider a slender ring with $\tilde{\sigma} = 0.002$, initially perturbed using $k_1 = 5$, $k_2 = 6$, and $a_1 = a_2 = 0.03$. Unsteady computations are performed using M1 and M2. In order to observe the guidelines of Section 4, for M1 we use $K = 2$, $\alpha = 1.25$, and $N = 1024$; for M2, we use $K = 2.5$ and $N = 1024$.

Results of the simulations are shown in Fig. 7, which shows the evolution of peak curvature and the spatial distribution of curvature along the filament arc length at the end of the computations. The results clearly show that the peak curvature changes substantially with time and that significant differences in the curvature occur along the filament centerline. In addition, Fig. 7 shows that the predictions of M1 and M2 remain in close agreement


FIG. 7.   Left: Evolution of the maximum curvature of a perturbed vortex ring with $\tilde{\sigma} = 0.002$, $a_1 = a_2 = 0.03$, $k_1 = 5$, and $k_2 = 6$. Right: Spatial distribution of the curvature along the filament axis at $t = 0.75$. Plotted are results obtained using M1 with $K = 2$, $\alpha = 1.25$, and $N = 1024$ and M2 with $K = 2.5$ and $N = 1024$.

with each other throughout the computations. This agreement provides strong support to the validity of the optimized constructions.

6. CONCLUSIONS

The construction of three improved schemes for the simulation of slender vortex filaments is discussed. The schemes are based on the discretization of the filament centerline into desingularized Lagrangian elements and transport of these elements using the centerline velocity. In the first scheme, the centerline velocity is obtained via logarithmic extrapolation of two thin-tube velocity estimates with artificially large cores. In the second scheme, the centerline velocity is obtained by combining one thin-tube velocity prediction with a local correction formula. Meanwhile, the third scheme relies on a combination of local/distant thin-tube predictions evaluated on fine/coarse resolution levels. The performance of the improved schemes is analyzed in light of unsteady calculations of a perturbed slender vortex ring in three dimensions. The computations show that the predictions of the new schemes are in close agreement with each other and with results of previous models. The computations also show that, by only requiring adequate resolution of the filament centerline and allowing large integration time steps, the improved schemes defeat the stiffness of the equations of motion.

We close by noting that the present constructions differ from those of Zhou [9], who used an asynchronous splitting technique to treat stiff local terms and nonstiff nonlocal terms, and Hou *et al.* [12], who relied on a recasting of the filament equations of motion using generalized curvature coordinates. It appears to be possible, in principle, to seek further enhancement of performance by combining some of the present techniques with those outlined in [9] and [12]. Possible extensions also include the generalization of the present methods to accommodate core size variations along the filament centerline.

ACKNOWLEDGMENTS

REFERENCES

1. A. J. Chorin, The evolution of a turbulent vortex, *Comm. Math. Phys.* **83**, 517 (1982).

2. O. M. Knio and A. F. Ghoniem, Numerical study of a three-dimensional vortex method, *J. Comput. Phys.* **86**, 75 (1990).

3. R. Klein and O. M. Knio, Asymptotic vorticity structure and numerical simulation of slender vortex filaments, *J. Fluid Mech.* **284**, 275 (1995).

4. R. Klein, O. M. Knio, and L. Ting, Representation of core dynamics in slender vortex filament simulations, *Phys. Fluids* **8**, 2415 (1996).

5. O. M. Knio, L. Ting, and R. Klein, Interaction of a slender vortex filament with a rigid sphere: Dynamics and far-field noise, *J. Acoust. Soc. Am.* **103**, 83 (1998).

6. A. Callegari and L. Ting, Motion of a curved vortex filament with decaying vortical core and axial velocity, *SIAM J. Appl. Math.* **35**, 148 (1978).

7. R. Klein and L. Ting, Theoretical and experimental studies of slender vortex filaments, *Appl. Math. Lett.* **8**, 45 (1995).

8. J. T. Beale and A. Majda, High order accurate vortex methods with explicit velocity kernels, *J. Comput. Phys.* **58**, 188 (1985).

9. H. Zhou, On the motion of slender vortex filaments, *Phys. Fluids* **9**, 970 (1997).

10. A. S. Almgren, T. Buttke, and P. Colella, A fast adaptive vortex method in 3 dimensions, *J. Comput. Phys.* **113**, 177 (1994).

11. J. K. Salmon, M. S. Warren, and G. S. Winckelmans, Fast parallel tree codes for gravitational and fluid dynamical N-body problems, *Int. J. Supercomput. Appl.* **8**, 129 (1994).

12. T. Y. Hou, I. Klapper, and H. Si, Removing the stiffness of curvature in computing 3-D filaments, *J. Comput. Phys.* **143**, 628 (1998).