Parallel Transport Approach to Curve Framing

Andrew J. Hanson and Hui Ma
Department of Computer Science
Lindley Hall 215
Indiana University
Bloomington, IN 47405

January 11, 1995

Abstract

We propose an algorithm for generating a moving coordinate frame on a space curve based on the concept of parallel transport. Such algorithms can be used for creating ribbons, tubes, and camera orientations that are smoothly varying and controlled by the curve geometry itself. The more familiar Frenet frame approach suffers from ambiguity and sudden orientation changes when the curve straightens out momentarily. We compare the properties of alternative framing methods and point out when the parallel transport approach has unique advantages. We discuss a variety of implementation issues and illustrate the application of the algorithm to ribbons and tubes based on open and closed curves, as well as to the generation of moving camera orientations.

1 Introduction

We attack the problem of associating moving coordinate frames to three-dimensional space curves in ways that are well-understood mathematically and that have optimal behavior for certain classes of computer graphics applications. Classical differential geometry typically treats moving frames using the Frenet frame formalism because of its close association with a curve's curvature and torsion, which are coordinate-system independent [2, 9, 4]. The Frenet frame, unfortunately, has the property that it is undefined when the curve is even momentarily straight (has vanishing curvature), and exhibits wild swings in orientation around points where the osculating plane's normal has major changes in direction. We propose an alternative approach, the parallel-transport frame method (see Bishop [1]), which has a mathematically sound foundation and more appropriate behavior for computer graphics; in cases where the Frenet frame has desirable properties, a hybrid strategy is also feasible.

Typical computer graphics applications of the parallel-transport frame include the generation of ribbons and tubes from 3D space curves, and the generation of forward-facing camera orientations given an appropriate initial camera path. If the curve is coarsely refined, but is smooth enough to generate appropriate frame control points from the parallel-transport frame algorithm, the resulting frames can be used as control points for any desired degree of smooth spline interpolations using the methods of Shoemake [7] and Schlag [6]. Rotating camera orientations relative to a stable forward-facing frame can be added by various techniques such as that of Shoemake [8].

In Section 2, we introduce the basic mathematics of moving frames on space curves, emphasizing the parallel transport frame; the properties of the Frenet and parallel-transport frames are compared. Our algorithms are given in Section 3, where applications to the generation of ribbons, tubes, and camera frames are described. Appendices contain proofs and derivations of useful formulas.

2 The Differential Geometry of Space Curves

Our first goal is to define moving coordinate frames that are attached to a curve in 3D space. We will assume that our curves are defined in practice by a discrete sequence of points connected by straight line segments; thus numerical derivatives can be defined at each point.

2.1 Frenet Frames

The *Frenet frame* (see, e.g., [2, 4]) is defined as follows: If $\vec{\mathbf{x}}(t)$ is any thrice-differentiable space curve with non-vanishing second derivative, its tangent, binormal, and normal vectors at a point on the curve are given by

$$\begin{aligned} \vec{\mathbf{T}}(t) &= \frac{\vec{\mathbf{x}}'(t)}{\|\vec{\mathbf{x}}'(t)\|} \\ \vec{\mathbf{B}}(t) &= \frac{\vec{\mathbf{x}}'(t) \times \vec{\mathbf{x}}''(t)}{\|\vec{\mathbf{x}}'(t) \times \vec{\mathbf{x}}''(t)\|} \\ \vec{\mathbf{N}}(t) &= \vec{\mathbf{B}}(t) \times \vec{\mathbf{T}}(t) \ . \end{aligned} \tag{1}$$

The Frenet frame obeys the following differential equation in the parameter t (which is the origin of the requirement for one more order of differentiability beyond the second derivative):

$$\begin{bmatrix} \vec{\mathbf{T}}'(t) \\ \vec{\mathbf{N}}'(t) \\ \vec{\mathbf{B}}'(t) \end{bmatrix} = v(t) \begin{bmatrix} 0 & \kappa(t) & 0 \\ -\kappa(t) & 0 & \tau(t) \\ 0 & -\tau(t) & 0 \end{bmatrix} \begin{bmatrix} \vec{\mathbf{T}}(t) \\ \vec{\mathbf{N}}(t) \\ \vec{\mathbf{B}}(t) \end{bmatrix} \tag{2}$$

where $v(t) = \|\vec{\mathbf{x}}'(t)\|$ is the scalar magnitude of the curve derivative (often reparameterized to be unity, so that t becomes the arclength s), $\kappa(t)$ is the scalar curvature, and $\tau(t)$ is the

torsion. These quantities can in principle be calculated in terms of the parameterized or numerical local values of $\vec{\mathbf{x}}(t)$ and its first three derivatives as follows:

$$\begin{aligned} \kappa(t) &= \frac{\|\vec{\mathbf{x}}'(t) \times \vec{\mathbf{x}}''(t)\|}{\|\vec{\mathbf{x}}'(t))\|^3} \\ \tau(t) &= \frac{(\vec{\mathbf{x}}'(t) \times \vec{\mathbf{x}}''(t)) \cdot \vec{\mathbf{x}}'''(t)}{\|\vec{\mathbf{x}}'(t) \times \vec{\mathbf{x}}''(t)\|^2} \ . \end{aligned} \tag{3}$$

If we are given a non-vanishing curvature and a torsion as smooth functions of t, we can theoretically integrate the system of equations to find the unique numerical values of the corresponding space curve $\vec{\mathbf{x}}(t)$ (up to a rigid motion).

The Frenet equations permit an alternate form [9] that is equivalent to the equations of motion of a physical gyroscope being acted upon by a force $\vec{\mathbf{F}}_{\text{Frenet}}$:

$$\begin{aligned} \vec{\mathbf{T}}' &= \vec{\mathbf{F}}_{\text{Frenet}} \times \vec{\mathbf{T}} \\ \vec{\mathbf{N}}' &= \vec{\mathbf{F}}_{\text{Frenet}} \times \vec{\mathbf{N}} \\ \vec{\mathbf{B}}' &= \vec{\mathbf{F}}_{\text{Frenet}} \times \vec{\mathbf{B}} \ , \end{aligned} \tag{4}$$

where

$$\vec{\mathbf{F}}_{\text{Frenet}} = \tau\vec{\mathbf{T}} + \kappa\vec{\mathbf{B}} \ . \tag{5}$$

Intuitively, the Frenet frame's normal vector $\vec{\mathbf{N}}$ always points toward the center of the osculating circle [9]. Thus, when the orientation of the osculating circle changes drastically or the second derivative of the curve becomes very small, the Frenet frame behaves erratically or may become undefined.

2.2 Parallel Vector Fields

Before going on to introduce parallel transport frames as an alternative to the Frenet frame, let us spend a moment looking at parallel vector fields on curves in general; such vector fields are typically constructed by parallel transporting vectors along a curve.

Given a space curve $\vec{\mathbf{x}}(s)$ parameterized by arclength s, we define a vector field $\vec{\mathbf{V}}(s)$ to be *normal* if it is everywhere perpendicular to the curve's tangent $\vec{\mathbf{T}}(s) = \vec{\mathbf{x}}'(s)$. A normal vector field $\vec{\mathbf{V}}(s)$ is said to be *parallel* to the curve $\vec{\mathbf{x}}(s)$ if its derivative is tangential along the curve; that is, $\vec{\mathbf{V}}'(s) \parallel \vec{\mathbf{T}}(s)$. Such a vector field turns only as much as is necessary for it to remain normal.

More generally, we can define an arbitrary vector field $\vec{\mathbf{V}}$ along a curve $\vec{\mathbf{x}}$ to be parallel if its normal component is parallel and its tangential component is a constant multiple of the unit tangent field of $\vec{\mathbf{x}}$.

We call a curve $\vec{\mathbf{y}} = \vec{\mathbf{x}} + \vec{\mathbf{V}}$ a *parallel curve* of $\vec{\mathbf{x}}$.

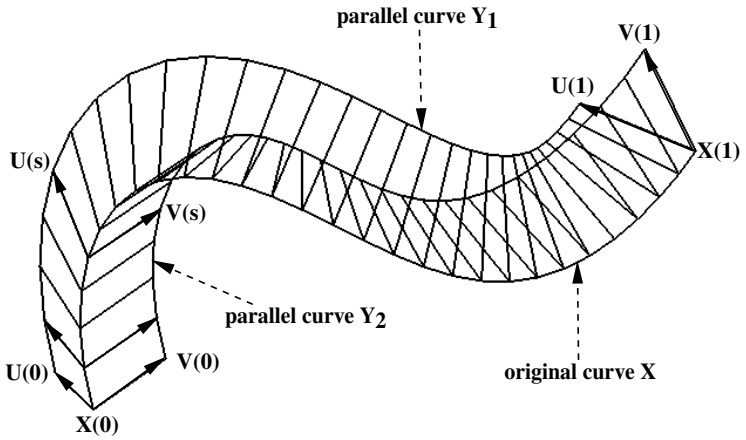

Figure 1: Properties of parallel vector fields.

Properties of Parallel Vector Fields. A curve $\vec{\mathbf{x}}$, a parallel normal vector field $\vec{\mathbf{V}}$, and the corresponding parallel curve $\vec{\mathbf{y}}$ have the following key properties:

(a) $\vec{\mathbf{V}}$ has constant length.

(b) $\vec{\mathbf{V}}$ is perpendicular to both $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$.

(c) $\vec{\mathbf{V}}$ is locally a segment of minimum length between the two curves if $\|\vec{\mathbf{V}}\|$ is sufficiently small.

(d) An initial normal vector $\vec{\mathbf{V}}_0$ at a point $\vec{\mathbf{x}}(s_0)$ generates a *unique* parallel field $\vec{\mathbf{V}}(s)$ on $\vec{\mathbf{x}}(s)$ such that $\vec{\mathbf{V}}(s_0) = \vec{\mathbf{V}}_0$.

(e) If normal vectors $\vec{\mathbf{V}}_0$ and $\vec{\mathbf{U}}_0$ generate parallel fields $\vec{\mathbf{V}}(s)$ and $\vec{\mathbf{U}}(s)$ respectively, the angle between $\vec{\mathbf{V}}(s)$ and $\vec{\mathbf{U}}(s)$ is constant along the curve. That is $\vec{\mathbf{V}}(s)\cdot\vec{\mathbf{U}}(s) = \vec{\mathbf{V}}_0\cdot\vec{\mathbf{U}}_0$ for all s.

These properties are illustrated in Figure 1.

2.3 Parallel Transport Frames

The *parallel transport frame* is an alternative approach to defining a moving frame that is well-defined even when the curve has vanishing second derivative. Because of the property (e) of parallel vector fields, we can parallel transport an orthonormal frame along a curve simply by parallel transporting each component of the frame.

The parallel transport frame is based on the observation that, while $\vec{\mathbf{T}}(t)$ for a given curve model is unique, we may choose any convenient arbitrary basis $(\vec{\mathbf{N}}_1(t), \vec{\mathbf{N}}_2(t))$ for the

remainder of the frame, so long as it is in the normal plane perpendicular to $\vec{\mathbf{T}}(t)$ at each point. If the derivatives of $(\vec{\mathbf{N}}_1(t), \vec{\mathbf{N}}_2(t))$ depend only on $\vec{\mathbf{T}}(t)$ and not each other, we can make $\vec{\mathbf{N}}_1(t)$ and $\vec{\mathbf{N}}_2(t)$ vary smoothly throughout the path regardless of the curvature. We therefore have the alternative frame equations

$$\begin{bmatrix} \vec{\mathbf{T}}' \\ \vec{\mathbf{N}}_1' \\ \vec{\mathbf{N}}_2' \end{bmatrix} = v \begin{bmatrix} 0 & k_1 & k_2 \\ -k_1 & 0 & 0 \\ -k_2 & 0 & 0 \end{bmatrix} \begin{bmatrix} \vec{\mathbf{T}} \\ \vec{\mathbf{N}}_1 \\ \vec{\mathbf{N}}_2 \end{bmatrix} \ . \tag{6}$$

One can show (see, e.g, Bishop [1]) that

$$\kappa(t) = \left((k_1)^2 + (k_2)^2\right)^{1/2} \tag{7}$$

$$\theta(t) = \arctan\left(\frac{k_2}{k_1}\right) \tag{8}$$

$$\tau(t) = -\frac{d\theta(t)}{dt} \ , \tag{9}$$

so that k_1 and k_2 effectively correspond to a Cartesian coordinate system for the polar coordinates κ, θ with $\theta = -\int \tau(t)\, dt$. The orientation of the parallel transport frame includes the arbitrary choice of integration constant θ_0, which disappears from τ (and hence from the Frenet frame) due to the differentiation.

The parallel transport frame equations permit an alternate torque-like form analogous to Eqs. (4,5):

$$\begin{aligned} \vec{\mathbf{T}}' &= \vec{\mathbf{F}}_{\rm PT} \times \vec{\mathbf{T}} \\ \vec{\mathbf{N}}_1' &= \vec{\mathbf{F}}_{\rm PT} \times \vec{\mathbf{N}}_1 \\ \vec{\mathbf{N}}_2' &= \vec{\mathbf{F}}_{\rm PT} \times \vec{\mathbf{N}}_2 \ , \end{aligned} \tag{10}$$

where

$$\vec{\mathbf{F}}_{\rm PT} = \det \begin{bmatrix} k_1 & k_2 \\ \vec{\mathbf{N}}_1 & \vec{\mathbf{N}}_2 \end{bmatrix} = -k_2 \vec{\mathbf{N}}_1 + k_1 \vec{\mathbf{N}}_2 \ . \tag{11}$$

We may write these equations for an arbitrary vector in the form of a single determinant, with the result

$$\vec{\mathbf{V}}' = \det \begin{bmatrix} k_1 & k_2 & 0 & 0 \\ N_1^{(x)} & N_2^{(x)} & V^{(x)} & \hat{\mathbf{x}} \\ N_1^{(y)} & N_2^{(y)} & V^{(y)} & \hat{\mathbf{y}} \\ N_1^{(z)} & N_2^{(z)} & V^{(z)} & \hat{\mathbf{z}} \end{bmatrix} \ . \tag{12}$$

Both Eq. (12) and Eq. (6) suggest straightforward generalizations to N dimensions.

Some Important Properties. We observe that Eqs. (7) and (11) imply that the *magnitude* of the force in the parallel transport case is $\|\vec{\mathbf{F}}_{\mathrm{PT}}\| = \kappa$, and thus is less than the Frenet force of Eq. (5), which gives $\|\vec{\mathbf{F}}_{\mathrm{Frenet}}\| = (\kappa^2 + \tau^2)^{1/2}$. We see that the parallel transport frame in some sense washes out the torsion, leaving the minimal possible changes in direction required by the curvature alone. An informative way of visualizing these properties would be to plot the paths of the equivalent quaternions on the 3-sphere in four-dimensional space [5].

As with the Frenet equations, we can begin with a pair of functions $(k_1(t), k_2(t))$ and an initial frame, and then integrate any alternate form of the frame equations to find the curve $\vec{\mathbf{x}}(t)$ up to a rigid motion.

These equations also give an abstract construction for a parallel-transported frame. In Section 3, we will present a much simpler numerical method. In Appendix C, we give a proof showing that the frame field generated by our algorithm correctly approximates the parallel-transport frame of an underlying smooth curve as the curve segment length approaches zero.

2.4 Comparison of Frenet and Parallel Transport Frames

The contrast between the properties of the parallel-transport (PT) frame and the Frenet frame is best seen by looking at some examples.

In Figure 2a, we show the frames for a convex plane curve; the Frenet and PT frames are identical. However, as soon as the curve has inflection points in the plane, as shown in Figure 2b, one sees that the Frenet frame's normal components instantly switch sign at each inflection point, while the PT frame has no such discontinuities; if the curvature remains zero along a straight line segment, the Frenet frame provides no prescription for defining a smooth transition from the frame coming into the straight segment and the (possibly radically different) frame leaving the straight segment.

Next, we look at a non-planar curve drawn on a "roof-top," which exhibits momentarily vanishing curvature and a radical change in the normal to the osculating circle; again, as illustrated in Figures 3a,b, the parallel transport frame is well-behaved and the Frenet frame is not smooth enough to be used as the basis for a ribbon or tube construction.

Finally, we look at a cylindrical helix, which has the property that the torsion is a constant along the whole curve. In this case the Frenet frame returns to the same orientation each time the helix passes through the same axial line on the tube, as shown in Figure 4a. This behavior is in fact desirable in some applications. However, at the same time, it hides an essential property: non-zero torsion means the parallel transport frame is rotating with a constant angular velocity that washes out the torsion in such a way as to *reduce* the total change in the frame orientation at the end of one circuit! Viewed in terms of the quaternion picture of rotations, the path traveled in the space of unit quaternions will be shorter. Figure 4b illustrates the way in which the parallel-transport frame changes by a constant rotation in the normal plane with each cycle around the cylinder.

We note that, since we can create a closed curve by attaching a planar curve to two

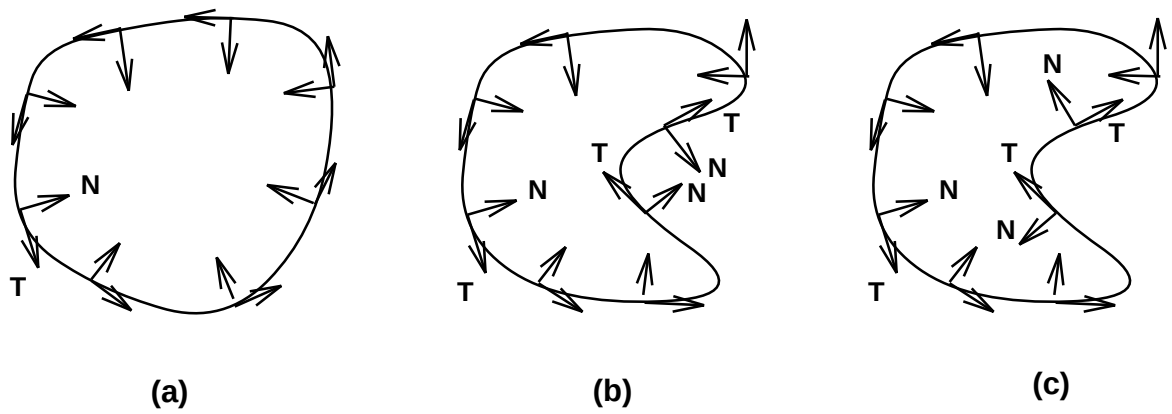

Figure 2: Comparing the Frenet and parallel-transport frames on plane curves. (a) Convex curve; both frames are identical, with the third component of the frame pointing out of the paper. (b) The Frenet frame on a non-convex curve (with inflection points) reverses the direction of the normal fields at each inflection point, so the direction of $\vec{\mathbf{B}}$ is into the paper on the indented portion of the curve. (c) A parallel-transport frame on the same non-convex curve maintains continuity in the direction of the third component of the frame throughout the curve.

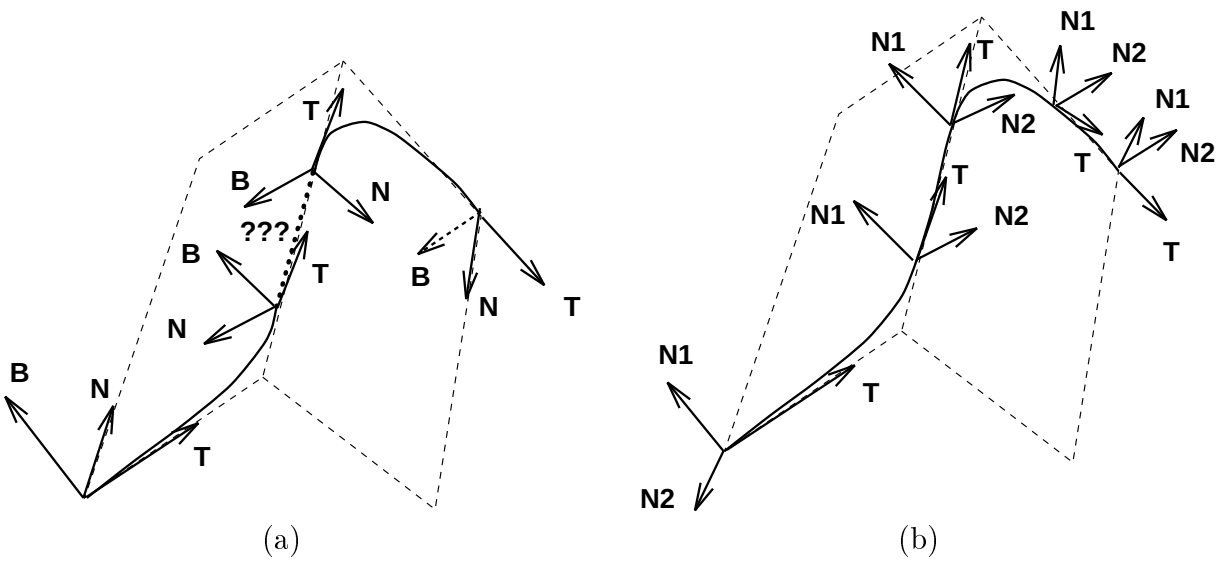

Figure 3: Comparing the Frenet and parallel-transport frames on a "roof-top." (a) The Frenet frame becomes undefined on the straight line at the peak, then changes abruptly as the curve descends the right side. (b) The parallel transport frame is smooth throughout. (Note: the initial orientation of the normal plane is arbitrary.)

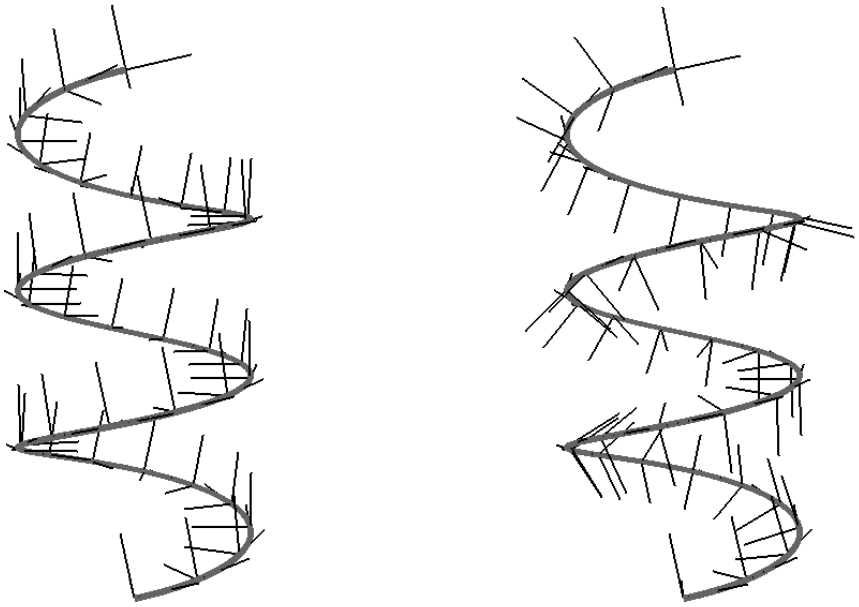

(a) (b)

Figure 4: (a) The Frenet frame of a 3D helix, which has constant torsion; the frame is identical after each turn. (b) The parallel transport frame on a helix, showing how the torsion produces a constant angular velocity or "spin" about the moving tangent vector. The total amount of spin experienced after each turn depends on the pitch of the helix, which determines the torsion.

points of the helix tangent to a single plane, it is possible for closed curves to have parallel transport frames that do not match up after one full circuit of the curve; we will discuss a simple correction procedure for this situation below (see Section 3.1).

3 Algorithms and Implementation Issues

In practice, we never have smooth curves in numerical applications, but only piecewise linear curves that are presumed to be approximations to differentiable curves. We will need to compute the tangents to a curve given by the set of points $\{\vec{\mathbf{x}}_i\}$. Without making any assumptions about the shape of the curve approximated by the points, the best we can do is to compute the tangent using a formula involving neighboring points such as

$$\vec{\mathbf{T}}_i = \frac{\vec{\mathbf{x}}_{i+1} - \vec{\mathbf{x}}_{i-1}}{\|\vec{\mathbf{x}}_{i+1} - \vec{\mathbf{x}}_{i-1}\|} \ .$$

If we are willing to make some assumptions, such as taking any three neighboring points to represent the arc of a circle, we may compute the tangent to that circle at the middle point;

this *osculating circle* [9], whose center lies on the intersection line of the planes perpendicular to each chord is discussed in detail in the Appendix B.

If the curve is locally straight, i.e., $\vec{\mathbf{x}}''(t) = 0$ or $\vec{\mathbf{T}}_{i+1} = \vec{\mathbf{T}}_i$, then there is no locally-determinable coordinate frame component in the plane normal to $\vec{\mathbf{T}}$; a non-local definition must be used to decide on the remainder of the frame once $\vec{\mathbf{T}}$ is determined, and this is what the parallel-transport algorithm provides.

The basic steps of the parallel transport algorithm for piecewise linear curves can be formulated as follows (see diagram in Figure 5):

INPUT: (1) A list of unit tangent vectors $\{\hat{\mathbf{T}}_i\}$, $i = 0, \ldots, N$;
(2) An initial normal vector $\vec{\mathbf{V}}_0$, $\vec{\mathbf{V}}_0 \perp \hat{\mathbf{T}}_0$.

OUTPUT: A list of parallel-transported normal vectors $\{\vec{\mathbf{V}}_i\}$,
$i = 1, \ldots, N$, $\vec{\mathbf{V}}_i \perp \hat{\mathbf{T}}_i$.

ALGORITHM:

for $i \leftarrow 0$ **to** $N - 1$ **step** 1
  $\vec{\mathbf{B}} \leftarrow \hat{\mathbf{T}}_i \times \hat{\mathbf{T}}_{i+1}$;
  if $\|\vec{\mathbf{B}}\| = 0$ **then**
    $\vec{\mathbf{V}}_{i+1} \leftarrow \vec{\mathbf{V}}_i$;
  else
    $\hat{\mathbf{B}} \leftarrow \vec{\mathbf{B}}/\|\vec{\mathbf{B}}\|$;
    $\theta \leftarrow \arccos(\hat{\mathbf{T}}_i \cdot \hat{\mathbf{T}}_{i+1})$; // $0 \le \theta \le \pi$
    $\vec{\mathbf{V}}_{i+1} \leftarrow \mathbf{R}(\hat{\mathbf{B}}, \theta) * \vec{\mathbf{V}}_i$;
    // Rotate by angle θ about $\hat{\mathbf{B}}$ (see Appendix A)
  end if
end for

Remarks.

- If $\hat{\mathbf{T}}_i$ is nearly parallel to $\hat{\mathbf{T}}_{i+1}$, then θ is close to zero, so the rotation matrix $\mathbf{R}$ is close to the identity, making the value of $\vec{\mathbf{B}}$ irrelevant.
- Since all actions are rotations, the length of any transported vector is preserved automatically.
- In fact, the input of the algorithm is not restricted to normal vectors; the method correctly parallel transports any vector. What happens, in effect, is that the normal component of the vector is parallel transported, while the tangential component is repeatedly rotated to coincide in direction with the current tangent vector.

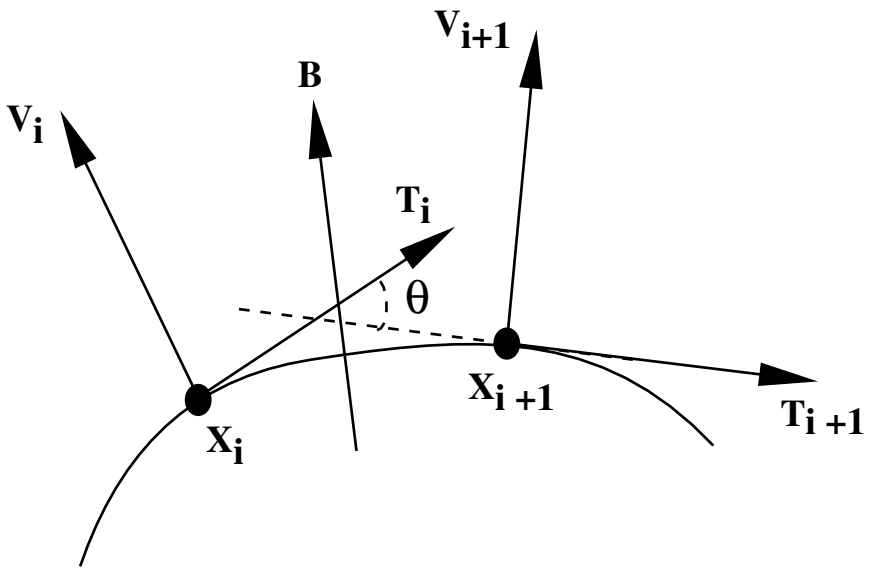

Figure 5: Diagram illustrating the geometric quantities used in the parallel transport algorithm.

- Given an initial frame with one vector in the direction of the curve tangent and two normal vectors, we can thus construct the parallel-transport frame field by applying the algorithm to each of the normal components separately. Orthonormality is automatically preserved due to property (e) in Section 2.2. Alternatively, since the third component of the triad is dependent on the other two, we may parallel transport only one component and compute the second by taking the cross-product of the first with the tangent vector.

- A common method for the generation of tubes is what might be called the "projection method." In this approach, one takes the current unit normal vector $\hat{\mathbf{V}}_i$, the current unit tangent vector $\hat{\mathbf{T}}_i$, and the next unit tangent vector $\hat{\mathbf{T}}_{i+1}$, and computes $\hat{\mathbf{V}}_{i+1}$ as follows:

$$\hat{\mathbf{V}}_{i+1} = \frac{\hat{\mathbf{V}}_i - \hat{\mathbf{T}}_{i+1}(\hat{\mathbf{V}}_i \cdot \hat{\mathbf{T}}_{i+1})}{\|\hat{\mathbf{V}}_i - \hat{\mathbf{T}}_{i+1}(\hat{\mathbf{V}}_i \cdot \hat{\mathbf{T}}_{i+1})\|} \ .$$

  Parallel transport explicitly preserves the length of the current binormal component of the vector, while this method may change the length due to the required renormalization. Therefore projection is not equivalent to the parallel transport method. In fact, the projection method always fails when the angle is greater than or equal to $\pi/2$, while the parallel-transport algorithm handles this case naturally.

- The long chain of matrix multiplications may incur some numerical error; in some situations, this error can be reduced by representing the initial frame as a quaternion (see, e.g., [7]). Since each rotation $\mathbf{R}_i$ can be expressed directly in terms of a quaternion using the same parameters $q = (\cos\frac{\theta}{2}, \hat{\mathbf{B}}\sin\frac{\theta}{2})$, one can carry out the numerical

computation of the frame change over the entire curve by quaternion multiplication instead of 3×3 matrix multiplication.

3.1 Closed Curves and Spinning

For closed curves, we let $\vec{\mathbf{x}}_{N+1} = \vec{\mathbf{x}}_0$. The Frenet frame, which is defined only by local curve properties, returns to its initial value. In contrast, the parallel transport frame will in general *not* return to its initial orientation. The angular difference α between the initial and final frames is determined by the torsion,

$$\alpha = -\oint \tau(s)\, ds \mod (2\pi)$$

where s is the arclength.

We can *heuristically* create a continuous frame that is aligned after one trip around the closed curve by adding an additional "spin" around the tangent direction at each vertex. For example, if the curve is described be a set of points $\{\vec{\mathbf{x}}_i\}_{i=0}^{N+1}$, where $\vec{\mathbf{x}}_{N+1} = \vec{\mathbf{x}}_0$, and the partial curve lengths are

$$L_i = \sum_{j=1}^{i} \|\vec{\mathbf{x}}_j - \vec{\mathbf{x}}_{j-1}\|, \qquad 1 \le i \le (N+1) \ ,$$

then for $i = 1, \ldots, N$ we can choose

$$\begin{aligned} \alpha_i &= \alpha \frac{L_i}{L_{N+1}} \\ \vec{\overline{\mathbf{V}}}_i &= \mathbf{R}(\hat{\mathbf{T}}_i, \alpha_i) * \vec{\mathbf{V}}_i \ . \end{aligned}$$

Examples are shown in Figure 6.

Hybrid frames. If the application requires, we can add additional spin to the framing field; for example, if the Frenet frame and the torsion are well-defined throughout an entire curve segment, one can undo the entire torsion by removing the mod(2π) in the equation above, and the resulting frame is the same as the Frenet frame if the initial frame is chosen as the Frenet frame.

3.2 Sweeping Tubes and Ribbons

To generate a ribbon or tube, we can use the parallel transport method to create a complete structure by sweeping an initial cross-section composed of vectors $\{\vec{\mathbf{V}}_0^k\}; k = 0, \ldots, K$ through a curve. A thickness function $\{r_i^k\}; i = 0, \ldots, N,\ k = 0, \ldots, K$ can also be applied as

$$\vec{\overline{\mathbf{V}}}_i^k = r_i^k \vec{\mathbf{V}}_i^k \ .$$

Examples are shown in Figure 7.

Smoothing corners. Creating tessellations of ribbons and tubes is complicated by the fact that the outside corners can be handled easily without introducing self-intersections, while the inside corners may have colliding normal line segments if the lengths of the normal segments are too long. For the purposes of this paper, we will assume that the burden is on the user to supply a curve that is sufficiently detailed and sufficiently smooth, as well as a tube or ribbon cross-section that does not unreasonably cause self-intersections when swept along the curve.

With this assumption, there is at least one natural way of smoothing out a structure in the context of the parallel transport frame.

We take two points on the curve, $\vec{\mathbf{X}}_i$ and $\vec{\mathbf{X}}_{i+1}$, the corresponding tangents $\hat{\mathbf{T}}_i$ and $\hat{\mathbf{T}}_{i+1}$, and a pair of normal vectors $\vec{\mathbf{V}}_i$ and $\vec{\mathbf{V}}_{i+1}$. The smoothed swept structure, either a ribbon or a tube, is then defined by linearly interpolating along the line segment $\vec{\mathbf{X}}_{i+1} - \vec{\mathbf{X}}_i$ and performing a constant-angular-velocity rotation (i.e., a Slerp in the terminology of Shoemake [7]) from $\vec{\mathbf{V}}_i$ to $\vec{\mathbf{V}}_{i+1}$ to define the intervening smooth curve points.

There are undoubtedly many other solutions to this problem.

3.3 Parallel Transport Camera Frames

In many (though not all) applications of camera animation, it is desirable to have the camera gaze direction pointing forward along a space curve throughout the motion. Typical examples would include a flight looking through the front window of an airplane cockpit, riding on a roller coaster, or sliding down a bannister. Other applications require the camera gaze direction to remain in some fixed skew orientation relative to the camera path, e.g., a passenger looking out an airplane window. All such applications are easily accommodated using the parallel transport mechanism applied to an initial camera orientation. An example is shown in Figure 8. In this figure, we started with a closed curve, the trefoil knot, and a parallel transport frame that did not return to its initial value after one circuit; we then used the method described in Section 3.1 to adjust the frame field. The resulting frame is distinct from the Frenet frame for this curve.

Smooth Orientation Changes. If the discrete points on the camera path are not closely spaced, using the discrete rotations of the basic parallel transport algorithm will result in unacceptably jerky camera motion. This is easily corrected by reinterpreting the camera frame fields at each curve vertex as *key frames* for a Shoemake-style quaternion interpolation [7, 6]; then the camera orientations will move smoothly throughout the curve, passing through or near the vertex frame fields, depending on the particular spline chosen.

In addition, certain classes of camera rolls and spins can be handled naturally as well. The simplest such case is the one in which the additional motions take place with respect to the parallel-transport frame of the curve. More complex motions can be handled by a variety of interpolation methods (see, e.g., [8]).

4 Conclusion and Future Directions

In this paper, we have introduced a coherent method for generating smooth moving frames based on the geometry of a space curve. Our principal new tool is the parallel-transport frame, supplemented when appropriate by the classical Frenet frame. The tools we have introduced provide a number of mathematically well-defined options for producing ribbons, tubes, and camera orientation sequences automatically from the intrinsic geometry of a given space curve.

Possible topics for future investigation include the treatment of parallel transport on higher dimensional manifolds such as surfaces and volumes, parallel transport along curves restricted to such manifolds (that is, curves in curved ambient spaces), and a variety of problems having to do with creating mathematically satisfactory smoothed corners for tubes and ribbons derived from piecewise linear curves.

Acknowledgments

This work was supported in part by NSF grant IRI-91-06389. We thank Ji-Ping Sha for useful discussions and Bruce Solomon for bringing reference [1] to our attention.

References

[1] Bishop, R. L. There is more than one way to frame a curve. *Amer. Math. Monthly 82*, 3 (March 1975), 246–251.

[2] Eisenhart, L. P. *A Treatise on the Differential Geometry of Curves and Surfaces.* Dover, New York, 1909 (1960).

[3] Foley, J., van Dam, A., Feiner, S., and Hughes, J. *Computer Graphics, Principles and Practice*, second ed. Addison-Wesley, 1990. page 227.

[4] Gray, A. *Modern Differential Geometry of Curves and Surfaces.* CRC Press, Inc., Boca Raton, FL, 1993.

[5] Hanson, A. J., and Ma, H. Visualization flow with quaternion frames. In *Proceedings of Visualization '94* (1994), IEEE Computer Society Press. .

[6] Schlag, J. Using geometric constructions to interpolate orientation with quaternions. In *Graphics Gems II*, J. Arvo, Ed. Academic Press, 1991, pp. 377–380.

[7] Shoemake, K. Animating rotation with quaternion curves. In *Computer Graphics* (1985), vol. 19, pp. 245–254. Proceedings of SIGGRAPH 1985.

[8] Shoemake, K. Fiber bundle twist reduction. In *Graphics Gems IV*, P. Heckbert, Ed. Academic Press, 1994, pp. 230–236.

[9] Struik, D. J. *Lectures on Classical Differential Geometry.* Addison-Wesley, 1961.

A Rotation Matrix

The rotation matrix used in the algorithm in the main text is [3]

$$\mathbf{R}(\theta,\hat{\mathbf{n}}) = \begin{bmatrix} c+(n_1)^2(1-c) & n_1n_2(1-c)-sn_3 & n_3n_1(1-c)+sn_2 \\ n_1n_2(1-c)+sn_3 & c+(n_2)^2(1-c) & n_3n_2(1-c)-sn_1 \\ n_1n_3(1-c)-sn_2 & n_2n_3(1-c)+sn_1 & c+(n_3)^2(1-c) \end{bmatrix} , \tag{13}$$

where $c = \cos\theta$ and $s = \sin\theta$ and $\hat{\mathbf{n}}\cdot\hat{\mathbf{n}} = 1$.

B Computing the Tangent of Osculating Circle

Suppose we are given three points $(\vec{\mathbf{x}}_0,\, \vec{\mathbf{x}}_1,\, \vec{\mathbf{x}}_2)$ in a Euclidean space of arbitrary dimension. Then we may compute the center and radius of the osculating circle, as well as the tangent direction at $\vec{\mathbf{x}}_1$ as follows. First use a Gram-Schmidt procedure to compute the direction $\vec{\mathbf{u}}$ perpendicular to the chord $(\vec{\mathbf{x}}_1 - \vec{\mathbf{x}}_0)$:

$$\vec{\mathbf{u}} \;=\; (\vec{\mathbf{x}}_2 - \vec{\mathbf{x}}_1) - (\vec{\mathbf{x}}_1 - \vec{\mathbf{x}}_0)\frac{(\vec{\mathbf{x}}_2 - \vec{\mathbf{x}}_1)\cdot(\vec{\mathbf{x}}_1 - \vec{\mathbf{x}}_0)}{(\vec{\mathbf{x}}_1 - \vec{\mathbf{x}}_0)\cdot(\vec{\mathbf{x}}_1 - \vec{\mathbf{x}}_0)} \,. \tag{14}$$

Next, note that the center $\vec{\mathbf{x}}_c$ may be written as a vector from chord midpoint $\vec{\mathbf{a}} = (\vec{\mathbf{x}}_1 + \vec{\mathbf{x}}_0)/2$ in the perpendicular direction

$$\vec{\mathbf{x}}_c = \vec{\mathbf{a}} + t\hat{\mathbf{u}}$$

where $\hat{\mathbf{u}} = \vec{\mathbf{u}}/\|\vec{\mathbf{u}}\|$ and t is to be computed. Then, since each of the vectors $(\vec{\mathbf{x}}_0,\, \vec{\mathbf{x}}_1,\, \vec{\mathbf{x}}_2)$ is a distance R from $\vec{\mathbf{x}}_c$,

$$\begin{aligned} R^2 &= \|\vec{\mathbf{x}}_c - \vec{\mathbf{x}}_1\|^2 && (15)\\ &= \|\vec{\mathbf{x}}_c - \vec{\mathbf{x}}_2\|^2 && (16)\\ &= \|\vec{\mathbf{a}} - \vec{\mathbf{x}}_1\|^2 + 2t\hat{\mathbf{u}}\cdot(\vec{\mathbf{a}} - \vec{\mathbf{x}}_1) + t^2 && (17)\\ &= \|\vec{\mathbf{a}} - \vec{\mathbf{x}}_2\|^2 + 2t\hat{\mathbf{u}}\cdot(\vec{\mathbf{a}} - \vec{\mathbf{x}}_2) + t^2 && (18) \end{aligned}$$

Since $\hat{\mathbf{u}}\cdot(\vec{\mathbf{a}} - \vec{\mathbf{x}}_1) = 0$ by construction and the t^2 terms cancel, we find

$$t = \frac{\|\vec{\mathbf{a}} - \vec{\mathbf{x}}_2\|^2 - \|\vec{\mathbf{a}} - \vec{\mathbf{x}}_1\|^2}{2\hat{\mathbf{u}}\cdot(\vec{\mathbf{x}}_2 - \vec{\mathbf{a}})} \tag{19}$$

The value of R follows at once, and the direction of the tangent vector may be computed by applying Gram-Schmidt to $(\vec{\mathbf{x}}_0 - \vec{\mathbf{x}}_1)$:

$$\vec{\mathbf{T}} \;=\; (\vec{\mathbf{x}}_0 - \vec{\mathbf{x}}_1) - (\vec{\mathbf{x}}_c - \vec{\mathbf{x}}_1)\frac{(\vec{\mathbf{x}}_c - \vec{\mathbf{x}}_1)\cdot(\vec{\mathbf{x}}_0 - \vec{\mathbf{x}}_1)}{(\vec{\mathbf{x}}_c - \vec{\mathbf{x}}_1)\cdot(\vec{\mathbf{x}}_c - \vec{\mathbf{x}}_1)} \tag{20}$$

C Formal Properties: Correctness of Continuous Limit

For any twice-differentiable curve $\vec{\mathbf{x}}(s)$ with tangent vector $\vec{\mathbf{T}}(s)$ and any corresponding coordinate frame $(\vec{\mathbf{T}}(s), \vec{\mathbf{N}}(s), \vec{\mathbf{B}}(s))$, the following formula holds for some functions $k_1(s)$, $k_2(s)$, $\tau(s)$ because of the orthonormality constraints:

$$\begin{bmatrix} \vec{\mathbf{T}}'(s) \\ \vec{\mathbf{N}}'(s) \\ \vec{\mathbf{B}}'(s) \end{bmatrix} = \begin{bmatrix} 0 & k_1 & k_2 \\ -k_1 & 0 & \tau \\ -k_2 & -\tau & 0 \end{bmatrix} \begin{bmatrix} \vec{\mathbf{T}}(s) \\ \vec{\mathbf{N}}(s) \\ \vec{\mathbf{B}}(s) \end{bmatrix} \quad . \tag{21}$$

Theorem 1. Let $\vec{\mathbf{V}}(s)$ be a unit normal vector field on the curve $\vec{\mathbf{x}}(s)$, let $(\vec{\mathbf{T}}(s), \vec{\mathbf{N}}(s), \vec{\mathbf{B}}(s))$ be any framing on the curve, and let $\alpha(s)$ be the angle between $\vec{\mathbf{V}}(s)$ and $\vec{\mathbf{N}}(s)$. Then $\vec{\mathbf{V}}(s)$ is parallel if and only if, for any s_1 and s_2,

$$\alpha(s_2) = (\alpha(s_1) - \int_{s_1}^{s_2} \tau(s)\, ds) \mod (2\pi) \; . \tag{22}$$

Proof:

$$\vec{\mathbf{V}} = \cos(\alpha)\vec{\mathbf{N}} + \sin(\alpha)\vec{\mathbf{B}}$$

$$\vec{\mathbf{V}}' = \cos(\alpha)\vec{\mathbf{N}}' - \alpha' \sin(\alpha)\vec{\mathbf{N}} + \sin(\alpha)\vec{\mathbf{B}}' + \alpha' \cos(\alpha)\vec{\mathbf{B}}$$

Plugging in Eq. (21), and remembering that, by definition, $\vec{\mathbf{V}}$ is parallel if and only if both the coefficients of $\vec{\mathbf{N}}$ and $\vec{\mathbf{B}}$ are zero, we find

$$-\sin(\alpha)\alpha' - \sin(\alpha)\tau = 0$$

$$+\cos(\alpha)\alpha' + \cos(\alpha)\tau = 0$$

which is equivalent to

$$\alpha' = -\tau$$

or

$$\alpha(s_2) = (\alpha(s_1) - \int_{s_1}^{s_2} \tau(s)\, ds) \mod (2\pi) \; .$$

QED.

Remark. If the frame is a Frenet frame, τ is the classical torsion.

Proof of smooth limit. First, define the *norm of a tessellation* to be the maximum length of any line segment in the curve. We wish to show that, in the limit as the norm of a curve's tessellation approaches zero, the result of parallel transporting a normal vector using our algorithm approaches the parallel transport vector field on a smooth curve.

We assume that the input $\{\vec{\mathbf{T}}_i\}$ to the algorithm are the actual tangents of the smooth curve.

Let

$$\vec{\mathbf{B}}_i = \frac{\vec{\mathbf{T}}_i \times \vec{\mathbf{T}}_{i+1}}{\|\vec{\mathbf{T}}_i \times \vec{\mathbf{T}}_{i+1}\|}$$

$$\vec{\mathbf{N}}_i = \vec{\mathbf{T}}_i \times \vec{\mathbf{B}}_i \ .$$

Then $(\vec{\mathbf{T}}_i, \vec{\mathbf{N}}_i, \vec{\mathbf{B}}_i)$ gives a local frame at $\vec{\mathbf{x}}_i$.

Let

$$\vec{\mathbf{V}}_i = \cos(\alpha_i)\vec{\mathbf{N}}_i + \sin(\alpha_i)\vec{\mathbf{B}}_i$$

The algorithm preserves $\sin(\alpha_i)\vec{\mathbf{B}}_i$, and rotates $\cos(\alpha_i)\vec{\mathbf{N}}_i$ to $\cos(\alpha_i)(\vec{\mathbf{T}}_{i+1} \times \vec{\mathbf{B}}_i)$, hence

$$\vec{\mathbf{V}}_{i+1} = \cos(\alpha_i)(\vec{\mathbf{T}}_{i+1} \times \vec{\mathbf{B}}_i) + \sin(\alpha_i)\vec{\mathbf{B}}_i \ .$$

Projecting out the components of $\vec{\mathbf{V}}_{i+1}$ in the direction of $\vec{\mathbf{N}}_{i+1}$ and $\vec{\mathbf{B}}_{i+1}$, we find, with $\cos(\theta_i) = \vec{\mathbf{B}}_{i+1} \cdot \vec{\mathbf{B}}_i$,

$$\begin{aligned}
\vec{\mathbf{V}}_{i+1} \cdot \vec{\mathbf{N}}_{i+1} &= \cos(\alpha_i)(\vec{\mathbf{T}}_{i+1} \times \vec{\mathbf{B}}_i) \cdot \vec{\mathbf{N}}_{i+1} + \sin(\alpha_i)\vec{\mathbf{B}}_i \cdot \vec{\mathbf{N}}_{i+1} \\
&= \cos(\alpha_i)(\vec{\mathbf{N}}_{i+1} \times \vec{\mathbf{T}}_{i+1} \cdot \vec{\mathbf{B}}_i) + \sin(\alpha_i)(\vec{\mathbf{B}}_i \cdot \vec{\mathbf{T}}_{i+1} \times \vec{\mathbf{B}}_{i+1}) \\
&= \cos(\alpha_i)(\vec{\mathbf{B}}_{i+1} \cdot \vec{\mathbf{B}}_i) + \sin(\alpha_i)(\vec{\mathbf{B}}_{i+1} \times \vec{\mathbf{B}}_i \cdot \vec{\mathbf{T}}_{i+1}) \\
&= \cos(\alpha_i)\cos(\theta_i) - \sin(\alpha_i)\sin(\theta_i) \\
&= \cos(\alpha_i + \theta_i) \ , \\
\vec{\mathbf{V}}_{i+1} \cdot \vec{\mathbf{B}}_{i+1} &= \cos(\alpha_i)(\vec{\mathbf{T}}_{i+1} \times \vec{\mathbf{B}}_i) \cdot \vec{\mathbf{B}}_{i+1} + \sin(\alpha_i)\vec{\mathbf{B}}_i \cdot \vec{\mathbf{B}}_{i+1} \\
&= \cos(\alpha_i)(\vec{\mathbf{B}}_i \times \vec{\mathbf{B}}_{i+1} \cdot \vec{\mathbf{T}}_{i+1}) + \sin(\alpha_i)(\vec{\mathbf{B}}_i \cdot \vec{\mathbf{B}}_{i+1}) \\
&= \cos(\alpha_i)\sin(\theta_i) + \sin(\alpha_i)\cos(\theta_i) \\
&= \sin(\alpha_i + \theta_i) \ .
\end{aligned}$$

Therefore $\alpha_{i+1} = (\alpha_i + \theta_i) \mod (2\pi)$. By induction, for any indices $i_1 < i_2$,

$$\alpha_{i_2} = \alpha_{i_1} + \sum_{j=i_1}^{i_2-1} \theta_j \mod (2\pi) \ .$$

Now as the norm of the tessellation approaches zero, the discrete frame field $\{(\vec{\mathbf{T}}_i, \vec{\mathbf{N}}_i, \vec{\mathbf{B}}_i)\}_{i=0}^{N}$ approaches a framing $(\vec{\mathbf{T}}, \vec{\mathbf{N}}, \vec{\mathbf{B}})$ on the smooth curve. Since $\vec{\mathbf{B}}_i \perp \vec{\mathbf{T}}_{i+1}$ and $\vec{\mathbf{B}}_{i+1} \perp \vec{\mathbf{T}}_{i+1}$, then

$$\vec{\mathbf{B}}' \perp \vec{\mathbf{T}} \ \Rightarrow \ k_2 = 0 \ \Rightarrow \ \vec{\mathbf{B}}' = -\tau\vec{\mathbf{N}} \ .$$

Since $\|\vec{\mathbf{B}}_{i+1} - \vec{\mathbf{B}}_i\| \approx |\theta_i|$, then

$$\tau = -\theta'$$

Fixing $\vec{\mathbf{x}}_{i_1}$ and $\vec{\mathbf{x}}_{i_2}$, and letting the norm of the tessellation approach zero, we have

$$\begin{aligned} \alpha_{i_2} &= (\alpha_{i_1} + \sum_{j=i_1}^{i_2-1} \theta_j) \mod (2\pi) \\ &\longrightarrow (\alpha_{i_1} + \int_{\vec{\mathbf{x}}_{i_1}}^{\vec{\mathbf{x}}_{i_2}} \theta'(s)\, ds) \mod (2\pi) \\ &= (\alpha_{i_1} - \int_{\vec{\mathbf{x}}_{i_1}}^{\vec{\mathbf{x}}_{i_2}} \tau(s)\, ds) \mod (2\pi) \ . \end{aligned}$$

Hence, by **Theorem 1**, the parallel transport algorithm approaches the parallel transport vector field as the norm of the tessellation goes to zero. QED.

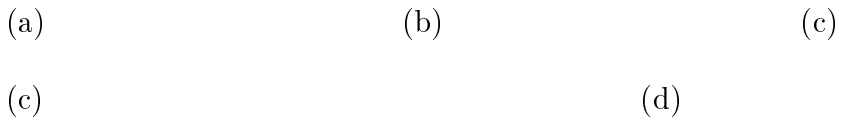

Figure 6: Parallel transport on closed curves. (a) A ribbon produced by parallel transport of an initial vector does not generally return to the same orientation after one circuit. (b) Closure can be enforced by distributing the angular deficit around the curve. (c) Adding multiples of π to the total axial rotation gives any desired amount of twisting (this example adds 6π).

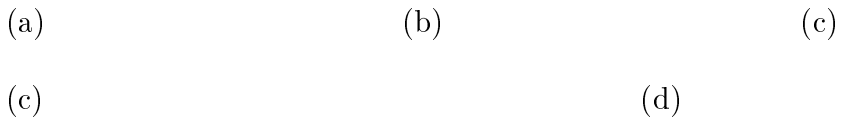

Figure 7: Creating ribbons and tubes using parallel transport. (a) A ribbon generated by a pair of vectors. (b) A tube with a circular cross-section. (c) A tube with a star-shaped cross-section and a varying radius.

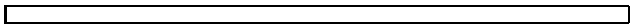

Figure 8: An application of the parallel transport frame to the generation of a moving camera orientation automatically determined by the geometry of the flight path itself. In this illustration, the orientation of the aircraft represents the camera orientation. When the ribbon appears blue, the top of the aircraft faces, while when the ribbon is green, we are looking at the bottom of the aircraft.