

Numerical Simulation and Immersive Visualization of Hairpin Vortices

H. M. Tufo,¹ P. F. Fischer,² M. E. Papka,² K. Blom³

August 20, 1999

Abstract

To better understand the vortex dynamics of coherent structures in turbulent and transitional boundary layers, we consider direct numerical simulation of the interaction between a flat-plate-boundary-layer flow and an isolated hemispherical roughness element. Of principal interest is the evolution of hairpin vortices that form an interlacing pattern in the wake of the hemisphere, lift away from the wall, and are stretched by the shearing action of the boundary layer. Using animations of unsteady three-dimensional representations of this flow, produced by the vtk toolkit and enhanced to operate in a CAVE virtual environment, we identify and study several key features in the evolution of this complex vortex topology not previously observed in other visualization formats.

1 Introduction

We present visualization results of numerically generated hairpin vortex formation and evolution in incompressible boundary layer flows. At moderate flow speeds, hairpin vortices provide an example of an observable, organized transition process from steady, two-dimensional laminar flows to unsteady, three-dimensional turbulent flows. Consequently, hairpin vortices are of interest in the study of transitional and turbulent boundary layers, where they have been frequently observed experimentally. In this study, the hairpin vortices are generated in the wake of a hemispherical roughness element embedded in a flat-plate boundary layer, following closely the experiments of Acarlar and Smith [1] and, to a lesser extent, those of Klebanoff et al. [13]. Numerical studies of hairpin vortices have also been undertaken by Singer and his colleagues under slightly different flow conditions [2, 20, 21].

The basic flow configuration is shown in Fig. 1. A time-independent velocity profile is prescribed across the upstream entrance of the domain. Experimental observations [1] indicate that the flow is symmetric about the plane $y=0$, so the flow is computed only in the half-domain shown. At sufficiently high nondimensional flow speeds, or Reynolds numbers, the steady boundary layer region near the plate is destabilized by the hemisphere, resulting in periodic shedding of hairpin vortices in the wake.

The simulation results are analyzed by using X-window-based software developed specifically for the numerical methods employed and using interactive software built on top of the “Visualization Toolkit” (vtk) to drive stereo visualization environments such as the CAVETM (Cave Automatic Virtual Environment) and ImmersaDeskTM. The resulting immersive visualization proceeds at a sufficiently high frame rate that the

¹Center on Astrophysical Thermonuclear Flashes, University of Chicago, Chicago, IL 60637.

²Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439.

³Electrical and Computer Eng., Iowa State University, Ames, IA, 50011

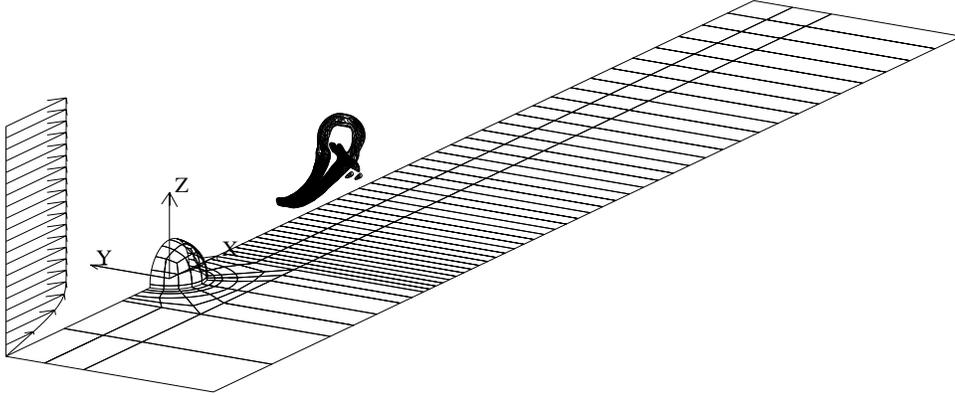


Figure 1: Computational domain showing inlet velocity profile, flat plate, hemisphere, and isolated hairpin vortex in hemisphere wake. For clarity, the vortex has been reflected about the symmetry plane.

hairpin evolution can be readily integrated by eye, thereby allowing one a comprehensive understanding of the dynamics of this complex flow.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of the numerical method and specific flow conditions considered. Section 3 describes the immersive visualization environment and software tools used. Section 4 presents a mixture of quantitative and qualitative visual results used to analyze this flow. We close in Section 5 with a summary of our results.

2 Numerical Method and Flow Simulation

This section describes the underlying discretizations, domain configuration, and boundary conditions for the the hairpin vortex simulations.

2.1 Spectral Element Method

The hairpin vortex simulation is based on numerical integration of the unsteady incompressible Navier-Stokes equations,

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla P + \frac{1}{Re} \nabla^2 \mathbf{u} \\ -\nabla \cdot \mathbf{u} &= 0, \end{aligned}$$

coupled with appropriate boundary conditions on the velocity, \mathbf{u} . Semi-implicit time stepping is employed in which the nonlinear convective terms are treated explicitly, while the viscous and pressure terms are treated implicitly. Spatial discretization is based on the spectral element method (SEM), which is a high-order weighted residual technique similar to the finite element method. Within each element, basis functions are based on tensor-products of N th-order Lagrange polynomials [5, 6, 14]. The nodes of the Lagrange polynomials are taken to be the Gauss-Lobatto-Legendre (GLL) quadrature points, so that high-order GLL quadrature can be substituted for the integrals required for the residual evaluation. The discretization is illustrated in Fig. 2, which shows a three-element mesh in \mathbb{R}^2 with the GLL grid for the case $N = 4$. Also shown is the reference (r, s) coordinate system used for all function evaluations. Functions in the mapped

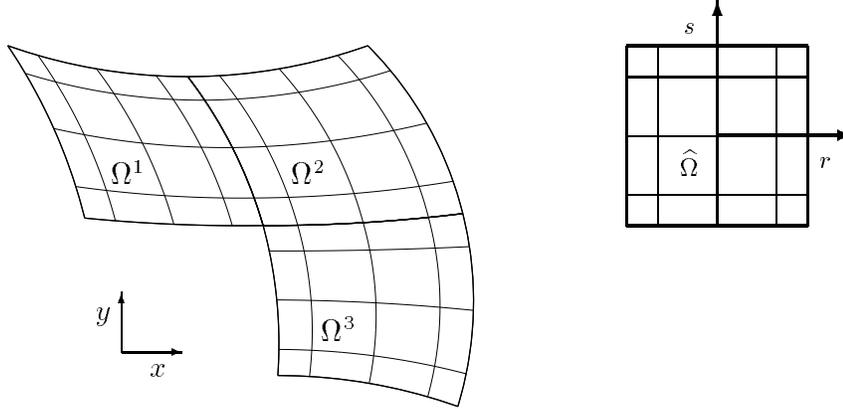


Figure 2: Spectral element discretization in \mathbb{R}^2 showing GLL nodal lines for $(K, N) = (3, 4)$.

coordinates are of the form

$$u(\mathbf{x}^k(r, s))|_{\Omega^k} = \sum_{i=0}^N \sum_{j=0}^N u_{ij}^k h_i^N(r) h_j^N(s), \quad (1)$$

where u_{ij}^k is the nodal basis coefficient; h_i^N is the Lagrange polynomial of degree N based on the GLL quadrature points, $\{\xi_j^N\}_{j=0}^N$; and $\mathbf{x}^k(r, s)$ is the coordinate mapping from the reference domain, $\hat{\Omega} := [-1, 1]^d$, to Ω^k . The use of the GLL basis for the interpolants leads to efficient quadrature for the weighted residual schemes and greatly simplifies operator evaluation in the case of deformed elements.

For problems having smooth solutions, such as the incompressible Navier-Stokes equations, convergence is exponential in N , despite the fact that only C^0 continuity is enforced across elemental interfaces. The resulting (minimal) numerical dissipation and dispersion are ideal for transitional flows of the type considered here, which can be very sensitive to such nonphysical effects. The rapid convergence is demonstrated in Table 1, which shows the computed growth rates of a small-amplitude Tollmien-Schlichting wave superimposed on plane Poiseuille channel flow at $Re = 7500$, following [5, 16]. The amplitude of the perturbation is 10^{-5} , so that the nonlinear Navier-Stokes results can be compared directly with linear theory to roughly five significant digits. Three error measures are computed: $error_1$ and $error_2$ are the relative amplitude errors at the end of the first and second periods, respectively, and $error_g$ is the error in the growth rate at a convective time of 50. From Table 1, it is clear that doubling the number of points in each spatial direction yields several orders of magnitude reduction in error, implying that just a small increase in resolution is required for very good accuracy. The significance of this is underscored by the fact that, in three dimensions, the effect on the number of gridpoints scales as the cube of the relative savings in resolution.

Table 1: Spatial convergence, O-S problem: $K = 15$, $\Delta t = .003125$

N	$E(t_1)$	$error_1$	$E(t_2)$	$error_2$	$error_g$
7	1.11498657	0.003963	1.21465285	0.037396	0.313602
9	1.11519192	0.003758	1.24838788	0.003661	0.001820
11	1.11910382	0.000153	1.25303597	0.000986	0.004407
13	1.11896714	0.000016	1.25205855	0.000009	0.000097
15	1.11895646	0.000006	1.25206398	0.000014	0.000041

2.1.1 Time Advancement

The Navier-Stokes time advancement is based on the second-order operator splitting methods developed in [15]. The convective term is expressed as a material derivative, and the resultant form is discretized via a stable second-order backward difference formula:

$$\frac{\tilde{\mathbf{u}}^{n-2} - 4\tilde{\mathbf{u}}^{n-1} + 3\mathbf{u}^n}{2\Delta t} = S(\mathbf{u}^n),$$

where $S(\mathbf{u}^n)$ is the linear symmetric Stokes problem to be solved implicitly, and $\tilde{\mathbf{u}}^{n-q}$ is the velocity field at time step $n - q$ computed as the explicit solution to a pure convection problem. The subintegration of the convection term permits values of Δt corresponding to convective CFL numbers of 2–5, thus significantly reducing the number of (expensive) Stokes solves.

The Stokes problem is of the form

$$\begin{bmatrix} \mathbf{H} & -\mathbf{D}^T \\ -\mathbf{D} & 0 \end{bmatrix} \begin{pmatrix} \mathbf{u}^n \\ \underline{p}^n \end{pmatrix} = \begin{pmatrix} \mathbf{B}\mathbf{f} \\ \underline{0} \end{pmatrix}$$

and is also treated by second-order splitting, resulting in subproblems of the form

$$H\underline{u}_i^n = \underline{f}_i,$$

for the velocity components ($i = 1, \dots, 3$), and

$$E\underline{p}^n = \underline{q}^n,$$

for the pressure. Here, H is a diagonally dominant Helmholtz operator representing the parabolic component of the momentum equations and is readily treated via Jacobi-preconditioned conjugate gradients; $E := \mathbf{D}\mathbf{B}^{-1}\mathbf{D}^T$ is the Stokes Schur complement governing the pressure; and \mathbf{B} is the (diagonal) mass matrix in the velocity space. E is a consistent Poisson operator and is effectively preconditioned by using the overlapping additive Schwarz procedure of Dryja and Widlund [4]. Further details of the discretization and solvers may be found in [5, 6].

2.1.2 Spectral Element Operators

The computational efficiency of spectral element methods derives from their tensor-product basis. To illustrate, we express the stiffness matrix for an undeformed element k in \mathbb{R}^2 as a sum of tensor products of one-dimensional operators,

$$A^k = \widehat{B}_y \otimes \widehat{A}_x + \widehat{A}_y \otimes \widehat{B}_x, \quad (2)$$

where \widehat{A}_* and \widehat{B}_* are the one-dimensional stiffness and mass matrices associated with the respective spatial dimensions. If $\underline{u}^k = u_{ij}^k$ is the matrix of nodal values on element k , then a typical matrix-vector product required of an iterative solver takes the form

$$\begin{aligned} (A^k \underline{u}^k)_{lm} &= \sum_{i=0}^N \sum_{j=0}^N (\widehat{B}_{y,mj} \widehat{A}_{x,li} u_{ij}^k + \widehat{A}_{y,mj} \widehat{B}_{x,li} u_{ij}^k) \\ &= \widehat{A}_x \underline{u}^k \widehat{B}_y^T + \widehat{B}_x \underline{u}^k \widehat{A}_y^T. \end{aligned}$$

The latter form illustrates how the tensor-product basis leads to matrix-vector products ($A\underline{u}$) being recast as *matrix-matrix* products, a feature central to the efficiency of spectral element methods. Similar forms result for other operators and for complex geometries.

The parallel implementation follows a standard message-passing-based single-program-multiple-data (SPMD) model [9] in which contiguous groups of elements are distributed to processors. Since iterative solvers are used, the principal communication kernel is the gather-scatter operation required for the residual vector assembly procedure. Because data is always stored on an element-by-element basis, the gather-scatter procedure required for residual evaluation is combined into a single communication phase wherein shared nodal values are exchanged and summed. This is a single local-to-local transformation, rather than separate gather and scatter phases common to many finite element implementations. The gather-scatter operation is implemented by using a stand-alone MPI/NX-based message-passing utility that supports a vector mode for problems having multiple degrees of freedom per vertex as well as a general set of commutative/associative operations [23]. The easy-to-use interface requires only two calls:

$$\mathbf{handle} = \mathbf{gs-init}(global\text{-}node\text{-}numbers, n), \quad \text{and} \quad \mathbf{ierr} = \mathbf{gs-op}(u, op, handle),$$

where *global-node-numbers()* associates the n local values contained in the vector $u()$ with their global counterparts, and *op* denotes the reduction operation performed on shared elements of $u()$. Communication overhead is further reduced through use of a recursive spectral bisection based element partitioning scheme to minimize the number of vertices shared among processors [18].

The spectral element code runs on a number of distributed-memory platforms, including the Intel Paragon, Cray T3E-600, SGI Origin 2000 and IBM SP. Simulations of the hairpin vortex problem have been run on 2048 333 MHz nodes of ASCI-Red in both in single- and dual-processor mode, with sustained performance of 319 GF being achieved for the latter [22].

2.2 Flow Simulation

For the hairpin vortex simulations, we considered several meshes and computational domains to test grid and domain independence, with our two primary production meshes having $K = 1021$ and $K = 1535$ spectral elements. All of the meshes tested have a unit-radius hemisphere located with its center at $(x, y, z) = (0, 0, .1)$. A short cylinder of height $dz = .1$ connects the hemisphere to the flat plate, which is located at $z = 0$. The cylinder was added so that a thin layer of elements could be placed near the plate to ensure adequate resolution of the boundary layer. The resultant 10% increase in the height of the hemisphere closely corresponds to height increases resulting from glue used to fasten the hemispheres in the experiments of Klebanoff et al. [13]. We consider a Reynolds number range of $Re_R := UR/\nu = 450\text{--}850$, where U is the free-stream velocity, R is the hemisphere radius, and ν is the kinematic viscosity of the fluid. Flow is in the positive x direction with incoming flow $\mathbf{u} = (u_B(y), 0, 0)$, where $u_B(y)$ represents a Blasius profile with $\delta_{99} = 1.2R$. The inflow boundary is located at $x = -8.4$ for the $K = 1021$ mesh and at $x = -10.0$ for the $K = 1535$ mesh. To reduce the computational burden, we impose reflection symmetry about the y -plane. Homogeneous Neumann boundary conditions are prescribed for the velocity at the outlet, located at $x = 30$. The upper boundary and the right boundary are also taken to be symmetry planes, corresponding to stress-free boundary conditions. These are located at $z = -6.5$ and $y = -6.4$ for $K = 1021$ and $z = -8.0$ and $y = -8.4$ for $K = 1535$. We note that, in addition to having different size domains and number of elements, the two meshes have different transitions between spherical and rectangular elements and between the inflow boundary and hemisphere, with the $K = 1535$ mesh having a smoother transition for both.

To transition between Reynolds numbers, we employ a sine ramp, $Re(t) = Re_i + (Re_f - Re_i) \sin(gt)$, where Re_i is the initial Reynolds number, Re_f is the final Reynolds number, and g is the growth rate, typically set such that duration of the ramp is 40–80 convective time units. Figure 3a shows the pressure trace at $\mathbf{x} = (2.4, 0.0, 1.2)$ for an 80 time unit transition between Reynolds number 675 and 700. Once

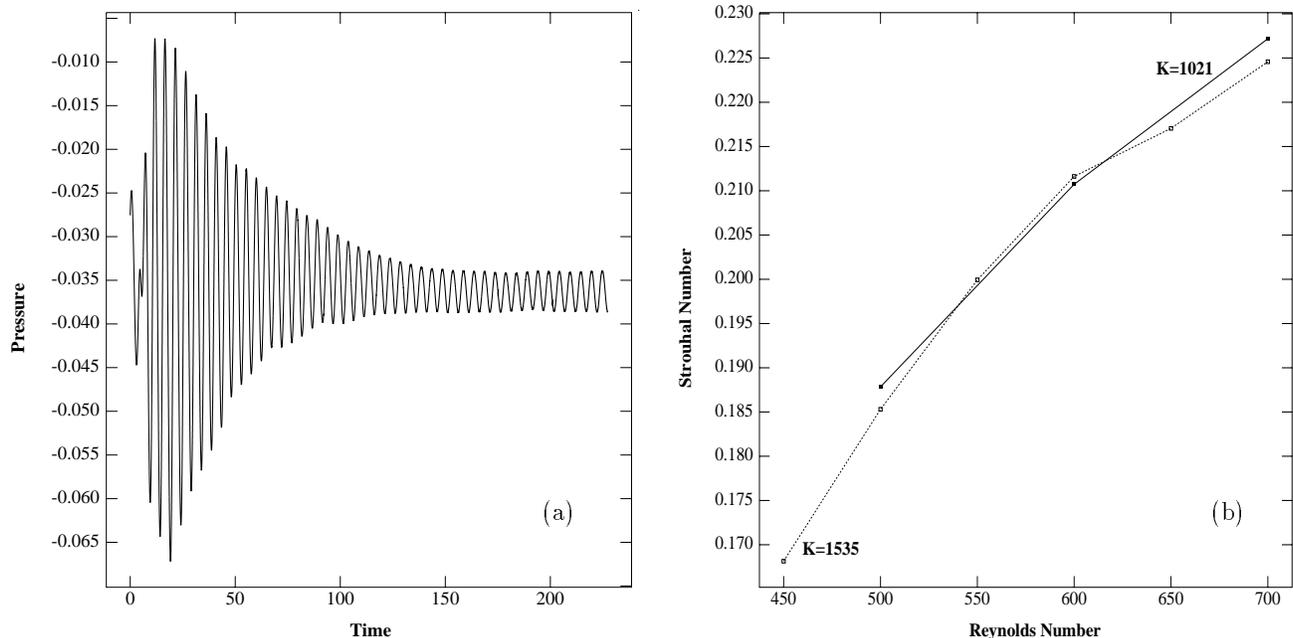


Figure 3: (a) Transition from $Re = 675$ to 700. (b) Strouhal number versus Reynolds number for the $K = 1021$ and $K = 1535$ meshes.

the transition has been made, the flow is driven until it settles into a steady periodic state. Of particular interest in these simulations is determination of the Strouhal number, $S = fk/U_k$, where f is the shedding frequency, $k = 1.1$ is the total height of the roughness element, and U_k is the tip velocity without the hemisphere present. U_k is calculated from a two-dimensional channel flow simulation on a domain identical in dimension to the hemisphere symmetry plane (centerplane) but without the hemisphere. The frequency, f , is determined from history plots similar to Fig. 3a. Strouhal numbers for the $K = 1021$ and 1535 meshes are shown in Fig. 3b. For this study $3 \leq N \leq 15$, with the maximum N considered for Reynolds number 450 being 9 and for Reynolds number 700 being 15. We note that the maximum discrepancy between the two meshes is less than 2% over the range of Reynolds numbers considered here and that spatial and temporal convergence checks were performed at several Reynolds numbers to verify convergence. We further note that our Strouhal numbers compare favorably to those in [1] and [13].

3 Visualization System

The visualization of the spectral element data is addressed through a multistage process. The first step uses a menu-driven X-window-based postprocessor developed specifically for spectral elements. It exploits the full accuracy of the high-order Lagrangian basis (1) when interpolating off-grid point values or computing derived quantities such as velocity gradients or vorticity. In addition, the postprocessor can map spectral element data onto unstructured hexahedral meshes of arbitrary density. This data is then processed for use in a second visualization package designed for immersive visualization of the general mesh data. The primary components of this second mode are built using the vtk library, which is an open-source software system for visualization that provides a high-level abstraction for constructing scientific visualization applications [19], and the CAVE library, which enables projection and exploration of immersive stereo images [3]. An overview of the entire visualization process is outlined in Fig. 4.

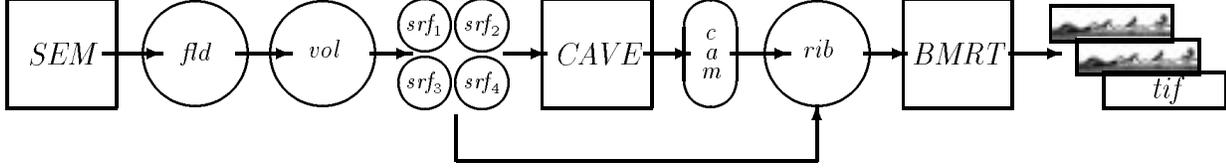


Figure 4: Visualization process: *SEM* produces numerical output (*fd*), which is translated by the *X*-postprocessor into vtk unstructured cell format (*vol*). These are processed into individual polygonal datasets (*srf*), representing isosurfaces of the vortex field at distinct threshold values, τ , or at different times. This data is then viewed interactively at low resolution in the *CAVE*, where camera paths are captured. The output (*cam*) is combined with (*srf*) to produce rib files (*rib*). These are processed by the *Blue Moon Rendering Tools (BMRT)* to yield the final high-resolution output (*tif*).

3.1 Vortex Identification

Vortex identification starts in the postprocessing phase and is based on the λ_2 method of Jeong and Hussain [10]. Identification of a vortex in viscous flows is challenging because the classic rules governing vortex dynamics generally apply only in the inviscid limit. In boundary layer flows, viscosity is non-negligible, and standard approaches such as integrating vortex lines or using pressure minima or vorticity maxima can lead to improper vortex identification. Jeong and Hussain have established a robust criterion for the identification of vortex (or coherent) structures in viscous flows based on the eigenvalues of the symmetric 3×3 tensor

$$M_{ij} := \sum_{k=1}^3 (\Omega_{ik}\Omega_{kj} + S_{ik}S_{kj}), \quad (3)$$

where

$$\Omega_{ij} := \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad S_{ij} := \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) \quad (4)$$

represent the symmetric and antisymmetric components of the velocity gradient tensor, $\nabla \mathbf{u}$. To minimize noise, the gradients are computed using the original polynomial description of the data, that is, by differentiating (1). Given the three (real) eigenvalues of M at each grid point, a vortex core is identified as any contiguous region having two negative eigenvalues. If the eigenvalues are sorted such that $\lambda_1 \leq \lambda_2 \leq \lambda_3$, then any region for which $\lambda_2 < 0$ corresponds to a vortex core. One advantage of this approach is that vortices can be identified as isosurfaces of a well-defined scalar field. Moreover, the criterion $\lambda_2(\mathbf{x}) < 0$ is scale invariant, so there is in principle no ambiguity in selecting which isosurface value to render. In practice, one usually biases the isosurface to a value that is below zero by a small fraction of the full dynamic range in order to avoid noise in regions where the velocity is close to zero.

3.2 Isosurface Extraction and Immersive Visualization

To develop an unsteady immersive rendering of the vortex evolution, we begin with a set of full-volume dumps of the primitive variables (u, v, w, p) at a selected number of time steps (typ. 150) covering one period of the vortex shedding cycle. For each time step, we compute λ_2 at each grid point with the *X*-window based postprocessor. The high-order accuracy of the spectral element basis (1) is essential for this step. The resultant scalar field is then interpolated onto an unstructured hexahedral mesh along with the pressure. This volume of data is then passed to a surface stripper (built on top of vtk) to extract isosurfaces at a given

threshold, τ , and produce a set of triangles corresponding to $\lambda_2 = \tau$. Typically, $-2 \leq \tau \leq -1$, out of a range of $[-30, 40]$.

Visualization within the CAVE environment requires real-time response, and, because of the time required to generate isosurfaces, interactive generation of the surfaces is impossible. Figure 5 plots the average time required for the generation of an individual isosurface at a given threshold versus Reynolds number. The rise in time reflects the increasing complexity of the flow field with increasing Re . In addition to isosurface extraction, smoothing and triangulation algorithms are applied as part of the visualization pipeline to improve performance in the CAVE. The resulting surface files are saved to disk in the vtk polydata format. This format contains all the information necessary to construct the geometric representation of the λ_2 isosurfaces with corresponding pressures at each vertex.

For immersive visualization, the surface files are loaded into another vtk-based application that is built on top of the CAVE library and is designed to run on all members of the CAVE family of display technology, including the ImmersaDesk. (The vtk classes for renderer and renderwindow have been extended to operate within the CAVE [7].) The CAVE is a $10' \times 10' \times 10'$ cube of rear-projected screens, illustrated in Fig. 5 (right), which allows the user to rapidly explore data from a number of different viewing angles. The ImmersaDesk is single $5'$ rear-projected screen. Both are capable of displaying immersive stereo images. The immersive application makes full use of the CAVE library's rich set of tools that enable users to explore and interrogate individual surface files, as well as a composited series of surface files for time- or parameter-dependent data. For the $Re = 700$ surface files a rate of 15 frames per second is achieved, which is more than sufficient to produce a realistic sense of motion to the user. Control over viewing location, frame rate, and playback are provided by the combination of a navigational wand and virtual menu. In addition, users can define a color table to map other scalars onto the λ_2 isosurfaces. At present, only the pressure is mapped.

Within the immersive environment it is also possible to generate high-quality renderings of the surface data. The user can navigate to points of interest within the data set and take virtual snapshots or virtual movies of the experience, capturing the needed information for high-resolution rendering. This is done either

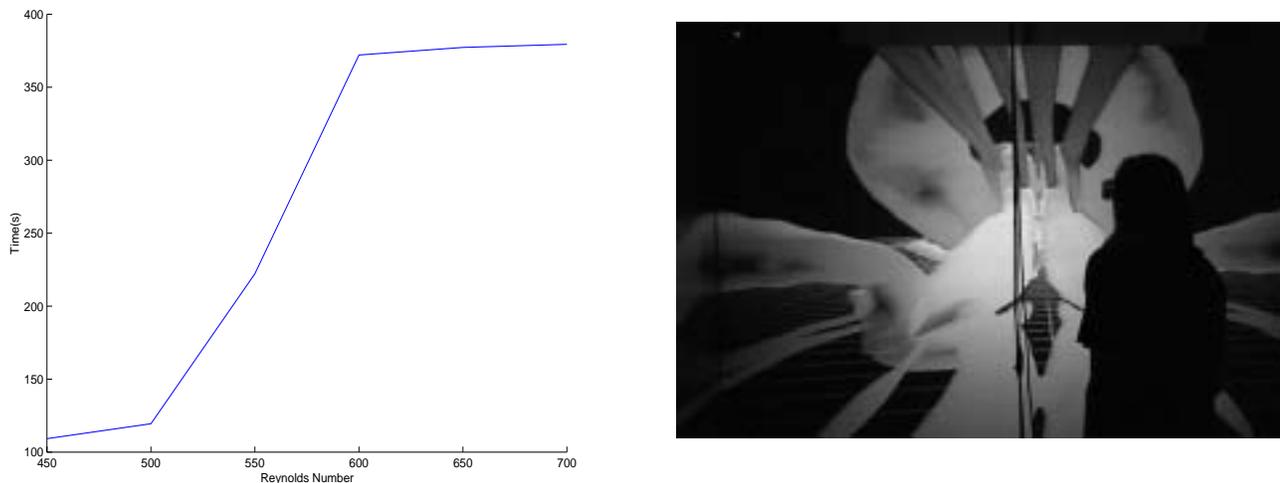


Figure 5: Two stages in the visualization process. Isosurfaces for a series of images are precalculated off-line. (Average calculation times versus Reynolds number are shown on the left.) Precomputed isosurfaces are then rendered at 15 frames per second in the CAVE (right) allowing interactive navigation of unsteady surface data and selection of specific views for high-resolution rendering.

by immediately generating a RenderMan RIB file using a class provided within vtk for a snapshot, or by capturing the camera position and orientation information in virtual movie mode to allow off-line rendering. (Virtual movie mode requires an additional offline step for generating the RIB files, because the size of the files would interfere with the required real-time response of the virtual environment.) Once the RIB files are created, they are processed by using *Blue Moon Rendering Tools* (BMRT), a freely available collection of rendering tools that adhere to the RenderMan interface standard [8]. BMRT enables users to capture points of interest from within the data set in a high-quality format suitable for publication or videotape.

4 Results

To put the dynamic results into perspective, we begin with a series of two-dimensional plots in the symmetry plane, $y = 0$. Figure 6 shows the results of a simulation at $Re = 700$ using 1021 elements of order $N = 11$. Time-averaged velocity profiles are shown in (a), and corresponding rms values, $u_{rms} := \langle u'^2 \rangle^{\frac{1}{2}}$, are shown in (b). Here $u' := u - \langle u \rangle$ is the fluctuating component of u , and $\langle \cdot \rangle$ denotes a single-period time average. The profiles have maximum values of 1.075 and 0.282 for u and u_{rms} , respectively. We see in (a) the Blasius profile at the inlet, the separated and recirculating wake region at $x = 2$, followed by a gradual recovery until the outlet, where the mean velocity profile is fuller than the inflow profile. From the rms profiles (b), it is clear that the flow upstream of the hemisphere is essentially steady. The immediate wake region at $x = 2$ exhibits remarkably little unsteadiness. The passage of the hairpin vortices is evidenced by the strong rms fluctuations in the wake region, which also reveal the lifting of the vortices away from the plate. Further downstream, there is a significant growth in activity near the wall, as can be seen by the peak in the rms profile at the outlet.

To indicate the structure of individual hairpin vortices at a fixed instance, we present contours of several quantities in Fig. 6 (c-e). Pressure contours are shown in (c). The vortex cores are readily identified by the low-pressure zones, shown in bold. Contours of spanwise vorticity (d) also show the hairpin vortices and, in particular, the movement of the heads and tips away from the plate. Finally, (e) shows contours of $\lambda_2 \in [-30, 0]$, revealing the intersection of the vortex tips with the symmetry plane and the presence of a steady horseshoe vortex at the base of the hemisphere near $x = -1$.

Results for $Re = 450$ are shown in Fig. 7. This is just above the critical Reynolds number at which the flow transitions from a steady to a steady-periodic state. In contrast to the $Re = 700$ case, the wake deficit in the symmetry plane profile (a) is very pronounced, even 30 radii downstream of the hemisphere. The rms profiles (b) grow with increasing x , although the peak rms value of .063 is a fraction of that for the $Re = 700$ case. The shear layer observed in (a) indicates that the growth in rms values in (b) can be interpreted as the onset of a Kelvin-Helmholtz instability: slight oscillations in the wall-normal velocity component translate into significant u' fluctuations in the shear-layer region where $\frac{\partial u}{\partial z}$ is large. The oscillations and wave like nature of the flow are quite evident in the contour plot of spanwise vorticity, shown at a fixed instant in the shedding cycle in (c).

While the views in Figs. 6 and 7 provide a fair amount of information about the hairpin vortex evolution, they fail to reveal any three-dimensional details. To see these, we rely on output from the visualization process described in Section 3. Figures 8 and 9 show $\lambda_2 = -1$ isosurfaces for $Re = 700$. We observe several vortex features, some of which we had not identified prior to viewing the unsteady animation in an immersive environment. Figure 8 shows the classic horseshoe vortex (a) upstream of the hemisphere, which is also commonly found upstream of end-mounted cylinders (as evidenced by snowdrift patterns at the bases of trees and telephone poles). Moving downstream, we see the interlacing of the hairpin vortex tails (b), as observed by Acarlar and Smith [1]; the hairpin head (c); and a vortex “bridge” (d), which is a common form

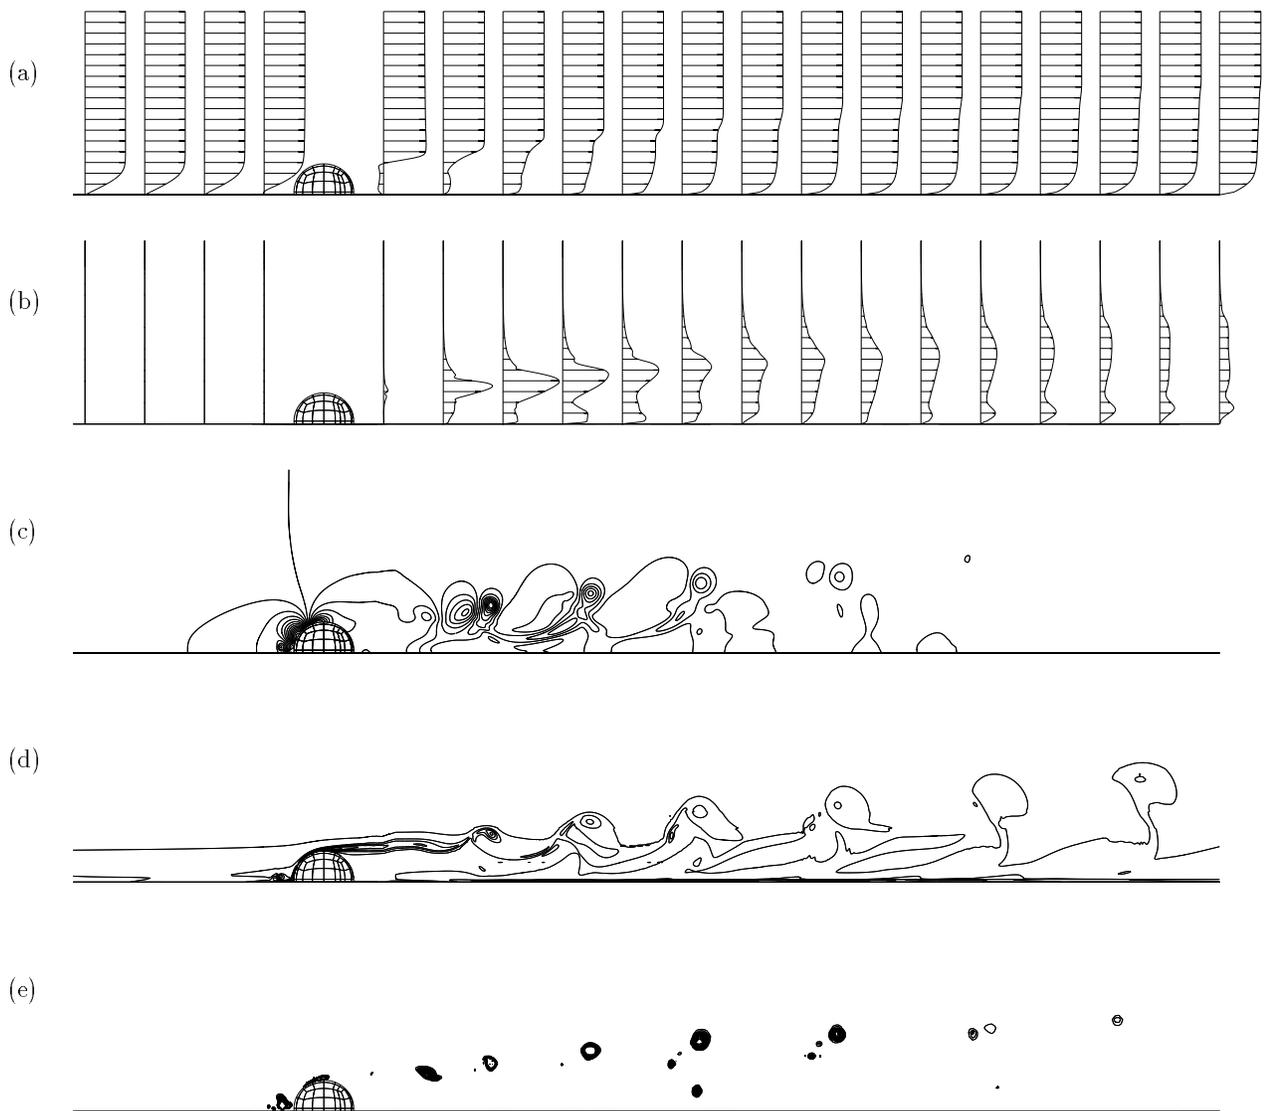


Figure 6: Symmetry plane data for $Re = 700$, $K = 1021$, $N = 11$ (a-b), and $N = 13$ (c-e): (a) mean velocity profiles, $\langle u \rangle$; (b) rms velocity $\langle u'^2 \rangle^{1/2}$; (c) pressure contours; (d) spanwise vorticity, ω_y ; (e) contours of $\lambda_2 < 0$.

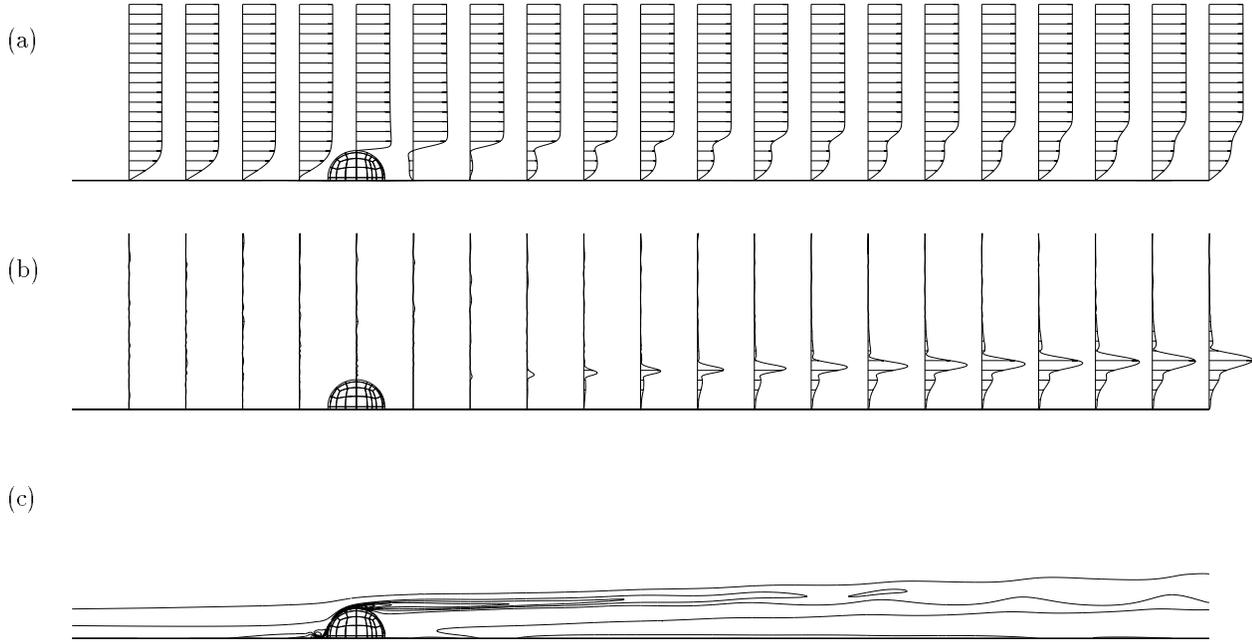


Figure 7: Symmetry plane data for $Re = 450$, $K = 1535$, $N = 9$: (a) mean velocity profiles, $\langle u \rangle$; (b) rms velocity $\langle u'^2 \rangle^{\frac{1}{2}}$; (c) spanwise vorticity, ω_y .

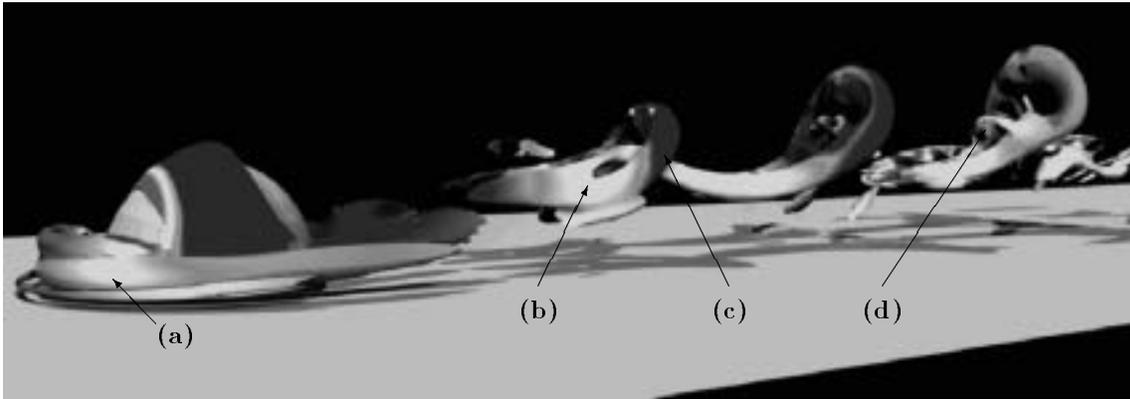


Figure 8: Key vortex structures at $Re = 700$: standing horseshoe vortex (a), interlaced tails (b), hairpin head (c), and bridge (d). Contours mapped onto $\lambda_2 = -1$ surface represent pressure (light=high, dark=low). Shadows on the plate provide additional perspective information.

of vortex reconnection in viscous flows [11, 12, 17, 24]. Acalar and Smith [1] note that the bridge-head structure eventually separates from the hairpin and lifts off as a separate vortex ring. At this Reynolds number the ring is so quickly dissipated by viscosity that the liftoff is not pronounced.

Because the bridge is quite thin and rather short lived, significant resolution in space and time is required to see it. We carefully chose the frame in order to present the bridge here. However, initial detection of the bridge and other similar unanticipated structures requires observation of a sequence of images from many viewing angles. Another structure detected as a result of such interactive viewing is shown in Fig. 9. The “spikes” seen jutting from the interior of the hairpin loop are readily visible in most of the higher Reynolds

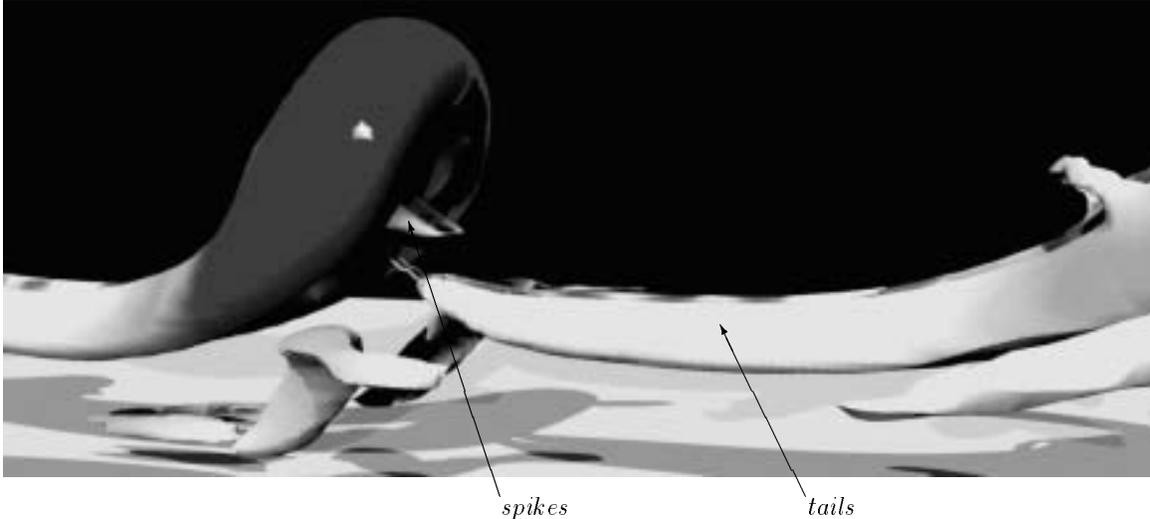


Figure 9: View showing distinction between “spikes” and tails of preceding vortex.

number computations. They appear at just about the time that the tails of the preceding vortex disappear as a result of stretching-induced dissipation. Initially, we believed the spikes were remnants of the dissipating tails. However, careful observation of the animation revealed that the spikes and tails briefly appear at the same time, indicating that they are not part of the same vortex structure. Downstream views of the head-spike structure seem to indicate that the spike formation is induced by the close proximity of the tails to the head, as seen in Fig. 8b.

A comparison between the hairpin vortices at $Re = 550$ and 700 is shown in Fig. 10. Few structural differences are seen over this range of Reynolds numbers. However, there is a noticeable difference in the distance that separates the successive hairpins, with shorter distances for $Re = 700$ corresponding to a higher Strouhal number (cf. Fig. 3). In addition, the $Re = 700$ case shows stronger evidence of tail-tail interlacing and the presence of a second horseshoe vortex upstream of the hemisphere.

We close this section with recently obtained results at $Re = 850$, shown in Fig. 11. The images reveal a cascade of vortices just a few diameters downstream of the hemisphere. At lower Reynolds numbers these vortices are not strong enough to yield the clearly defined structures that are present here. Clearly visible at this elevated Reynolds number are secondary vortices on either side of the hairpin, seen at the midpoint in the planform view. In addition, the liftoff of the hairpin head as a ring-vortex can be seen near the outflow in the profile view. The profile also shows the general inclination of the vortices, rising from the wall and moving downstream, commonly observed in turbulent boundary layers.

5 Discussion and Conclusion

We have examined the structure of hairpin vortices in the near wake region of a hemispherical roughness element at $Re_R = 700$ using spectral element simulations coupled with interactive immersive visualization. This is part of a wider investigation, with Reynolds numbers ranging from 450 (just above transition) to 1200, into the role hairpin vortices play in the transition between laminar and turbulent flow. We have shown excellent agreement between simulation and experimental data with regard to shedding frequency and clear identification of the principal features of the primary hairpin vortex and its evolution. We found that the combination of immersion and motion played a key role in the assimilation, integration, and understanding of this data set and believe that immersive visualization tools, such as those developed here, will be essen-

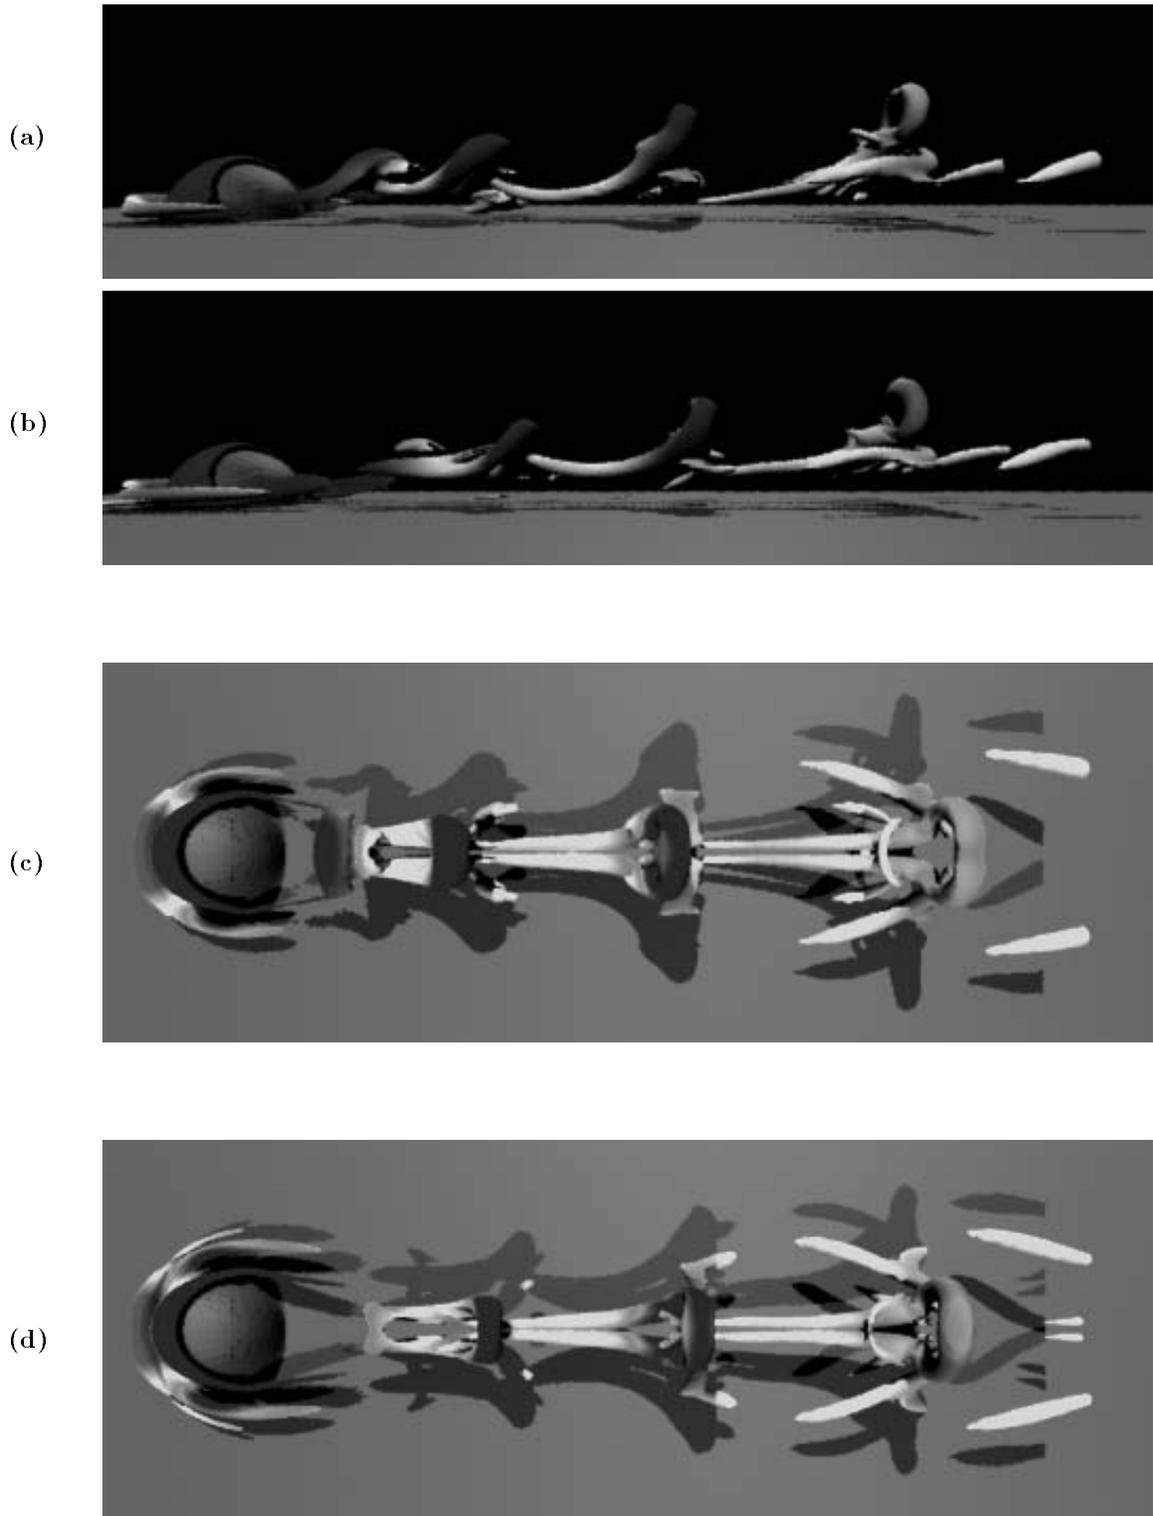


Figure 10: Comparison profile views of hairpin vortices at $Re = 550$ (a) and $Re = 700$ (b), and planform views at $Re = 550$ (c) and $Re = 700$ (d).

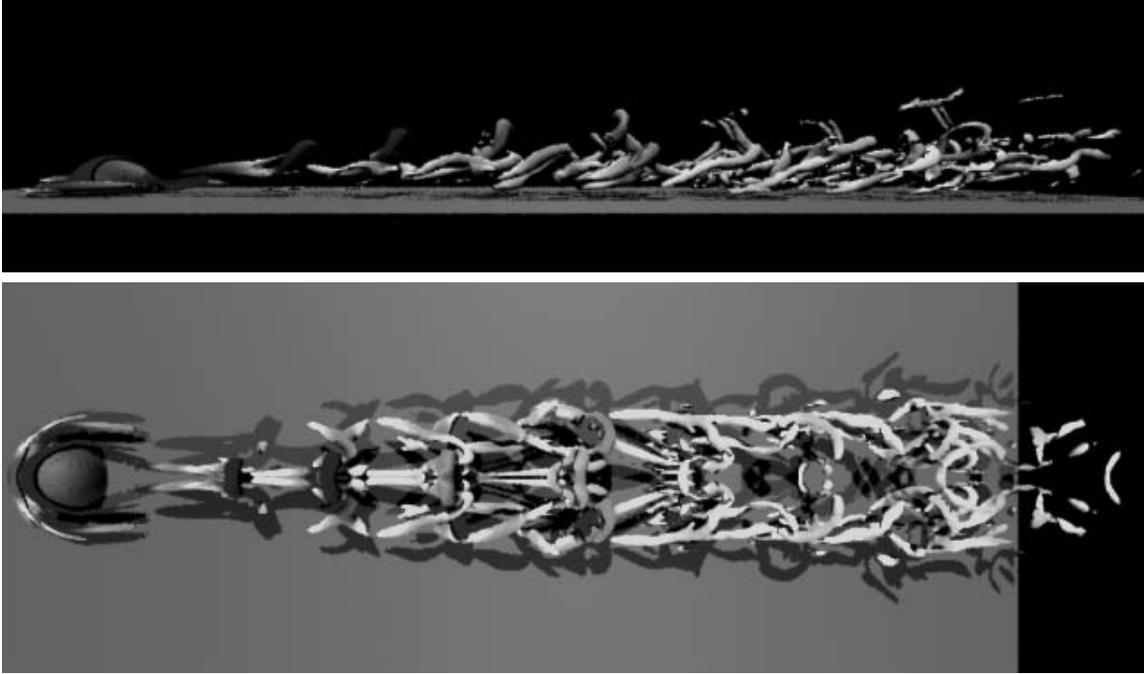


Figure 11: $Re = 850$ profile (top) and planform (bottom) views of $\lambda_2 = -1.1$ isosurfaces for the full computational domain.

tial tools for future investigation of flows where coherent structures play a key role in flow transition and development.

Further investigation is required into the nature of the secondary vortices that develop immediately downstream of the recirculation zone. Acalar and Smith observed two types of secondary vortices. The first formed in the wake of the hairpin and was entirely contained between the legs. It is an open question as to whether the bridge we observed corresponds to this structure and whether this is an example of classic vortex reconnection via bridging [12]. The second is a pair of vortices that form several diameters downstream of the hemisphere on either side of the primary hairpin vortex. We do indeed see the formation of such structures, which become more pronounced in recent calculations at higher Reynolds numbers. Singer and Joslin have shown that a cascade of such hairpins ultimately evolves into a turbulent spot [21]. Their calculations were in a plane channel, which allowed them to track the (isolated) vortex as it moved downstream, resulting in considerable computational savings. A similar approach could be used here, exploiting the temporal periodicity of the solution to provide a well-defined inflow boundary condition. It is clear that straightforward calculation of spatially developing flows from transition to turbulence in inhomogeneous geometries will require a significant increase in resolution to capture the range of scales present in both the solution and the geometry.

Exploration of extremely large data sets is difficult even with today’s state-of-the-art visualization tools. For example, to store the primitive variables (u, v, w, p) for a 150-frame “movie” requires 5–20 GB. Since we are interested in determining how vortex topology varies with Reynolds number, we anticipate having tens to hundreds of such “movies” in our database. There is a consequent need for advanced technology to archive, manipulate, and explore data sets in excess of a terabyte.

We are currently working on methods that facilitate such investigation. One promising path is to use subsampled or lower-resolution data sets for preliminary investigation and to record areas or paths where

high-resolution images are desired. These images can then be computed offline at full data set resolutions. We are developing this and other visualization software technology, such as automatic feature detection and tracking, to assist in the interrogation of large data sets such as considered here.

Acknowledgments

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, and by the Department of Energy under Grant No. B341495 to the Center on Astrophysical Thermonuclear Flashes at the University of Chicago.

References

- [1] M. S. Acalar and C. R. Smith, “A study of hairpin vortices in a laminar boundary layer: Part 1, hairpin vortices generated by a hemisphere protuberance”, *J. Fluid Mech.*, **175**, pp. 1–41 (1987).
- [2] D. C. Banks, T. Crockett, R. D. Joslin, B. A. Singer, “Parallel Rendering of Complex Vortical Flows”, <http://www.icas.edu/docs/hilites/banks/parallelRend.html>.
- [3] C. Cruz-Neira, D. J. Sandin and T. A. DeFanti, “Surround-screen projection-based virtual reality: The design and implementation of the CAVETM”, *Computer Graphics* (Proceedings of SIGGRAPH '93), ACM SIGGRAPH, August, pp. 135–142 (1993).
- [4] M. Dryja and O. B. Widlund, “An additive variant of the Schwarz alternating method for the case of many subregions”, *Technical Rep. 339*, Dept. Comp. Sci., Courant Inst., NYU (1987).
- [5] P. F. Fischer, “An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations”, *J. Comp. Phys.*, **133**, pp. 84–101 (1997).
- [6] P. F. Fischer, N. I. Miller, and H. M. Tufo, “An overlapping Schwarz method for spectral element simulation of three-dimensional incompressible flows”, in *Parallel Solution of Partial Differential Equations*, P. Bjorstad and M. Luskin, eds., Springer-Verlag, New York (1999).
- [7] Futures Lab Web page, www.mcs.anl.gov/fl.
- [8] L. Gritz and J. K. Hahn, “BMRT: A global illumination implementation of the RenderMan standard”, *J. Graphics Tools*, **1**(3), pp. 29–47 (1996).
- [9] M. T. HEATH, “The Hypercube: A Tutorial Overview”, in *Hypercube Multiprocessors 1986*, M. T. Heath, ed., SIAM, Philadelphia, pp. 7–10 (1986).
- [10] J. Jeong and F. Hussain, “On the identification of a vortex”, *J. Fluid Mech.*, **285**, pp. 69–94 (1995).
- [11] S. Kida, M. Takaoka, and F. Hussain, “Reconnection of two vortex rings”, *Phys. Fluids A*, **1**(4), pp. 630–632 (1989).
- [12] S. Kida and M. Takaoka, “Vortex reconnection”, *Annu. Rev. Fluid Mech.*, **26**, pp. 169–189 (1994).
- [13] P. S. Klebanoff, W. G. Cleveland, and K. D. Tidstrom, “On the evolution of a turbulent boundary layer induced by a three-dimensional roughness element”, *J. Fluid Mech.*, **92**, pp. 101–187 (1992).
- [14] Y. Maday and A. T. Patera, “Spectral element methods for the Navier-Stokes equations”, in *State of the Art Surveys in Computational Mechanics*, A. K. Noor, ed., ASME, New York, pp. 71–143 (1989).
- [15] Y. MADAY, A. T. PATERA, AND E. M. RØNQUIST, “An operator-integration-factor splitting method for time-dependent problems: application to incompressible fluid flow”, *J. Sci. Comput.*, **5**(4), pp. 263–292 (1990).
- [16] M. R. MALIK, T. A. ZANG, AND M. Y. HUSSAINI, “A spectral collocation method for the Navier-Stokes equations,” *J. Comput. Phys.*, **61** (1985) pp. 64–88.

- [17] M. V. Melander and F. Hussain, “Cross-linking of two antiparallel vortex tubes” *Phys. Fluids A*, **1**(4), pp. 633–636 (1989).
- [18] A. POTHEN, H. D. SIMON, AND K. P. LIOU, “Partitioning sparse matrices with eigenvectors of graphs”, *SIAM J. Matrix Anal. Appl.*, **11** 3 (1990) pp. 430-452.
- [19] W. Schroeder, K. Martin, and B. Lorensen, *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, Prentice Hall, Englewood Cliffs, N. J. (1998).
- [20] B. A. Singer “The formation and growth of a hairpin vortex”, *Instability, Transition, and Turbulence*, M. Y. Hussaini, A. Kumar, and C. L. Street, eds., Springer-Verlag, New York, pp.367 (1992).
- [21] B. A. Singer and R. D. Joslin, “Metamorphosis of a hairpin vortex into a young turbulent spot”, *Physics of Fluids*, **6** (11), pp. 3724–3736 (1994).
- [22] H. M. TUFO AND P. F. FISCHER, “Terascale spectral element algorithms and implementations”, SC’99, 1999.
- [23] H. M. TUFO, “Algorithms for Large-Scale Parallel Simulation of Unsteady Incompressible Flows in Three-Dimensional Complex Geometries”, Ph.D. Thesis, Brown University (1998).
- [24] N. J. Zabusky and M. V. Melander, “Three-dimensional vortex tube reconnection: Morphology for orthogonally-offset tubes”, *Physica D*, **37**, pp. 555–562 (1989).