
VAPOR Quick Start Guide

November, 2008

Web Links updated April 2013

Introduction

This guide is intended to provide the new user with a quick introduction to the use of VAPOR. Just work through the instructions below to get a basic understanding of the tool.

The VAPOR user interface is intended to enable scientists to readily visualize and analyze time-varying datasets that result from numerical simulation. These datasets can be extremely large: Often many terabytes of data result from months of supercomputer calculations. Simply reading and writing such data can take weeks, so visualization and analysis of this data would ordinarily involve enormous batch runs, with little hope for scientists to interactively work with their data. Using VAPOR, much of this work can now be performed interactively!

This guide provides a few examples to illustrate that such massive datasets do not need to be impenetrable. VAPOR permits users to interactively visualize large- and small-scale features of massive time varying datasets, without needing to process the entire dataset.

This guide is only an introduction to the basic features of VAPOR. Once you have worked through these examples, you will want to try VAPOR on your own data. Additional and advanced features are documented in the Users' Guide. Many of the capabilities of the user interface are described in "tool-tips"; i.e. when you hold the mouse for a second or two over a text input field or a slider, a description of the functionality will appear.

In this document there are frequent references to items in the VAPOR user interface. When a word or phrase in this text is literally a label in the user interface, that word or phrase is distinguished by underlining it.

Prior to working through this guide, you will need to download and install the example datasets from

<https://www.vapor.ucar.edu/page/vapor-download/#Example>.

There are both high-resolution and low resolution versions of these datasets. The guide assumes you are using the low-resolution data, however the instructions will not change much if you use the high-resolution data.

If you encounter problems following all the steps in this guide, you can quickly advance to any of the tasks below. For most of the illustrations in this guide, there is a VAPOR session file named "guideFigureN.vss", that can be used to put VAPOR into the same state as was used to generate Figure 'N' below. The session files are found in the quickstart examples directory, which is the subdirectory \$(VAPOR_SHARE)/examples/quickstart of the VAPOR install directory. They may also be downloaded from <https://www.vapor.ucar.edu/page/vapor-download/#Example>. To use such a session file, you need to do the following:

1. From the VAPOR file menu, click on Open Session and select the desired session file associated with one of the figures below. (Note: If you are on Windows, you may get a warning message that the log file cannot be opened. That's because the session file was created on a Unix platform where the path names do not have a drive letter. If you would like to have a log file, you can click on Edit→Edit Session Parameters, and specify a log file on your Windows system.)
2. From the VAPOR data menu, click on: Load Data into current Session and select the metadata (vdf file) associated with that figure. For examples in this guide, this is where you installed the example data you downloaded from the VAPOR website. (Note: In section 3 there are some sessions that require using two different datasets. To recover these sessions you need to first load the original dataset, and then click under the Data menu on: Merge (Import) a Dataset into Current Session.)
3. Click on each of the renderer tab(s) used in the illustration (DVR, Flow, Iso, or Probe). Under the heading Renderer Control, click the check-box for each renderer instance that is active.
4. If there is a mouse mode that is used in the illustration, (e.g. "Region Select", "Navigation", "Data Probe", or "Rake") you should select the appropriate mode by clicking on one of the mode select buttons below the VAPOR main menu bar.

You should see a similar image on your screen as is shown in the corresponding figure.

Prerequisites

Before you begin this guide, you need to have access to a (Linux, Windows, Mac or Irix) workstation on which VAPOR has been installed. Follow the installation instructions provided on the VAPOR website (<http://www.vapor.ucar.edu>). When you run VAPOR, you need to be directly connected to the graphics console of the system. You cannot run VAPOR via a remote shell. In some environments (check with your local administrator), VAPOR may be accessed remotely to a graphics system via Sun VirtualGL, SGI Vizserver or HP RGS.

You will also need to download and install example datasets from <https://www.vapor.ucar.edu/page/vapor-download/#Example>. At that site there are both high-resolution and low resolution versions of these datasets. The guide assumes you are using the low-resolution data; however the instructions do not change very much if you use the high-resolution data.

You may need to prepare your execution environment before running VAPOR. For Unix users, there are shell scripts in the bin/ subdirectory of the VAPOR installation directory. You will need to initialize your shell by sourcing vapor-setup.csh (in a c shell) or vapor-setup.sh (in a Bourne shell). Remember to re-source these whenever you change your shell. Windows and Mac users

need to run the VAPOR installers to set up the required paths. Macintosh users may need to run an X11 shell before launching the VAPOR user interface.

Section 1: Volume visualization

1.0 Launching VAPOR

From the command line issue the command “vaporgui”. (In Windows or Mac, you can double-click on the “vaporgui” icon on the desktop, or type “vaporgui” from a console window.) You should see a screen like the following:

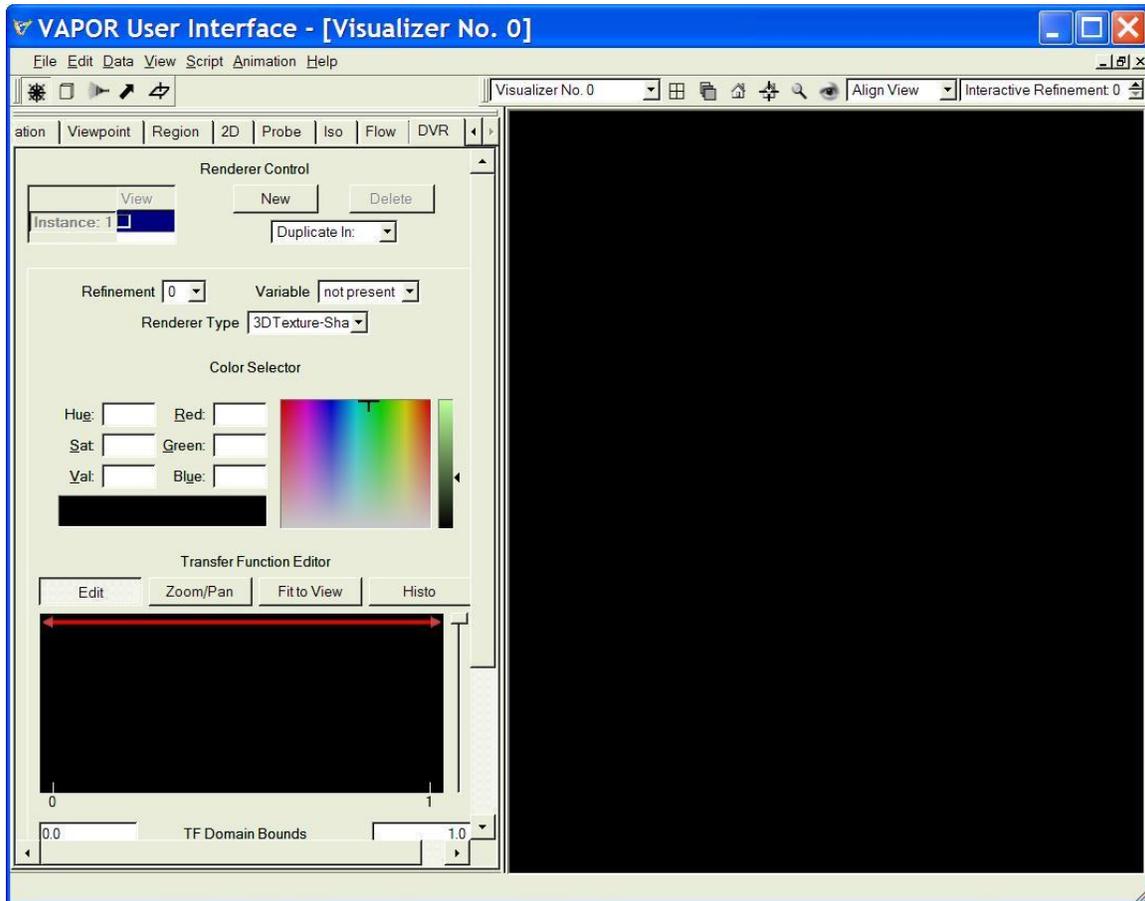


Figure 1.1: Initial VAPOR GUI window

The right side of this screen (mostly black) is where the 3D images are displayed. The left side of the screen has a set of tabs across the top, labeled Animation, Viewpoint, Region, 2D, Probe, Iso, Flow and DVR. We shall initially work with the DVR panel (this stands for “direct volume rendering”, referring to the fact that this panel controls the direct conversion of data to a volumetric representation). You may wish to resize the VAPOR window to make more room for visualizations.

1.1 Helpful tips: Tool-tips and context-sensitive help

You can quickly answer questions about things you see in the vapor user interface, using either “tool-tips” or “context-sensitive-help”. If you hold the mouse over a button, text-box or other widget for a few seconds, a *tool-tip* will appear, giving a one-sentence explanation of the widget. If you need a deeper technical explanation, click on the Help menu at the top of the screen, then click on the menu item labeled “? Explain This”. Then click on the item you are curious about. A yellow *context-sensitive help* box will appear such as shown below:

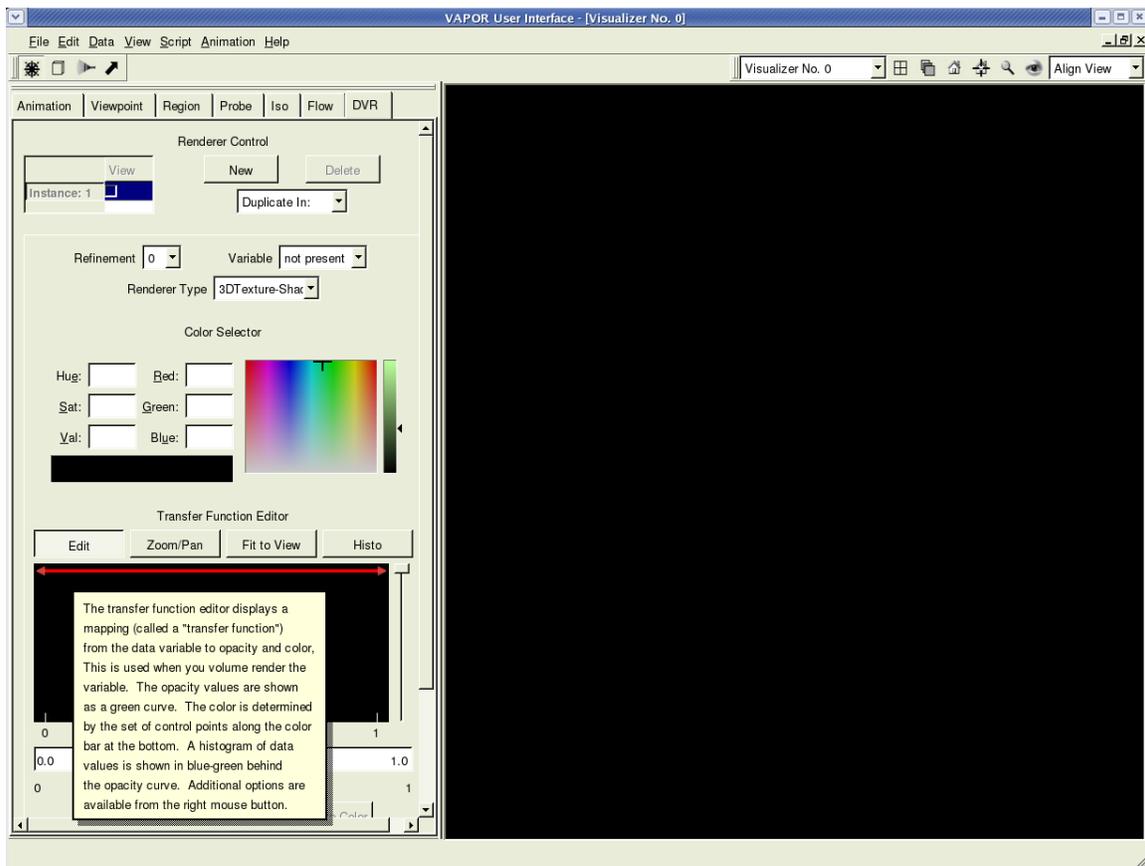


Figure 1.2: Context-Sensitive Help

1.2 VAPOR Renderers and Renderer Control

The DVR, Flow, Iso, 2D, and Probe tabs control the rendering of images in the visualizers. Each of these tabs has a Renderer Control section near the top, appearing as follows:

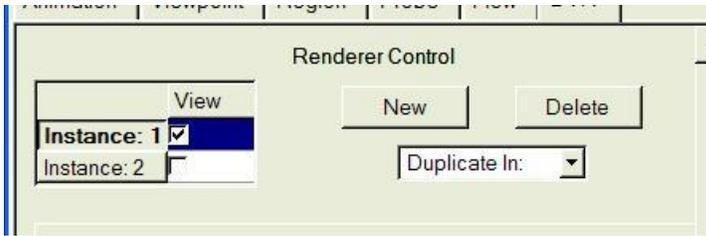


Figure 1.3: Renderer Control

Using the Renderer Control, you will be able to control any number of renderers in the current visualizer, but only after you have loaded a dataset. (We'll do this in the next paragraph.) The different renderers are identified as "Instance: 1", "Instance: 2", etc. When you click in the check-box, the specified renderer is *enabled*; i.e. the renderer will start drawing in the current visualizer. One of the renderers is *active*, in the sense that the parameters displayed in the tab apply to that renderer. You can change the active renderer by clicking over the instance name. Additional renderers can be created by using the "New" button or the "Duplicate In" selector. The "Duplicate In" selector constructs a duplicate of the renderer in the current visualizer, or in any other visualizer (when multiple visualizers are available). Multiple Flow, Iso, 2D, and Probe renderers can be enabled at the same time in the same visualizer. However, only one DVR renderer can be enabled at one time, because of limitations in hardware support for volume rendering.

1.3 Specifying a volume dataset

Click the Data menu (at the top of the screen) and you will see a list of various commands that can be applied to datasets. Click "Load a Dataset into current session". (You can alternatively load a dataset into default session, at this point). You will see a file selection dialog, asking you to "Choose the metadata file to load into current session":

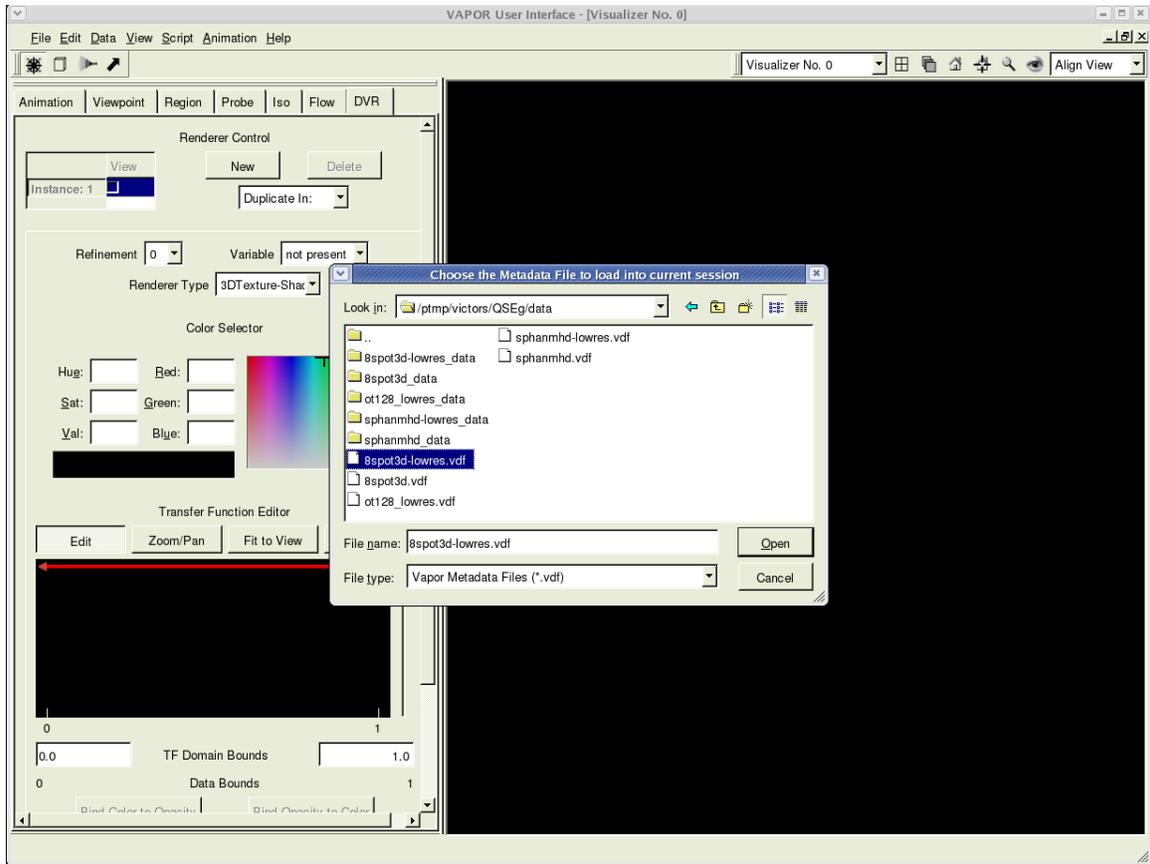


Figure 1.4: Loading a metadata file

This file selection dialog enables you to read a file of type “.vdf”, which stands for “Vapor Data Format”. These are *Metadata* files, not *data* files: .vdf files contain descriptions of data. Typically the referenced data is in many different files.

Choose the file “8spot3d-lowres.vdf” (or “8spot3d.vdf” if you have it) and click “Open”. This is a time-varying data set. It contains a vector field consisting of the three components (r_u , r_v , and r_w) of the cartesian momentum, resulting from simulation of flow in the solar interior, and the density variable, r_o . (This results from work by Mark Rast at NCAR and Univ. of Colorado).

Near the top-left of the DVR tab, under the heading “Renderer Control” there is a check-box, identified as “Instance 1:” Click in this check-box to enable a volume rendering of this data. The display will appear as follows:

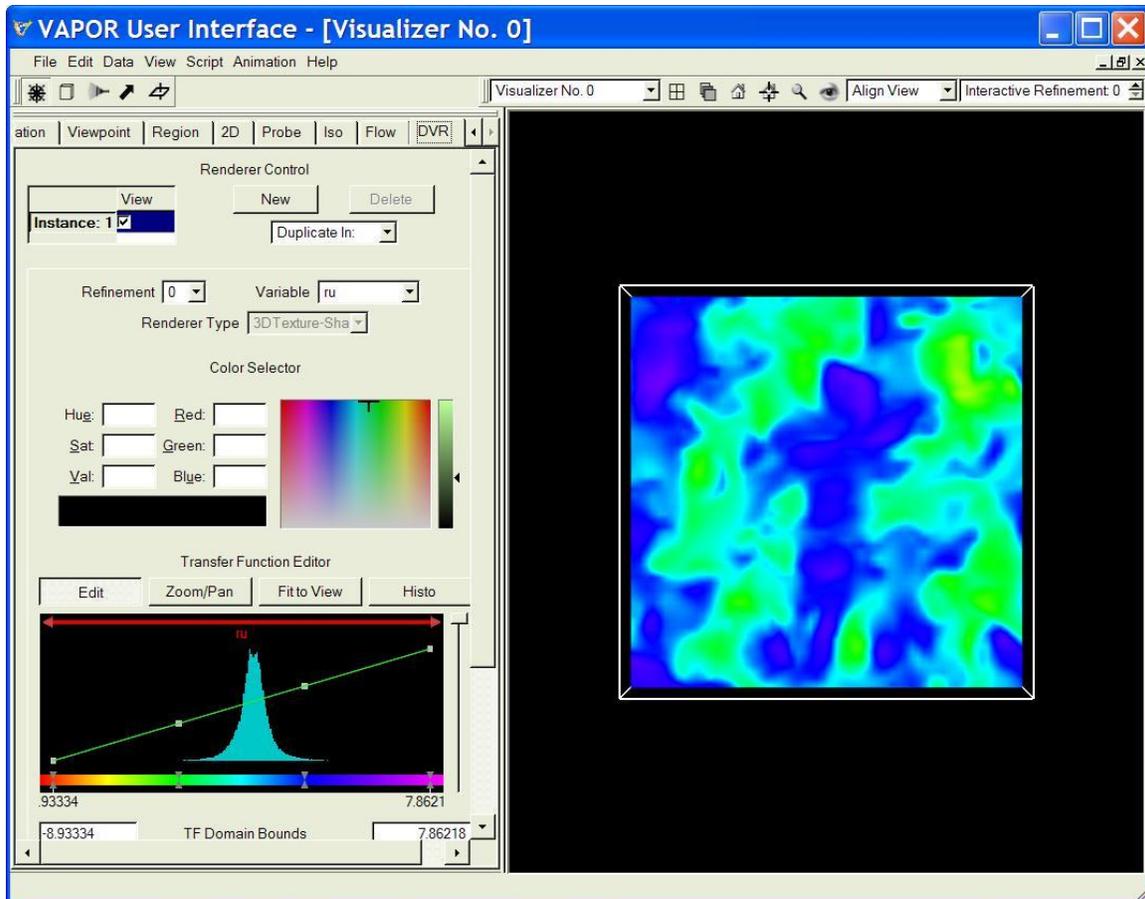


Figure 1.5: Volume rendering a variable

The image on the right is a 3D volume rendering of the variable 'ru' from the data set. Click and drag your left mouse in the right window and you can rotate the image. When you click and drag the right mouse button vertically, you zoom into or out of the scene. If you click and drag the middle mouse button, you translate the image vertically or horizontally. Return to the initial viewpoint by clicking the "home" button above the visualizer. Change the variable you are visualizing by selecting the variable name 'rw'

1.4 Designing a transfer function

The appearance (i.e. color and transparency) of the volume rendering is controlled by the Transfer Function Editor in the middle of the DVR panel, as shown below:

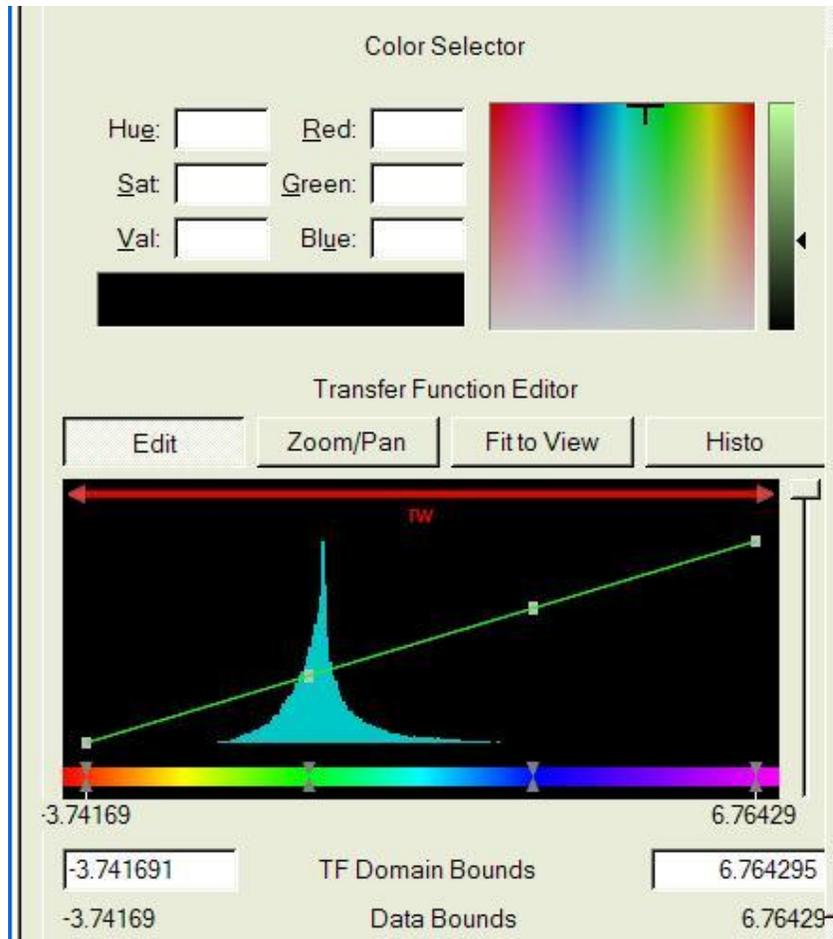


Figure 1.6: The VAPOR Transfer Function Editor

In the center of the transfer function editor window there is a blue-green histogram, indicating the values of the variable 'rw'. We shall use this histogram to construct a transfer function to visualize 'rw'.

The transfer function editor enables you to define a mapping from data values (in this case, values of the variable "rw") to both color and opacity. The color is controlled by a horizontal bar of colors at the bottom of the editor. Opacity is controlled by the diagonal green line in the center region. You initially will see four color control points (grey triangle pairs) on the color bar, and four opacity control points (grey squares) on the diagonal green line. Colors and opacity in the scene are interpolated between the control points.

First we make some changes in the opacity curve, to show some important features of the variable 'rw'. We will make the opacity curve in the transfer function editor look like the curve you see in Figure 1.7, below.

You move control points by clicking on them with the left mouse button and dragging them to their new locations. You create new control points by clicking with the right mouse at the desired location, and then selecting "New Opacity Control Point" from the menu of choices.

First drag the rightmost (fourth) opacity control point down to the bottom of the opacity graph. Drag the second and third opacity control points down to

about 1/10 of the distance from the bottom to top of the opacity window, on either side of the peak in the histogram. If you make mistakes, you can undo (or redo) your latest changes by selecting “Edit” from the main menu, and clicking to Undo (or Redo) the specified actions.

After you move these opacity control points, introduce two new opacity control points, slightly to the left and right of the peak in the histogram. This choice of opacity helps identify the data that is associated with either upward or downward flow, making the other data transparent. Your editor and your visualizer window should now have two bumps, one on either side of the histogram peak like the following:

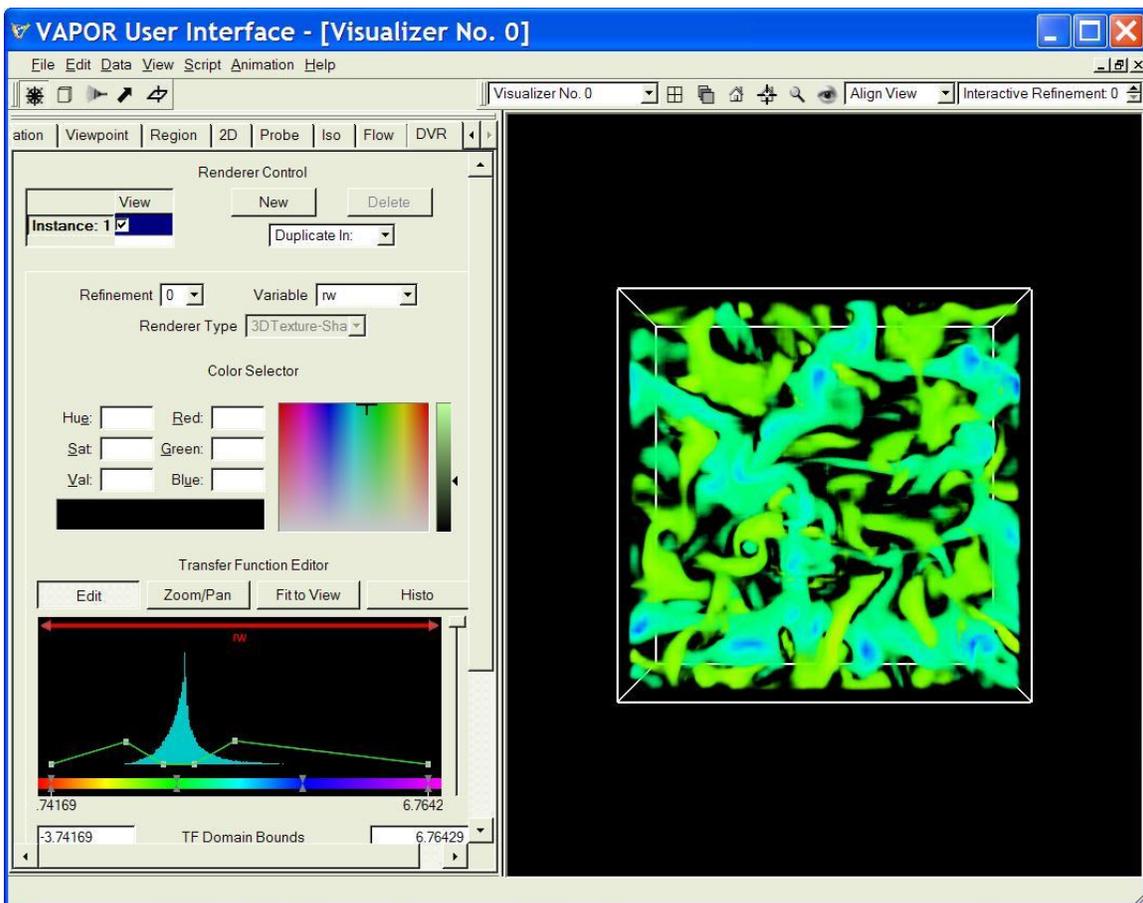


Figure 1.7: Using the transfer function opacity to emphasize positive and negative values

Now, introduce a new color control point by right-clicking slightly to the right of the histogram peak, and selecting “New Color Control Point”. This new color control point should be slightly to the right of the existing (second from left) color control point. Select the new color control point, by clicking on that point with the left mouse button. Then change the color by dragging the cross-hairs in the Color Selector to a sky-blue color. Then select the second color control point, and change the color of that color control point to orange. The resulting image should look like the following:

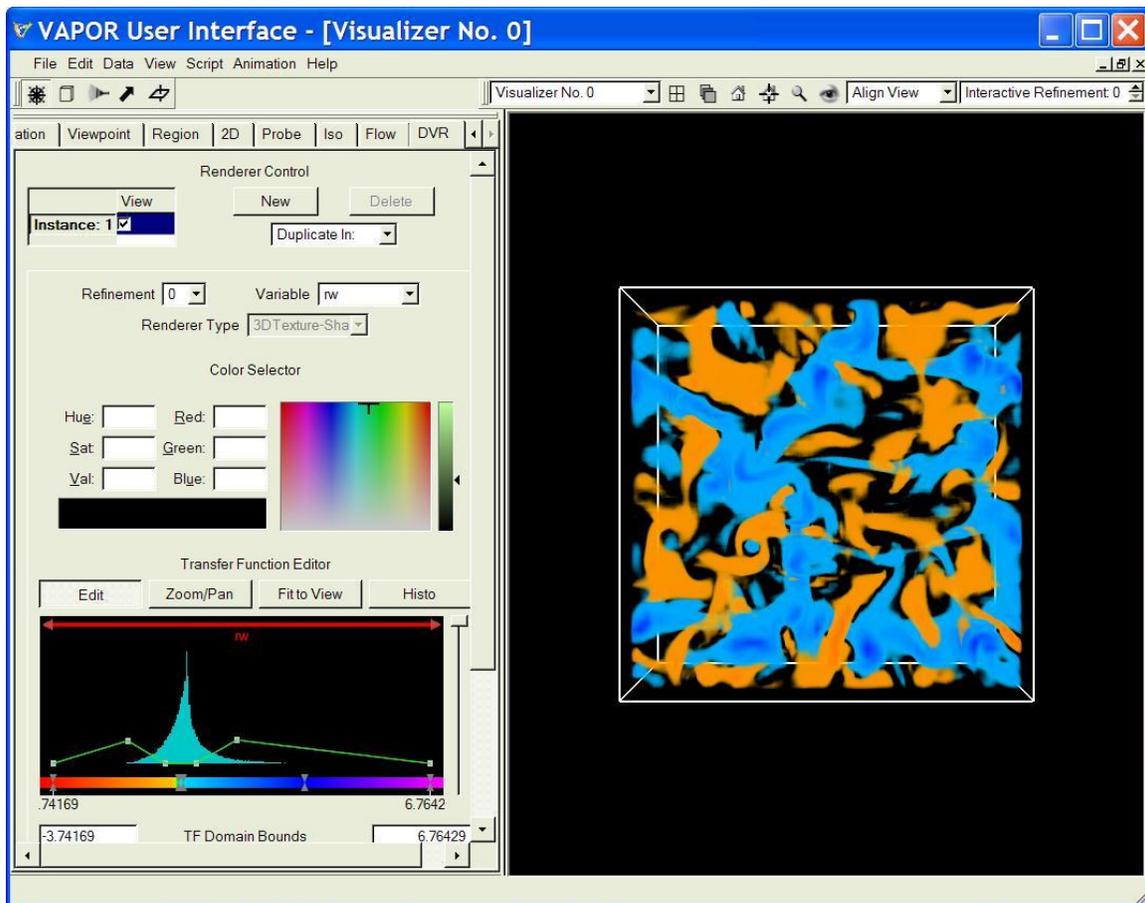


Figure 1.8: Editing the color map in the transfer function

With this choice of transfer function we have colored the vertical component of flow momenta so that upward (blue) and downward (orange) flow can be distinguished. Now is a good time to save the current setup. Click on File→. Save Session As, and specify the file “saved8spot.vss”. Later on, if you want to visualize this data again, you can use Open Session from the file menu to return VAPOR to this current state.

1.5 Animating the data.

This data set has 30 time steps. To control the time-stepping, click on the Animation tab, revealing a set of buttons similar to a VCR controller, as shown below:

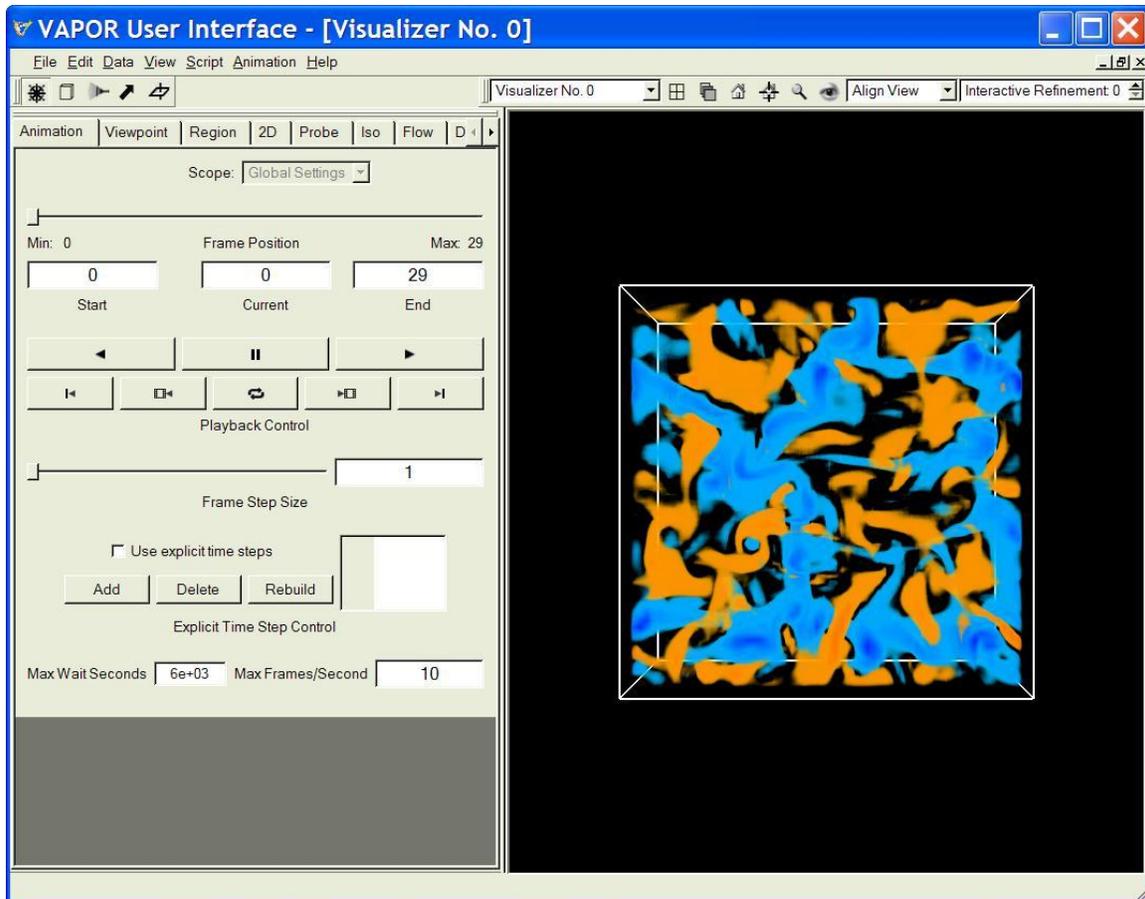


Figure 1.9: Animating the changing volume

Click on the “Play” button, labeled by a right-pointing triangle (▶). You will see the images slowly evolve from frame to frame, finally stopping at frame 29. You can click the left-pointing triangle to play backwards in time. You will note that the animation is faster the second time because the data is cached the first time it is read from disk. You can make it play slower or faster by setting the value of Max Frames/Second. Before going to the next section, reset the current frame position to “0” by sliding the frame position slider (at the top of the animation panel) to the far left.

1.6 Taking a closer look at the data.

So far we have been looking at this data at low-resolution (64x64x32) so that we can interact easily with the data. The data is available at the higher resolution 256x256x128 (and even at 512x512x256, if you have downloaded the high-resolution version of the sample datasets). If we were to load all the full resolution data into VAPOR, it would run much slower. However we can quickly examine small spatial regions of the higher-resolution data. We shall select a region containing a vortex and zoom into that region.

Return to the original viewpoint by clicking on the “Home” button above the visualizer. You will note that there is a swirl with a blue center to the left and

below the center of the visualizer. We shall isolate that swirl using interactive region selection. Click on the cube symbol at the top left of your screen, identified in its tooltip as Region Select Mode. You will see that the volume in the scene becomes enclosed by red lines, with cube-shaped handles on each side. These handles can be used to manipulate the bounds of the region that we visualize. If you ‘grab’ a handle by clicking on it with the left mouse, you can slide the whole region in the direction of the handle. If you ‘grab’ and slide with the right mouse, you will slide just the side of the region closest to the handle. You are initially prevented from sliding with the left mouse because the region is the entire domain and you cannot move the region bounds outside of the full data domain.

Click with the right mouse button on the handle at the top of the cube. The handle will change color to yellow. Drag it downward until the top of the region is just above the swirl. Then release the mouse. Similarly, drag the bottom, left, and right handles to narrow the region down to a small rectangular volume enclosing the vortex, as shown below:

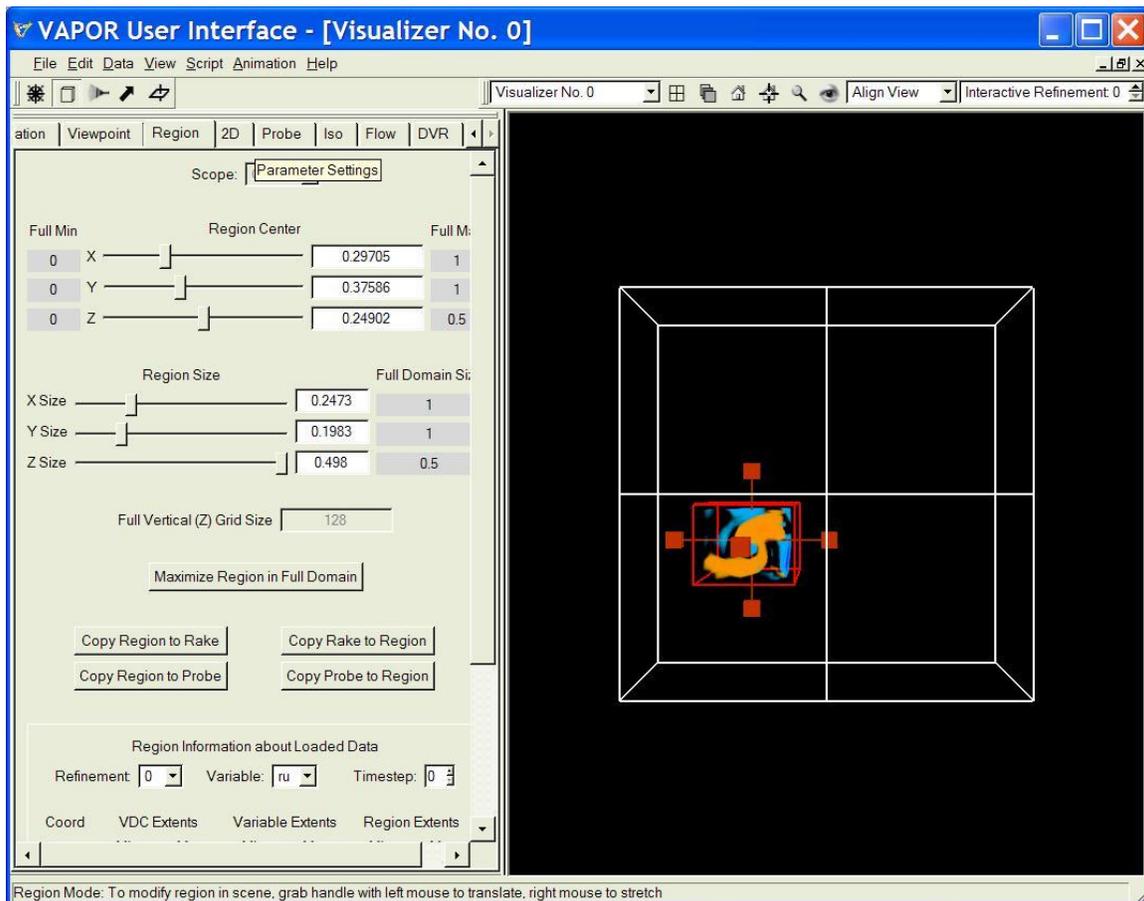


Figure 1.10: Using the region selection tool to specify a vortex

Return to “navigation” mode by clicking on the wheel button at the upper left of your screen. Zoom into this region by clicking on the magnifying-glass at the top right of your screen. Note that this vortex is quite fuzzy because the

resolution is so low. To increase the resolution, first click on the DVR tab. Near the top of the DVR panel, there is a Refinement selector currently set to “0”. The value of 0 is the base (lowest resolution) resolution of this data. Select refinement level “2”. It may take a few seconds to retrieve the higher resolution data, but then you will see a sharper image. You can zoom further into the volume by dragging the right mouse upward in the scene, obtaining a closer view of this vortex, like the following:

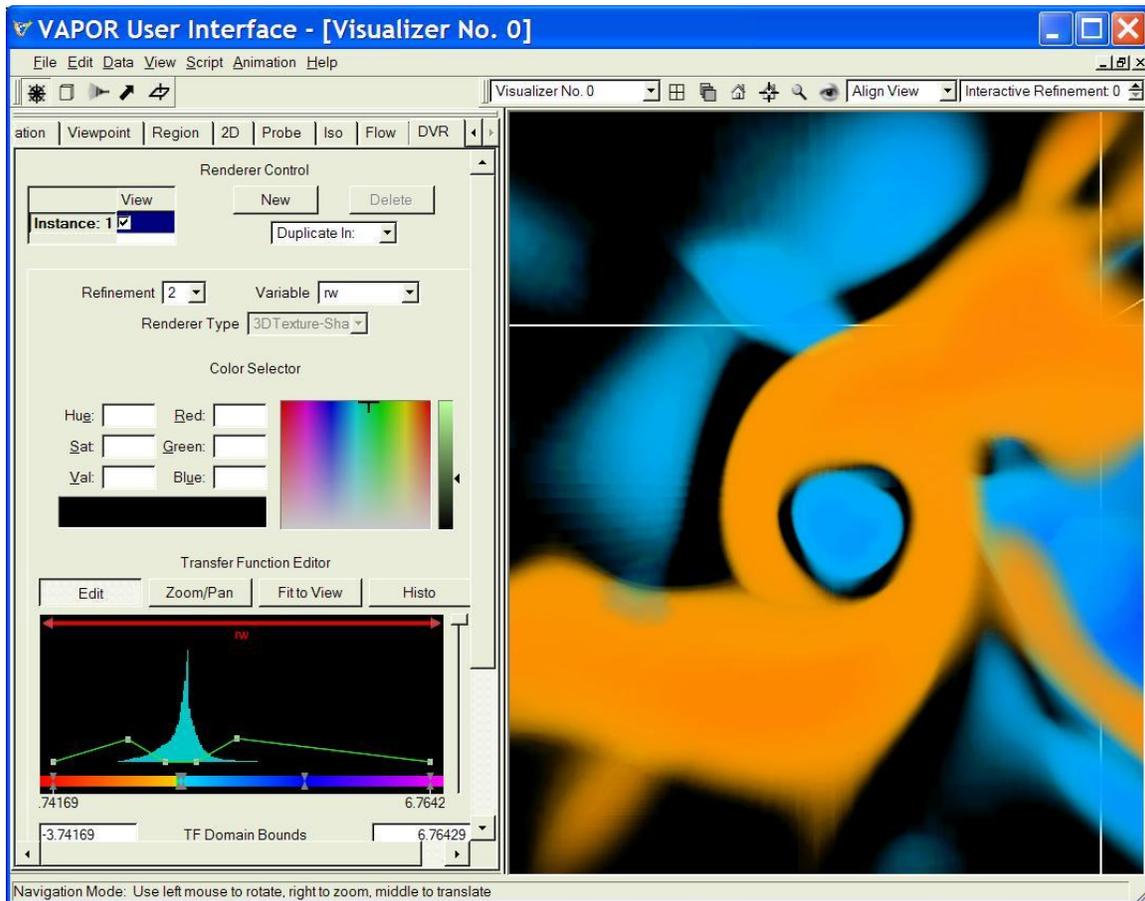


Figure 1.11: Zooming into the selected vortex

Save this session as “8spothighres.vss”. This session will be useful in the next section.

Section 2: Flow visualization

2.0 Getting started with flow integration

In this section we shall construct flow visualization based on the same dataset that we volume rendered in Section 1. VAPOR visualizes flow by performing a numerical integration along a vector field through a volume of data. The integration starts at user-specified starting points, called “seed points”, and tracks the successive positions in the volume that result from the integration. This process, of following points subjected to a vector field, is called “advection”. The successive advected positions can be visualized with various geometric entities.

We can take advantage of what we already know about the 8spot3d volume by opening the session file saved8spot.vss, saved in section 1.4. Click on **File**→**Open Session** and specify the saved8spot.vss file. Then, as in Section 1.2, load the 8spot3d-lowres.vdf data set (or the high-resolution version, if that is what you used) into this session. This time, do not enable the DVR visualization (yet). Instead, click on the **Flow** tab, and then enable flow rendering by clicking the checkbox labeled “**Instance: 1**” at the top of the flow panel. Resize VAPOR’s main window and use the right mouse button to zoom in, and you will see something like the following:

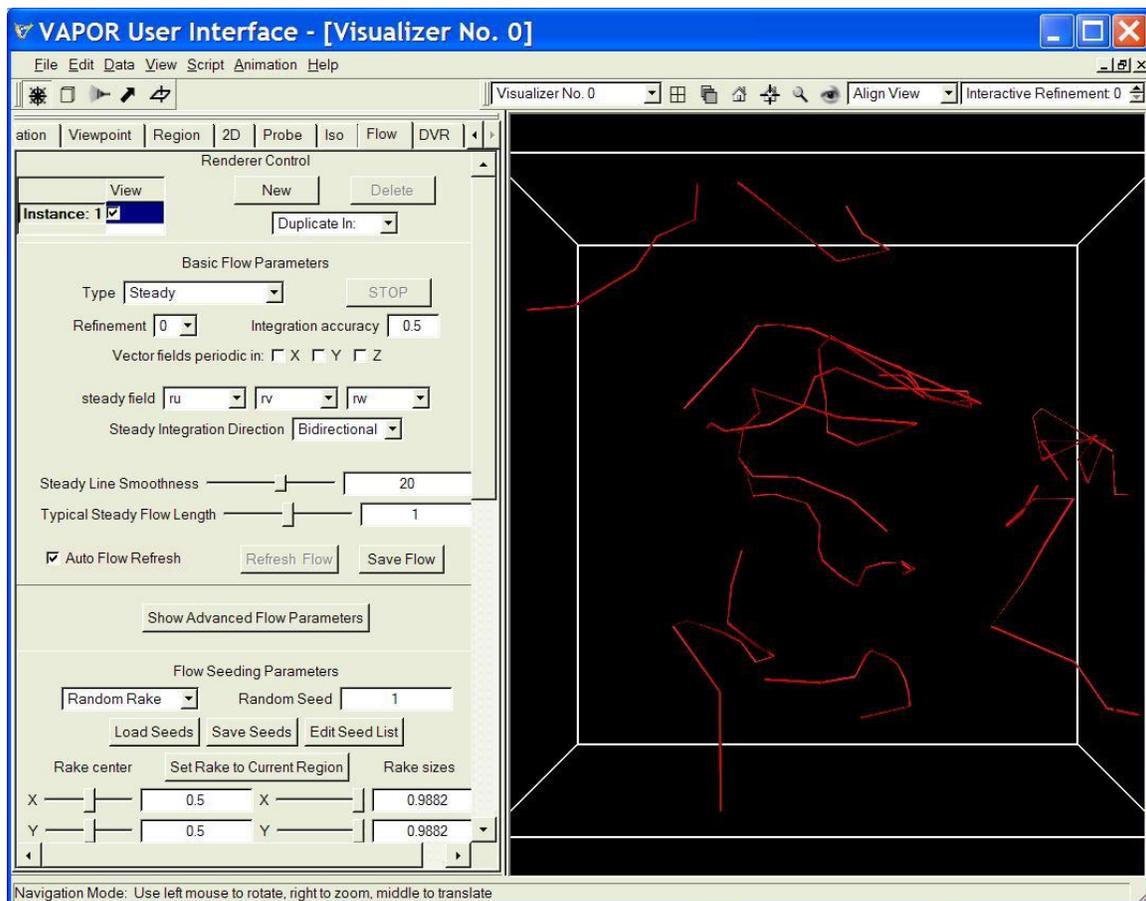


Figure 2.1: Random flow lines in a momentum field

This figure results from the default settings in the flow panel. The 10 steady flow lines result from integrating 10 randomly placed seed points, using the lowest resolution grid.

2.1 Initial Parameters for steady flow

Before we can make an interesting flow visualization, we need to establish some basic parameters. First, we must specify three variables that define the x, y, and z-components of the vector field that will be visualized. The current dataset happens to have four variables, r_0 , r_u , r_v , and r_w . The field we shall visualize has (x, y, z) components specified by (r_u, r_v, r_w) , so set the X, Y, and Z-coordinate selectors (labeled as steady field) to r_u , r_v , r_w (if they are not already set that way).

Second, we can choose how long or short we want the flow lines to be, and how smooth they should appear. Slide the Steady Line Smoothness slider to the right to a value of about 200, and set the Typical Steady Flow Length to about 2. These values are based on the average magnitude of the field, and will result in flow lines that are of length about 2 times the volume diameter, and having 400 sample points along their length.

Third, you need to select the flow type. By default, the flow type is *steady*. By steady we mean that the vector field does not change during the flow integration. The successive positions of a flow line in a steady flow depend only on the coordinates along the flow line, not on the elapsed time. You could also have selected the flow Type to be Unsteady or Field Line Advection, options that will be discussed later.

Fourth, one can select whether the flow integration goes forward, backward, or in both directions from the seed points. By default it is set to bidirectional.

You may also want to make the flow lines a bit thicker, or change their color. Scroll the flow panel down to the Color/Opacity mapping, and set the constant color of these lines by clicking on the “Select” button to the right of the words “Constant color, opacity values”. Select a purple shade. After making the above settings, your visualization will be similar to the following:

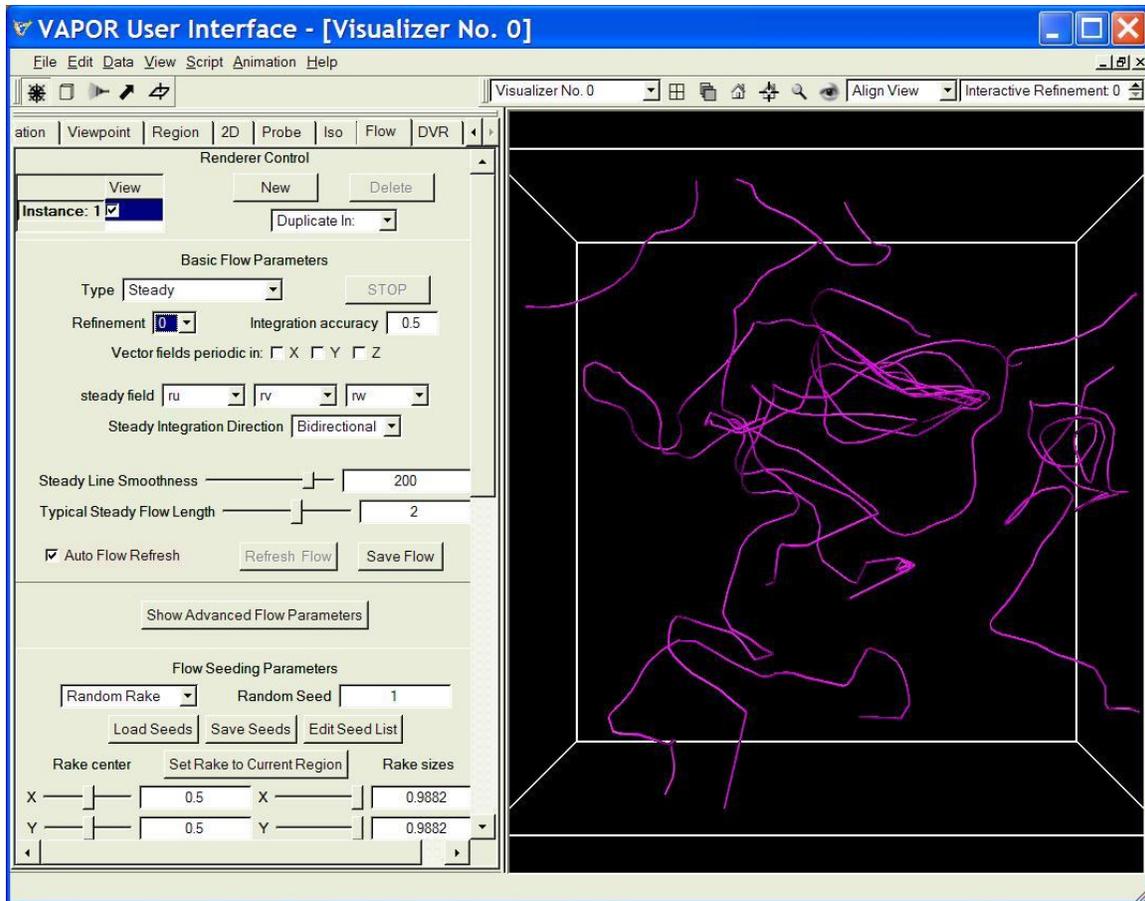


Figure 2.2: Setting the color and other properties of the flow visualization

2.2 Placing seeds with a rake

The problem with the above flow visualization is that there was no control over the positioning of flow seeds. To easily position lots of seed points, VAPOR provides a device called a “rake”, by analogy with rake-shaped devices that are used to inject smoke into wind tunnels. The default rake is random, providing a randomly generated set of 10 seed points, distributed over the full data volume. To manipulate the position and size of this rake, click on the rake icon near the top of your window. You will see red lines and red handles in the visualizer, similar to the region tool discussed in Section 1.6. You stretch or shrink the rake by dragging one of the red handles with the right mouse button. You translate the rake by dragging a handle with the left mouse button. You can also manipulate the rake by using the sliders and text boxes in the **Flow** panel. When the rake handles are released, the flow lines are recalculated based on the new rake. After you reduce the rake size and position the rake in the center of the volume you will see something like the following:

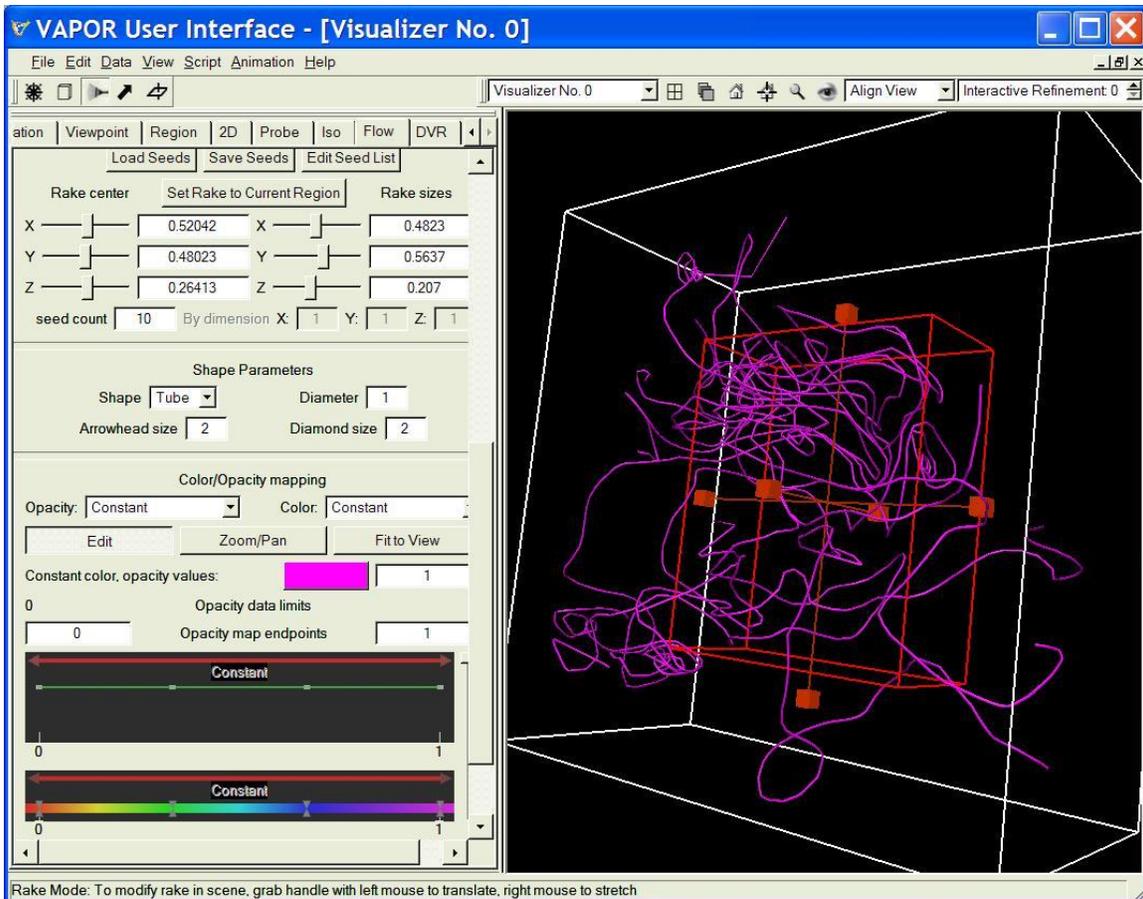


Figure 2.3: Using the rake tool to specify random seed points

We shall position the rake over the vortex that was explored in Section 1.6. To locate that vortex it will be useful to restore the session file where we found a high-resolution visualization of the vortex. From the main menu, click on **File**→**Open**, select and open the session file “8spothighres.vss”, and then load the “8spot3d-lowres.vdf” metadata file (or the high-resolution version, if that is what you have been using).

Click on the **DVR** tab, and, from the **Renderer Control** at the top of the **DVR** tab, click to enable **Instance: 1**. Now you will see the volume rendering of the vortex from section 1.6. We shall visualize the flow around this vortex. To do so, we position the flow rake over the region we have selected. Click on the **flow** tab and then click the button in the flow seeding parameters, entitled **Set Rake to Current Region**. Then, from the flow panel, apply the same flow settings that we set in section 2.1 (set the **Steady Line Smoothness** to 200; set the **Typical Steady Flow Length** to 2; set the **steady field** variables to ru,rv,rw; and set the constant color to purple). When you click to enable **Instance: 1** in the flow panel, you should see something like the following:

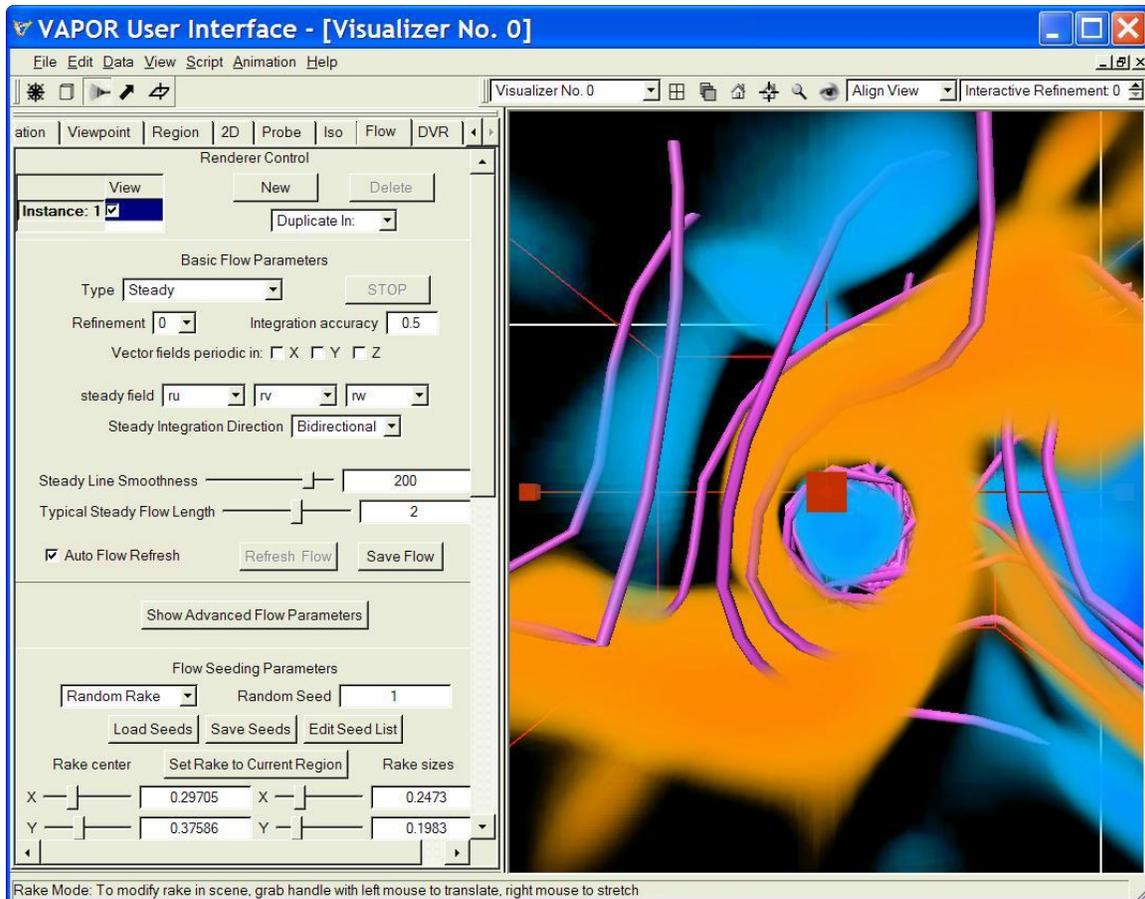


Figure 2.4: Inserting the rake into the vortex.

Your flow lines may look quite different than the lines illustrated above. The flow lines will vary quite a bit depending on the environment (e.g. random placement of flow seeds, data resolution, etc.)

2.3 Setting geometric parameters.

We are beginning to see the twisting of a vortex in the flow, but it is necessary to tweak the geometric parameters to really visualize what is going on. One can use color to show the direction of flow in the vortex. Within the Flow panel, scroll down to the section entitled "Color/Opacity Mapping". From the Color: selector, select Position along Flow. This causes the flow lines to be colored by mapping the length along the flow line to values between 0 and 1. The default color-mapping varies the color from red at the start, and the color changes to green, blue, and violet as the flow advances to its end. In this case many of the flow lines will leave the current region (in either direction) so colors near the start (red) and near the end (violet) may not show up. Make the field lines shorter, using the Typical Steady Flow Length slider, until the full range of colors is used.

The flow lines are somewhat obscured by the volume rendering in the visualizer. Click on the DVR tab. Note the vertical slider to the right of the Transfer Function Editor. This slider scales the opacity, making the volume more

or less transparent. Slide it about half-way down, so you can see the flow lines better.

To make sure your rake is concentrated on the most interesting part of the vortex, you can narrow it in the X- and Y- directions. Click on the **Flow** tab, scroll down to the **Flow Seeding Parameters**, and set the X-size and Y-size of the rake to 0.05 (by typing in the boxes labeled **Rake Size**, and then pressing enter). To see more lines in the region, set the **seed count** to 100, and under **Shape Parameters**, set the **Diameter** to 0.5.

Errors in flow integration can sometimes cause perfectly straight flow lines to appear. To obtain a more accurate visualization of these flows, increase the (integration) **Accuracy** (in the section of the Flow tab entitled “**Basic Flow Parameters**”) to about 0.9, or until the straight flow lines take on a more realistic shape. Increasing the refinement level under the Flow tab can have the same effect. Take care not to increase the **Accuracy** too much, though, as doing so can slow down visualization (especially during animation). Attempting to animate flows with a high refinement level can also be time-consuming.

This vortex can be visualized better from a side view. In navigation mode (click on the wheel near the top left of your screen), drag the left mouse upward in your screen to make the vortex point up and down. Drag downward with the middle-mouse button, to center the vortex in your visualizer. You will see something like the following:

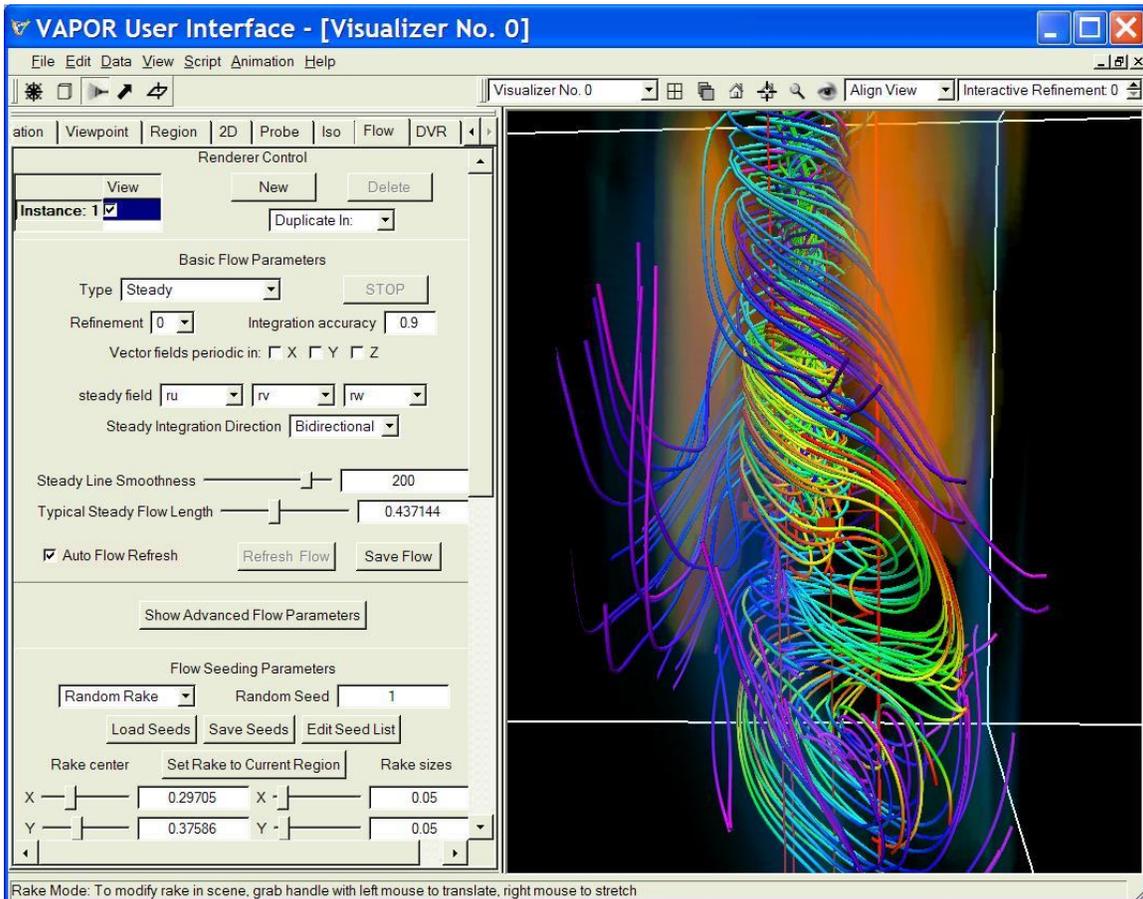


Figure 2.5: Flow in the vortex

Note that now you can see that the flow pulls the points upward (in the +Z direction) in the center of the vortex, while points on the outside of the vortex rotate outward and downward.

You may find it interesting to explore this vortex further. For example, in rake mode, one can position the flow rake in the interior of this vortex, or add more seeds to the rake. From the [Animation](#) panel you can see how this flow evolves over time. Note: Before you click “play” in the animation panel, you may need to lower the [Refinement](#) levels in both the [Flow](#) panel and in the [DVR](#) panel to 0. This example data set has only the lowest resolution data (refinement 0) for timesteps larger than 0. In general you can find out what refinement levels are available by checking the information at the bottom of the [Region](#) panel. The first time you try to render a refinement level that is not available, you will get a popup warning message.

2.4 Unsteady flow visualization

The steady flow visualization above captures the flow at various instants over time, but it does not show how particles in the flow would actually move as the flow changes. With unsteady flow, we integrate over changing time steps, simulating the result of releasing a particle in the evolving flow.

We begin again with the 8spot3d-lowres (or 8spot3d) data set. Load the session file “saved8spot.vss” and load the 8spot3d-lowres.vdf (or high-resolution, as appropriate) data into the current session. It will be helpful to see the coordinate axes in the scene, so, from the [Edit](#) menu, click [Edit Visualizer Features](#), and then click the checkbox [Show X,Y,Z \(R,G,B\) Arrows](#). Enable rendering in the DVR panel (by clicking on [Instance: 1](#)), and rotate the scene to give a side view (with +Z, the blue-colored axis, pointing up). The resulting scene appears approximately as follows:

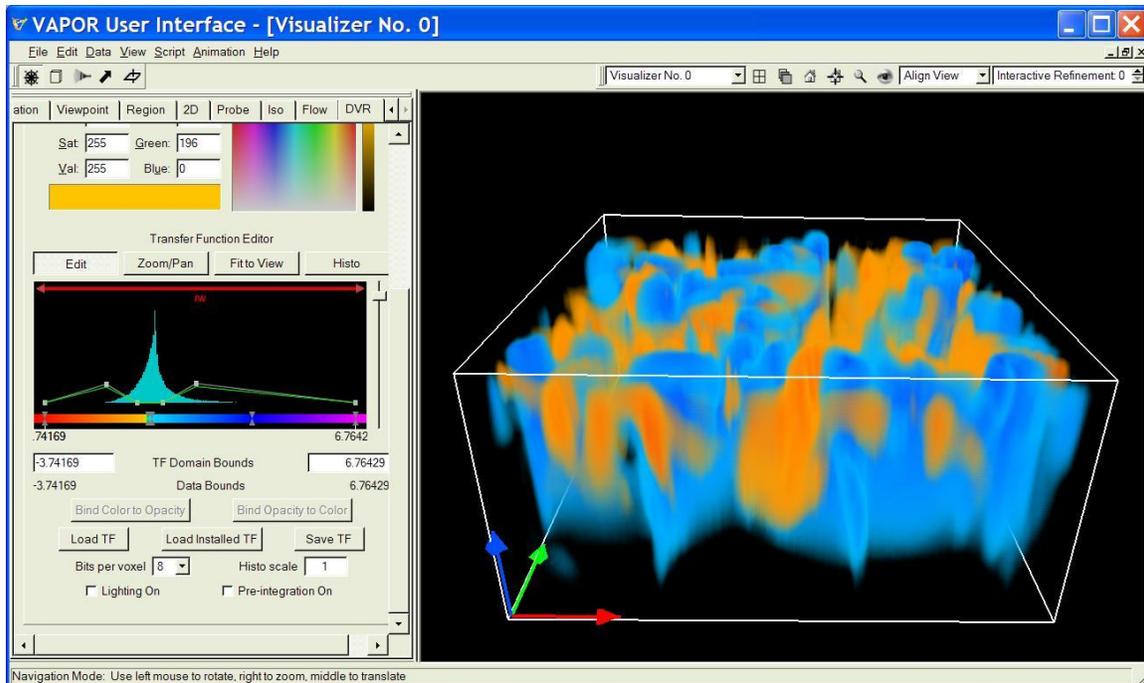


Figure 2.6: A side view of upward and downward flow

Note that we have stretched the screen horizontally and zoomed in to make a better image.

To set up the unsteady flow, click on the Flow tab. This time set the Flow Type selector to “Unsteady”. Set the unsteady field variables to (ru, rv, rw). In the Unsteady field scale factor enter the value 0.05. In general, to accurately represent an unsteady flow, you need to provide a field scale factor that matches the values from the simulation. A value of 0.05 corresponds to 1/20 seconds per time step. You can get a full explanation about setting this parameter by using the context-sensitive help on the Unsteady field scale factor text box.

Click the Instance: 1 check box at the top of the flow panel, to enable flow integration. You will have a few seconds to wait. Unsteady flow visualization can be slow because you need to access all the different time steps in the data whenever you make a change in the flow settings. To avoid these annoying delays, you may want to un-check the Auto Flow Refresh checkbox. With it un-checked, you can make several changes to the flow panel without delays. Later, when you are finished with your changes, you can click the Refresh Flow button to rebuild the flow lines. We shall make a number of changes to the flow settings to get a good unsteady flow visualization. With Auto Flow Refresh unchecked, when you make a change in the flow panel and press enter, the flow visualization will not update until you click Refresh Flow.

Instead of a random cube of seed points, this time we will use a regular 2D grid of seed points, positioned near the bottom of the volume. Scroll down to the Flow Seeding Parameters, select Nonrandom Rake (instead of Random Rake), and specify X,Y,Z dimensions of the rake (next to “By dimension”) as 6, 6, and 1. Click on the rake icon (above the flow tab), and then, dragging the top rake handle with the right mouse button, pull the top of the rake down nearly (but

not all the way) to the bottom of the scene. At the bottom of the Flow Seeding Parameters section of the Flow tab, set the Seed time start, end, increment values to 0, 0, 1, if it is not already set that way. This will cause the flow seeds to start the flow lines at only the first time step. After you refresh the flow, you will see the particles at the first time step. (It will probably help to lower the opacity slider in the DVR tab.) To see the paths followed by particles during the simulation, click on the Animation tab, and play to the end (time step 30), illustrating how these 36 particles move over time. Figure 2.7 illustrates the full paths traced by these particles.

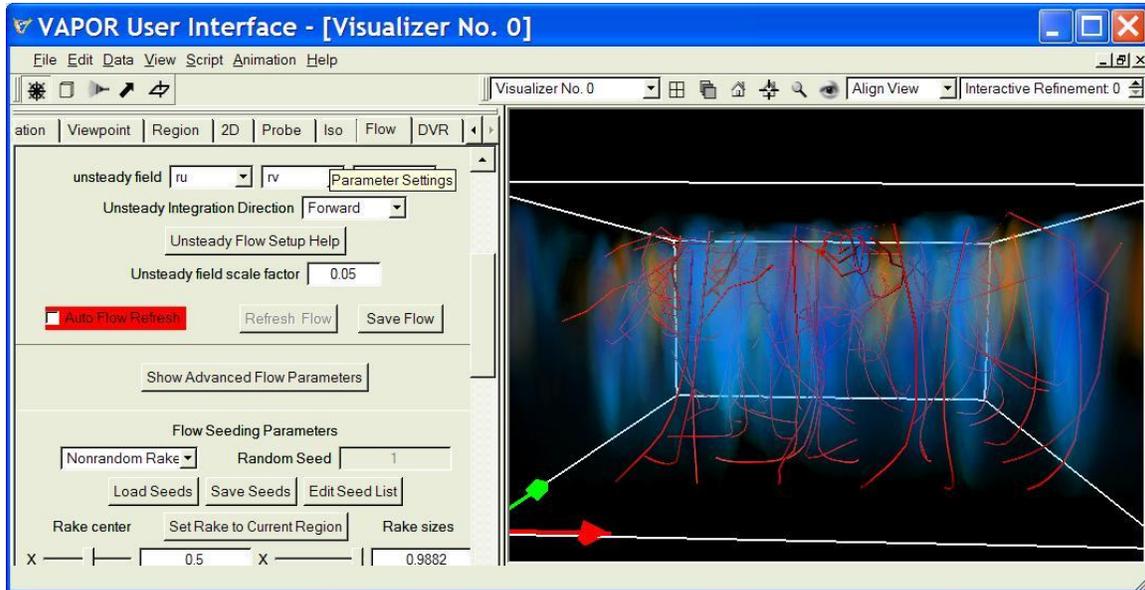


Figure 2.7: Unsteady flow lines

Next, we color-code the unsteady flow lines. Near the bottom of the Flow panel, click on the Color: selector, set it to Timestep. Near the bottom of the graphic parameters, there are (left and right) textboxes labeled Color map endpoints. Set them to 0 and 29, because these are the start and end time steps of this dataset.

Be sure to press the enter key after you make your editing changes, or they will not be saved. You will be able to see how the seed points start to move up from the bottom of the scene, as the color changes from red to green to blue and violet. Click the rake button to return to navigate mode and you will see the following:

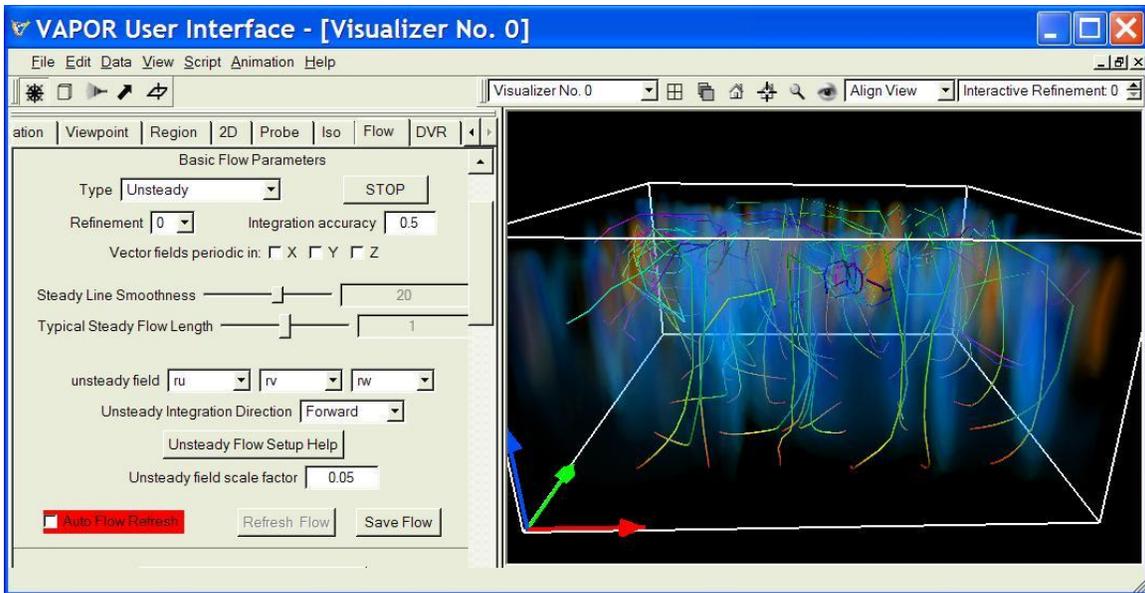


Figure 2.8: Color-mapped unsteady flow

We can animate this scene and see how the particles move over time. Note in the Shape Parameters (of the Flow tab) the items labeled timestep display interval min, max. We shall set these values so that, for each time step, you only see the portion of the flow near that time step. Set the min value to “-2” and the max value to “2”. With this setting, at each animation time step we shall just see the portion of a flow line from two time steps before, to two time steps after, the current time step. Set the Shape to Arrow, so that each arrow will represent the motion of a particle in one time step. Set the animation back to Frame Position 0, click Refresh Flow, and you will see a 6x6 array of red-orange arrows, starting near the bottom of your data cube, like the following:

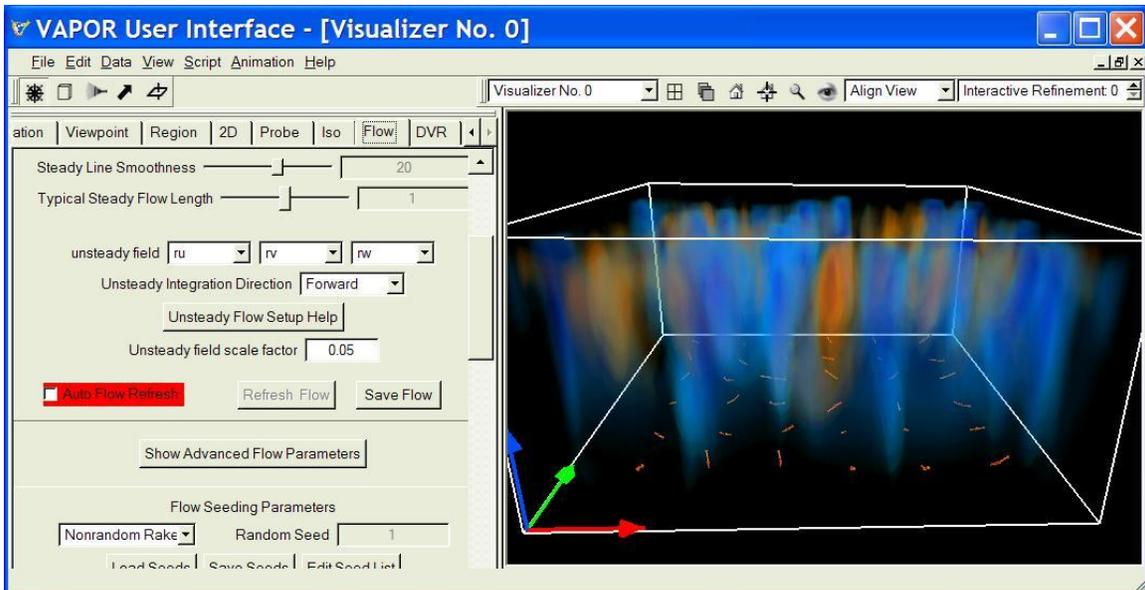


Figure 2.9: Using arrows to track particles in an unsteady flow

Now go to the Animation tab. Click on the “play” button and you will be able to follow the seed points as they are moved through the flow. You may want to slow down the Max Frames/Second to 1 second or so, so that you can easily track the motion.

Return to the Flow panel; we can add additional visual cues to this animation. Note that we are only seeding this flow at the very first time-step. We can inject new particles into the flow as it proceeds. At the bottom of the flow seeding parameters are three text boxes labeled Seed time start, end, increment. Set these values to 0, 28, and 4. This will result in additional flow lines starting every four time steps for the duration of the simulation. After you click Refresh Flow, return to the Animation panel and click the play button. Now you will see arrows starting near the bottom of the volume, with a new particle starting every four time steps. You can rotate or zoom into the scene to get a better understanding of the particle motion. The image below shows this animation at time step 8, just as the third particle injection is starting:

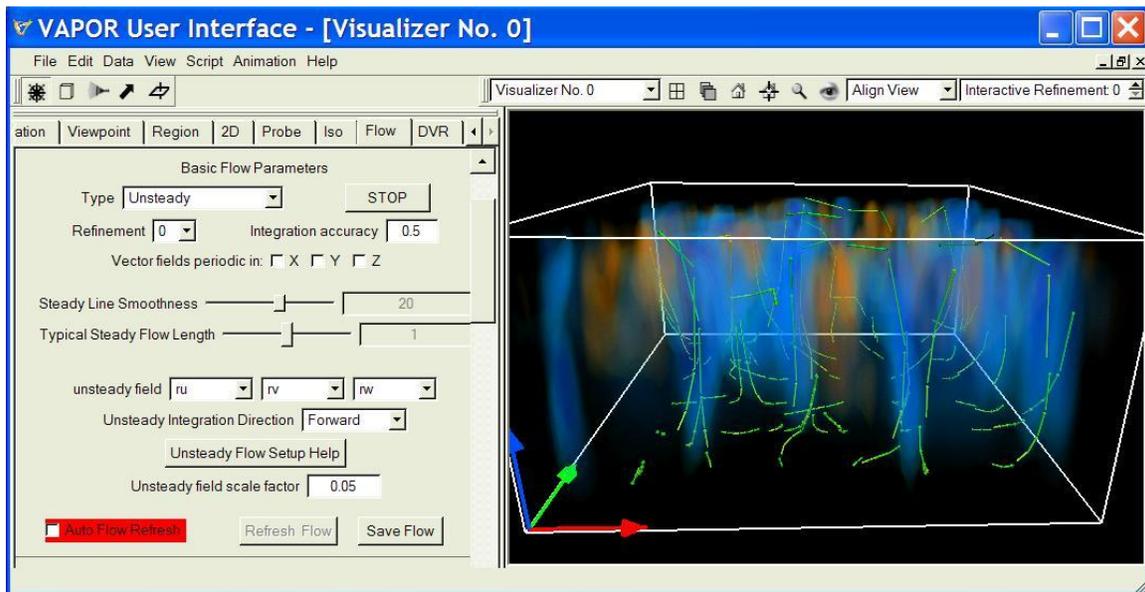


Figure 2.10: One time-step of an animating steady flow, using arrows.

If any of your adjustments to unsteady flows result in unreasonably long wait times, you may press the STOP button located in the “Basic Flow Parameters” section of the Flow tab to abort the calculation.

Section 3: Flow seed placement

So far we have been investigating the vector field using a rake to position seed points. Sometimes you will need better control over seed placement. VAPOR provides a couple of ways to place seeds more accurately. You can bias the distribution of randomly placed seeds, to obtain seeds that have larger or smaller field intensity. VAPOR also provides another tool, a “Data Probe” that allows you to specify exactly where a point is situated, so you can place the flow seeds individually based on field values at their position.

3.0: Visualizing a magnetic field

We shall illustrate this with another dataset. Start up vaporgui, and, from the Data menu select Load a data set into default session and then specify the metadata file “sphanmhd-lowres.vdf”, or the high-resolution version if you prefer. From the Edit→Edit Visualizer Features panel, enable the coordinate axes by checking the box labelled “Show X,Y,Z (R,G,B) Arrows”. From the DVR panel, set the Variable to “bc” and, enable the DVR rendering by clicking on Instance: 1. Zooming in some, you will see the following visualization:

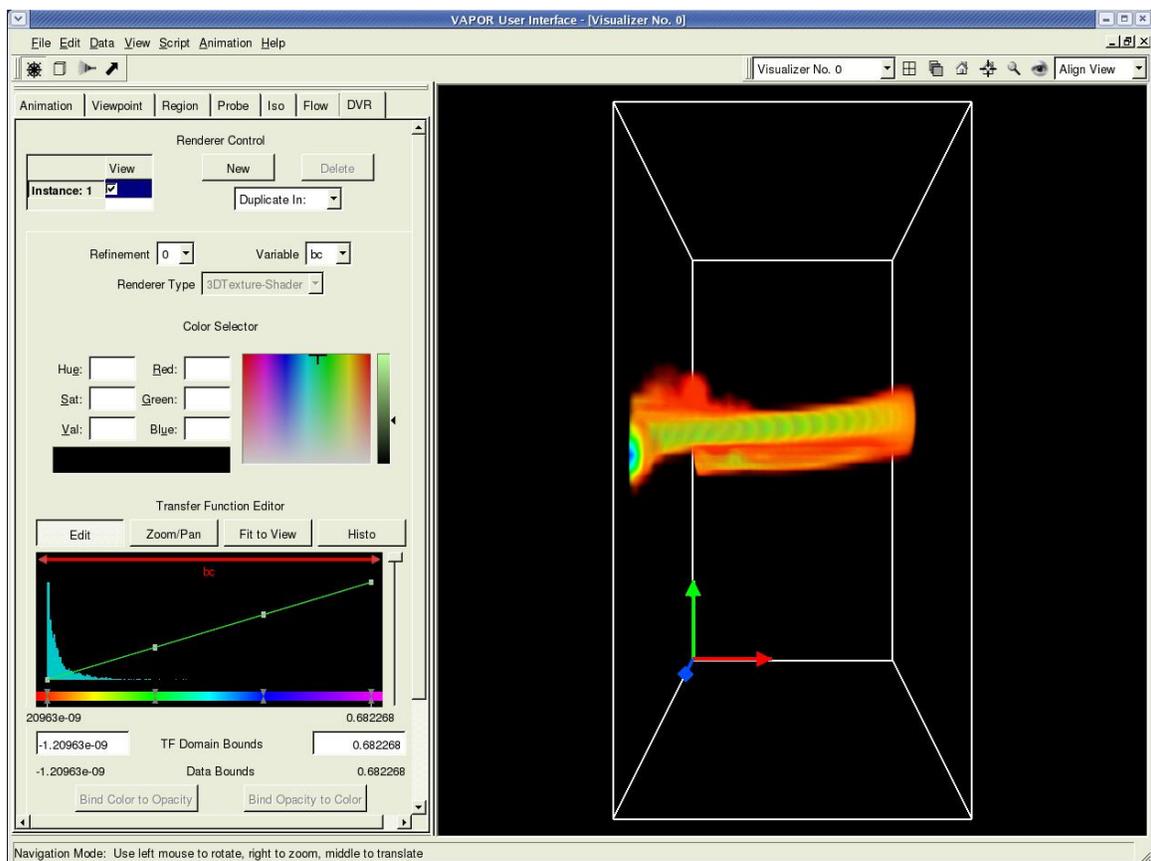


Figure 3.1: Volume visualization of the absolute field strength (bc) of a rising magnetic flux tube

This field is a magnetic field with components b_{1c} , b_{2c} , and b_{3c} . It results from a simulation (by Y. Fan of HAO/NCAR) of the buoyant rise of a twisted magnetic flux tube through the solar convective envelope towards the solar surface. The volume rendering shown in the domain corresponds to the absolute magnetic field strength (bc) of the rising flux tube.

In the **Flow** panel, set (b_{1c} , b_{2c} , b_{3c}) to be the x, y, and z-mapped steady field variables. Enable flow rendering by selecting **Instance: 1** at the top of the **Flow** tab. You will probably not see much of interest except for a few grey diamonds in the scene. That is because all of the interesting magnetic field values are concentrated along the tubular structure in the center of your visualization. Drag the “Steady Line Smoothness” slider to about 300, so the flow lines will be fairly smooth. Set the “Seed Count” (in the Flow Seeding Parameters) to 100 so we will see lots of flow lines. Using the rake tool, you can position the rake closer to the flux tube. (In the figure below, we have made the y-size of the rake equal to 0.175.) Dragging upward in the scene with the left mouse button, rotate the coordinate axes to the lower-left front corner and zoom in. You will obtain a side-view of the flux tube inside the rake, as follows:

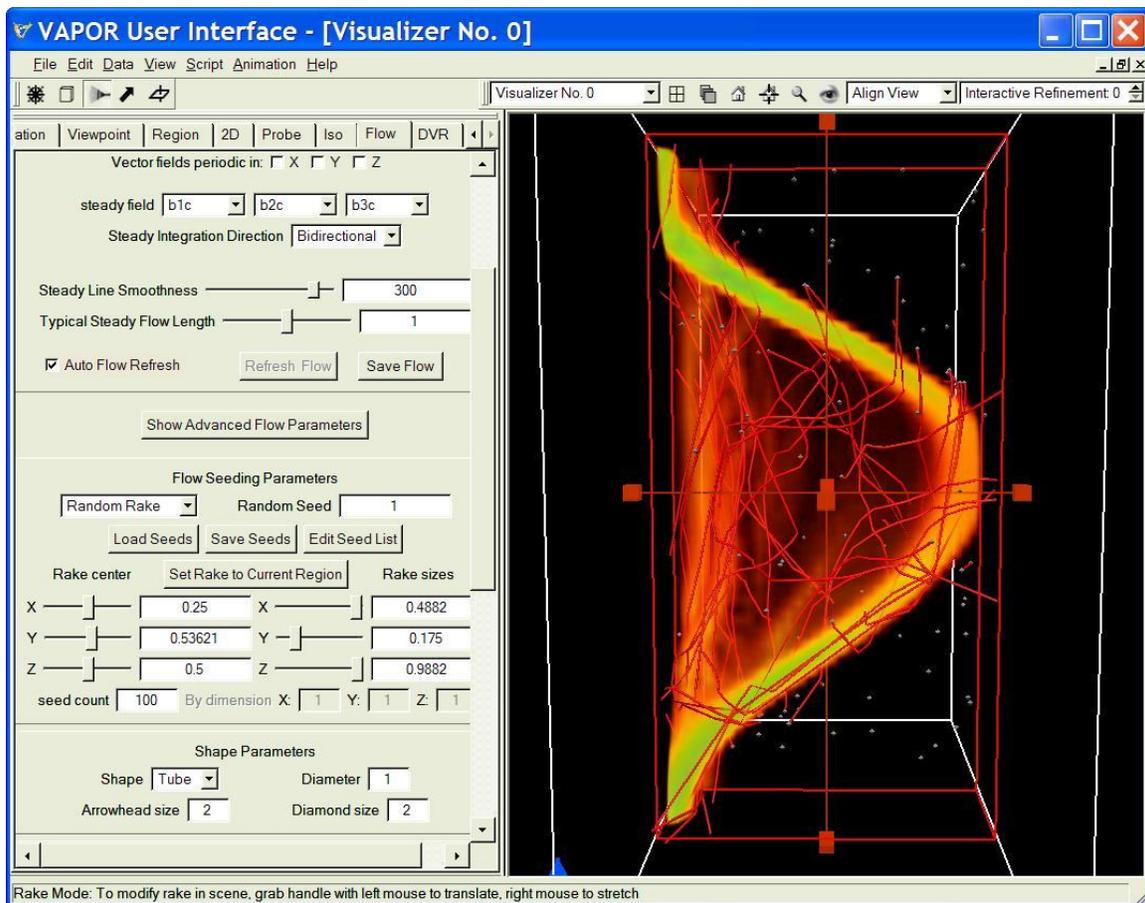


Figure 3.2: Using the rake to place seed points near a flux tube.

3.1 Biasing random seed distribution

Still, we have not managed to position many interesting flow lines. However there are a couple of advanced options that will help here. Click on Show Advanced Flow Parameters. Uncheck Automatically scale steady field, since we need manual control over field scaling in this example. Set the Steady field scale factor to 100 and set the Steady samples per flow line to 100. Also, set the Seed Distrib. Field variables to b1c, b2c, and b3c if they are not set that way already. The seed distribution field can be used to ensure that most of our seed points are placed in positions with higher field magnitude. Set the Seed Distribution Bias to a positive value (e.g. 7). Now the random seeds will be positioned near maximal values of the B-field. Lower the opacity scale in the DVR tab to about half-way up, and you will see the field lines are mostly in the flux tube, as follows:

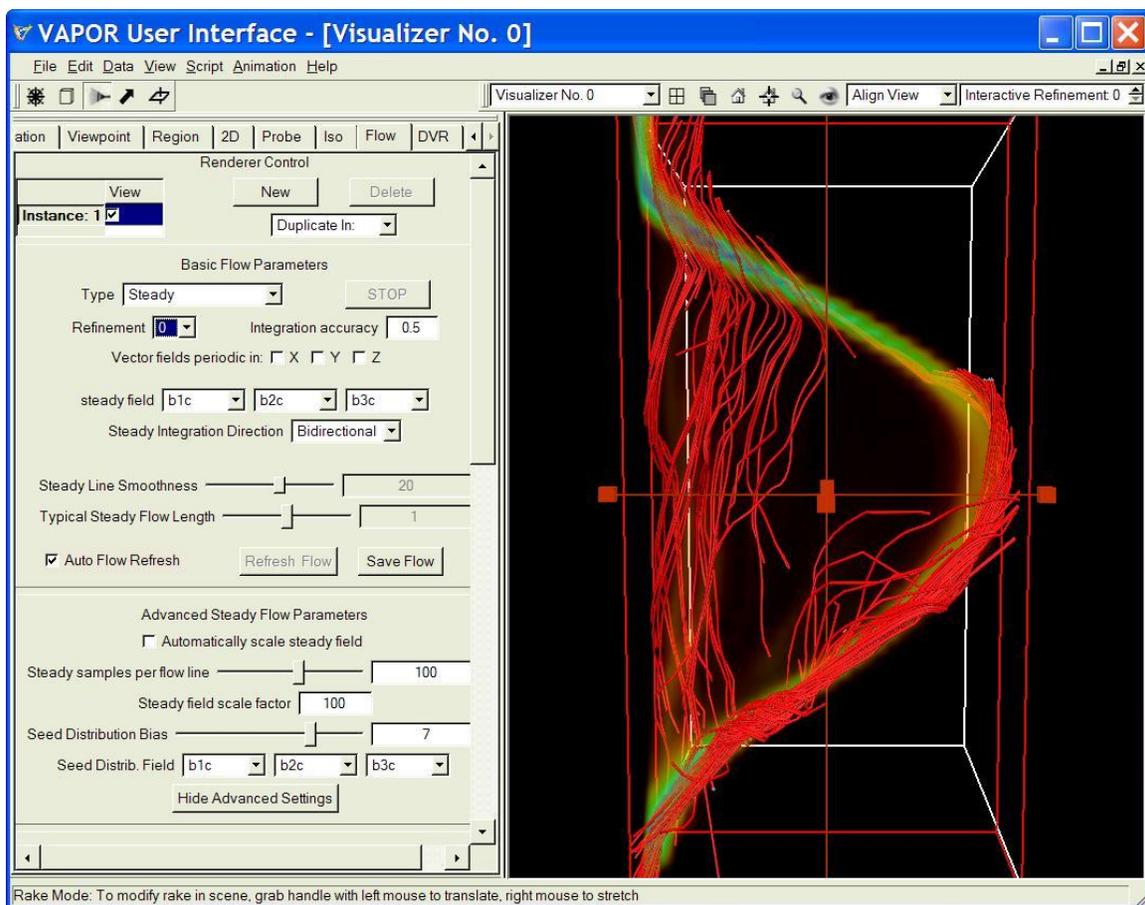


Figure 3.3: Flow lines resulting from biased randomly distributed seeds

3.2 Positioning the probe tool

While biasing the seed distribution is helpful, we can get even better control over seed placement by using the probe. First disable the flow, by un-checking the Instance: 1 checkbox in the Flow tab. Then, click on the bold arrow next to the rake icon near the top left of your screen. You will see a small red rectangle in the center of the scene, with six handles attached. (You may need to rotate the

scene a bit to see the rectangle.) The Probe tab will be visible on the left side of your screen. On the right side, near the bottom, of the probe tab, there is a list of Mapped Variable(s). (You will need to scroll down to see them) With the left mouse, click on “b1c”. Then, holding the control key down, click also on “b2c” and “b3c”, so that all three variables are selected. The result of selecting multiple variables is that the probe shows the root-mean square of the selected variables, i.e. the norm of the vector (b1c,b2c,b3c). Click the checkbox Instance: 1 at the top of the probe tab, to enable the probe rendering, and then rotate the visualization downward, and you will see an image like the following:

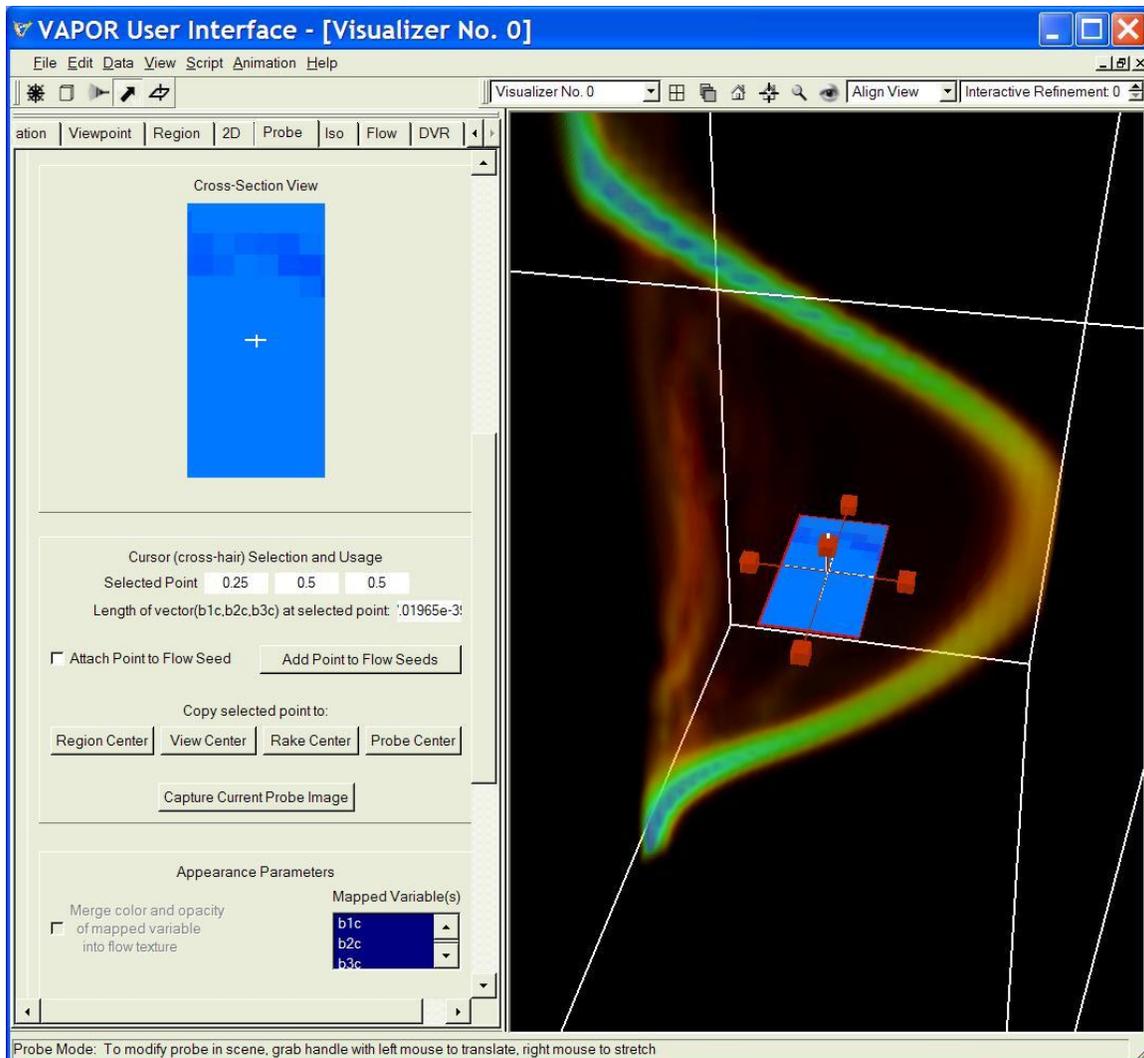


Figure 3.4: Placing a probe tool in the scene

The probe, shown as a blue rectangle in the probe panel, has been inserted in the scene, and has handles attached so you can move it around. Drag the handles with left mouse to position the probe where it cuts through the flux tube. You can alternatively position the probe by using the text boxes or sliders at the top of the probe panel, entitled Probe/Contour Center and

Probe/Contour Size A good starting position for the probe is at $x=0.44$, $y=0.56$, $z=0.5$, with x -size = 0.1, y -size = 0.1, z -size = 0, as shown below:

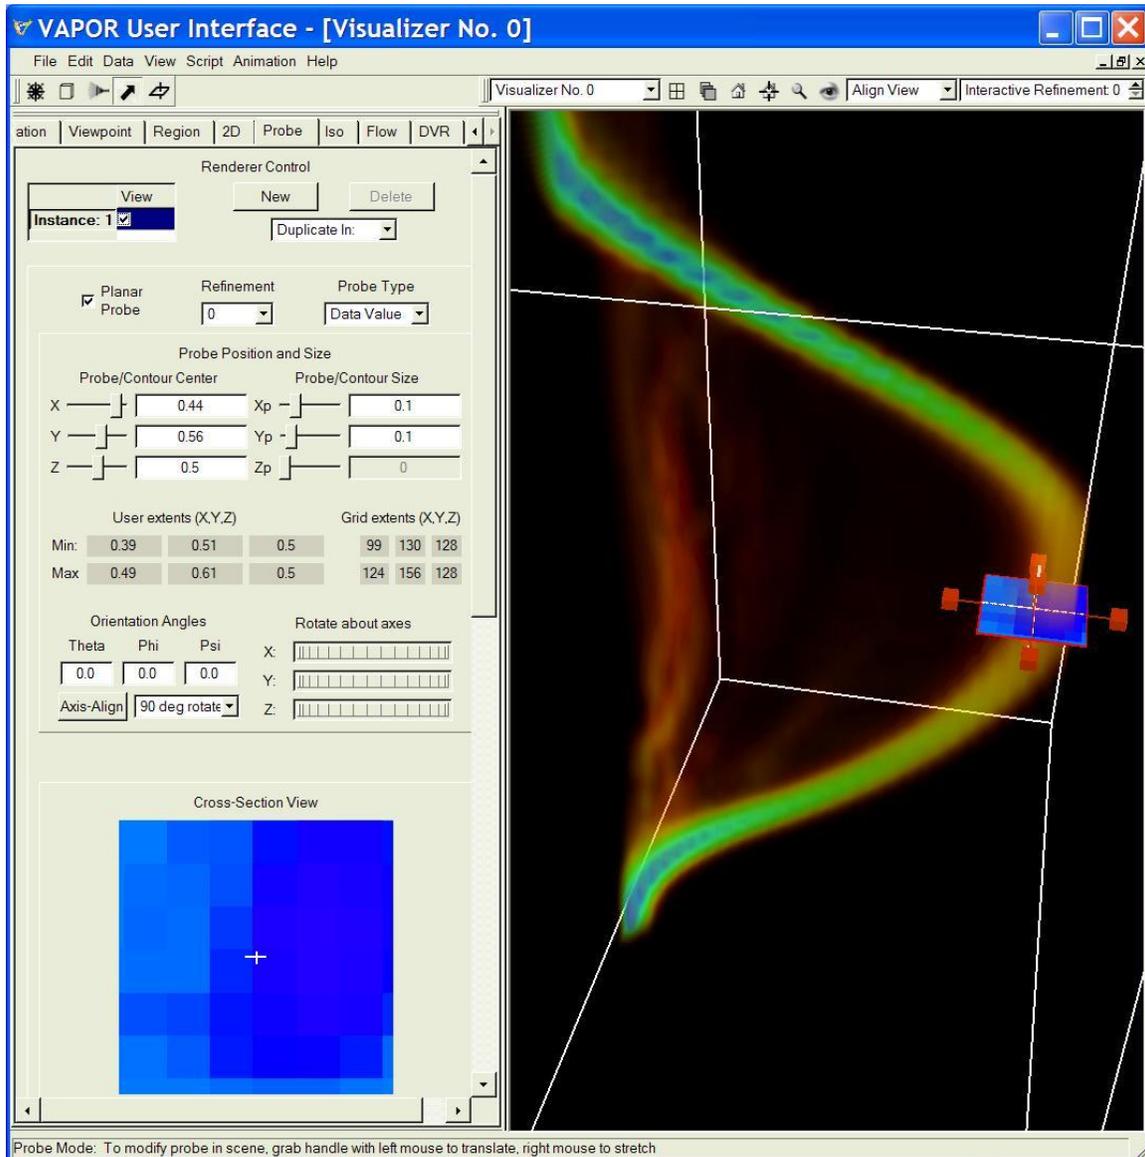


Figure 3.5: Inserting the probe in the flux tube

3.3 Color mapping the probe

The probe Cross-Section View looks pretty coarse, because we are using the lowest resolution data. Select Refinement 2 in the probe panel to get a better image.

If you scroll down to the bottom of the probe panel, you will see a Transfer Function Editor, just like the transfer function editor you used with the volume renderer in section 1. This transfer function controls the colors displayed in the plane of the probe, mapping the magnitude of the field (bc_1 , bc_2 , bc_3) to a color

and opacity. The transfer function editor also displays a histogram of the values in the probe.

For the purpose of understanding the magnetic field in the flux tube we are mostly interested in the large values of the field intensity. You can see what the values are in the probe by clicking the Histo button; they are mostly between 0 and 0.10. To show the larger values, type 0.05 as the left TF Domain Bound and 0.10 as the right TF Domain Bound. Click Fit to View and then click Histo to see the histogram over this range of values. You should see the following:

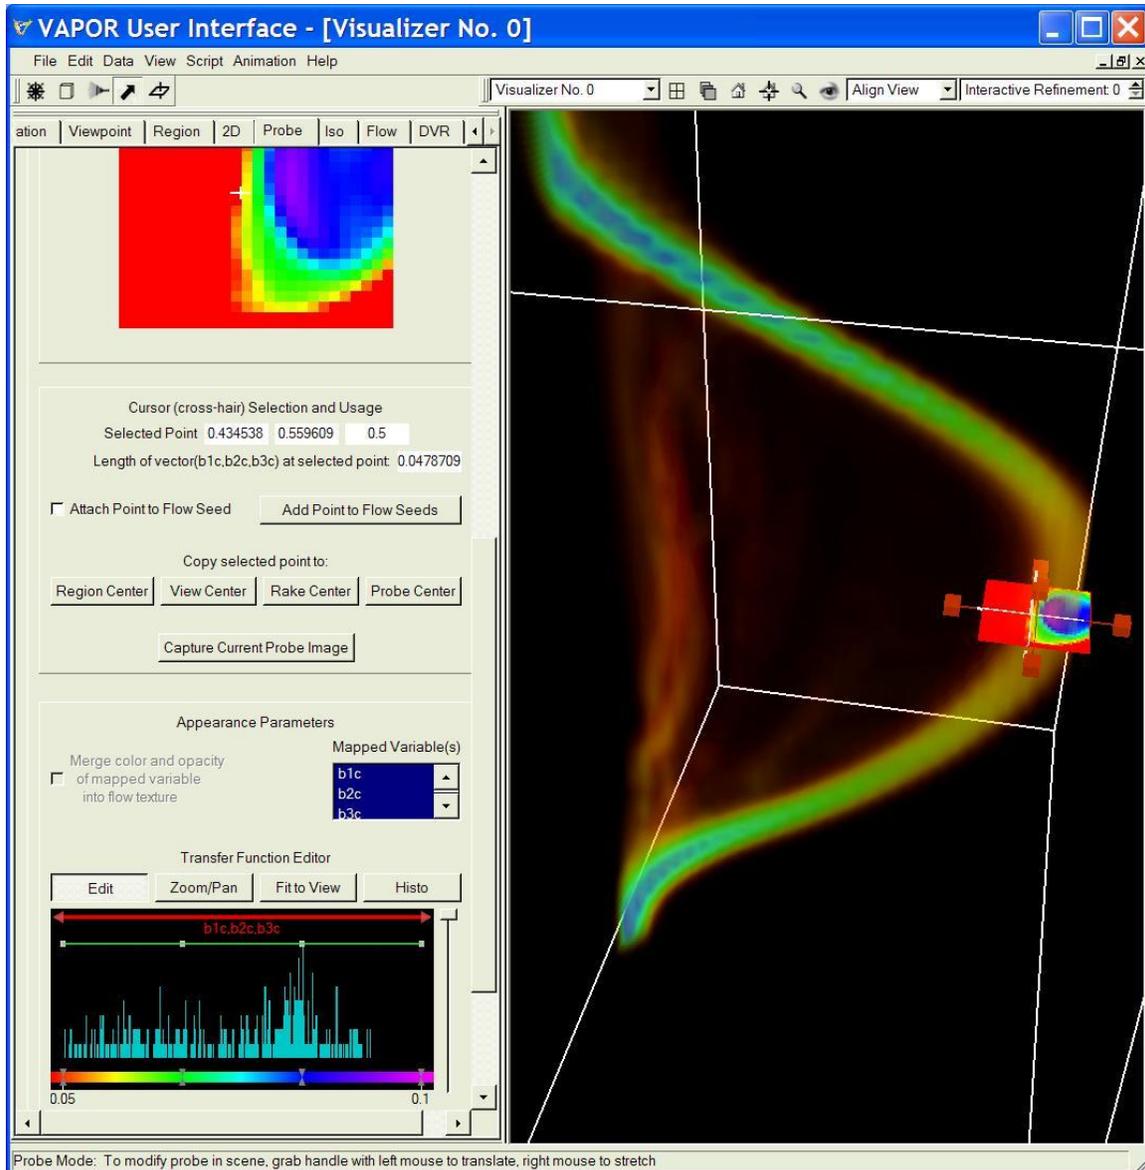


Figure 3.6: Color-mapping the probe to look at the flux tube cross-section

It will make it easier to navigate in this scene if the probe is the center of rotation. Click the left-mouse button in the center of the probe Cross-Section View (in the Probe tab), to establish the selected point position. Then click on

the button labeled Use selected point to center: View. Then, using your mouse in the scene, drag the right-mouse button upward to zoom in, closer to the probe. In order that the volume visualization will not interfere with your view, click on the DVR tab, and then slide the opacity scale in the DVR about halfway down (if it is not already). After you make these changes, the resulting image should be similar to the following:

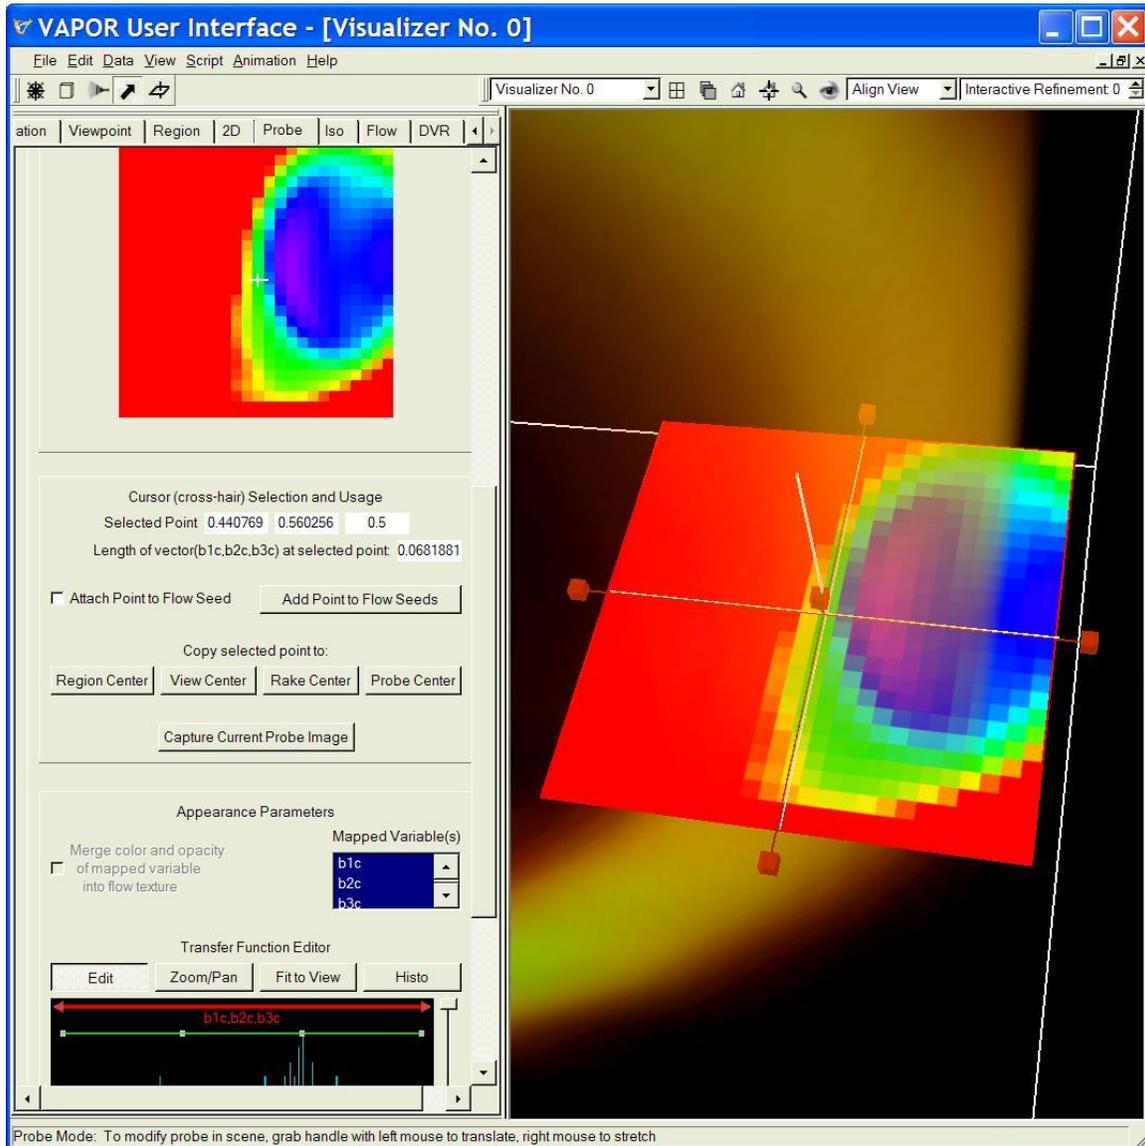


Figure 3.7: A good view of the probe in the flux tube

Now you can understand why it was so difficult to obtain useful flow seeds: The large values of the field are limited to the very small volume in the flux tube.

3.4 Using the probe to select flow seed points.

The probe tool can now be used to specify the location of seed points for field visualization. Before doing this, however, we need to prepare the flow panel to visualize the field lines that come from the probe. Click on the Flow tab. It should already be set to visualize the vector with components (b1c, b2c, b3c). Click on Show Advanced Flow Parameters (if they are not already shown). Uncheck the checkbox labeled “Automatically scale steady field”. Set the Steady field scale factor to 20.0 and set the Steady samples per flow line to 300. Under the Flow Seeding Parameters, select List of Seeds instead of the Random Rake. Select Refinement level 2 (so we will get more accurate tracking of field lines). Under Shape Parameters, set the Diameter to 0.5.

Click to enable this flow instance (if it is not already enabled). You are now prepared to visualize field lines associated with a list of seed points. However, because there are currently no seed points in that list, you won't see any field lines yet.

Next, return to the Probe panel. We shall specify field lines by moving the cross-hairs around the image in the probe tab. Click the check-box labeled Attach Point to Flow Seed. After waiting a few seconds for the necessary data to be retrieved, you will see a red field line that passes through the current cross-hair location in the visualization window. You can click various locations in the probe image and see the field lines that result. Whenever you are satisfied with the resulting field line, click Add Point to Flow Seeds. After you have selected several points from the probe and added them to the seeds, your flow visualization will look something like the following:

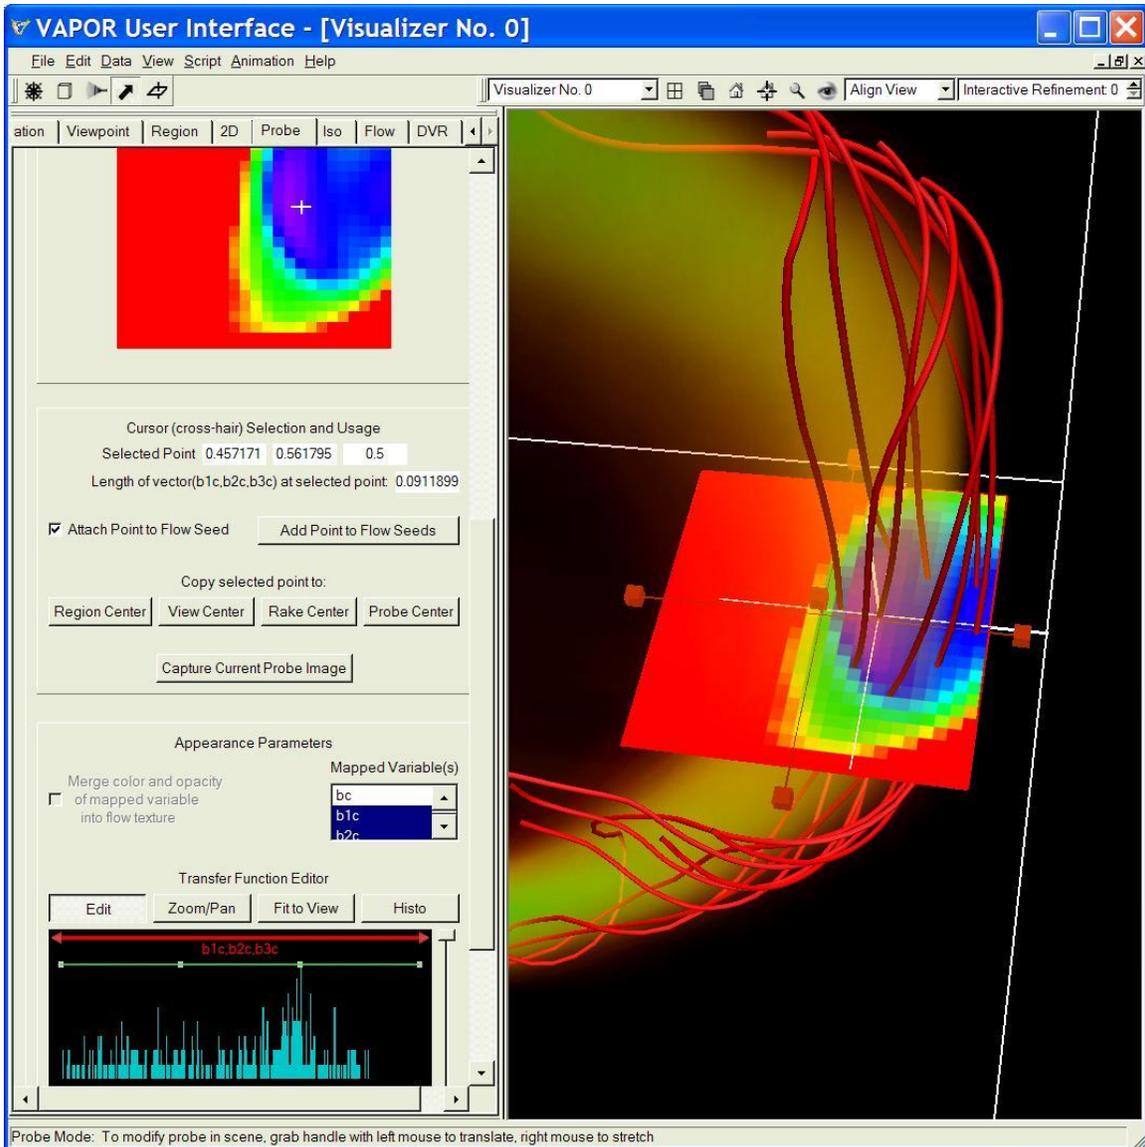


Figure 3.8: Magnetic field lines resulting from seeds selected in the probe

To make it easier to distinguish the various field lines, return to the flow tab, and scroll down to the bottom. Under **Color:** in the **Color/Opacity mapping**, select the color mapped entity to be **Seed Index**. This way each seed point determines a color based on its index. If you have a list of N seeds, the seeds in the seed list will have indices $0, 1, 2, \dots, N-1$. Below the color mapping window you will see **Color Data endpoints**. Set the left endpoint to 0 , and the right endpoint to $N-1$, where N is the number of flow seeds you have selected in the probe panel. You can look at (or modify) the values of these seed points by clicking on **Edit Seed List**. The resulting screen should like the following:

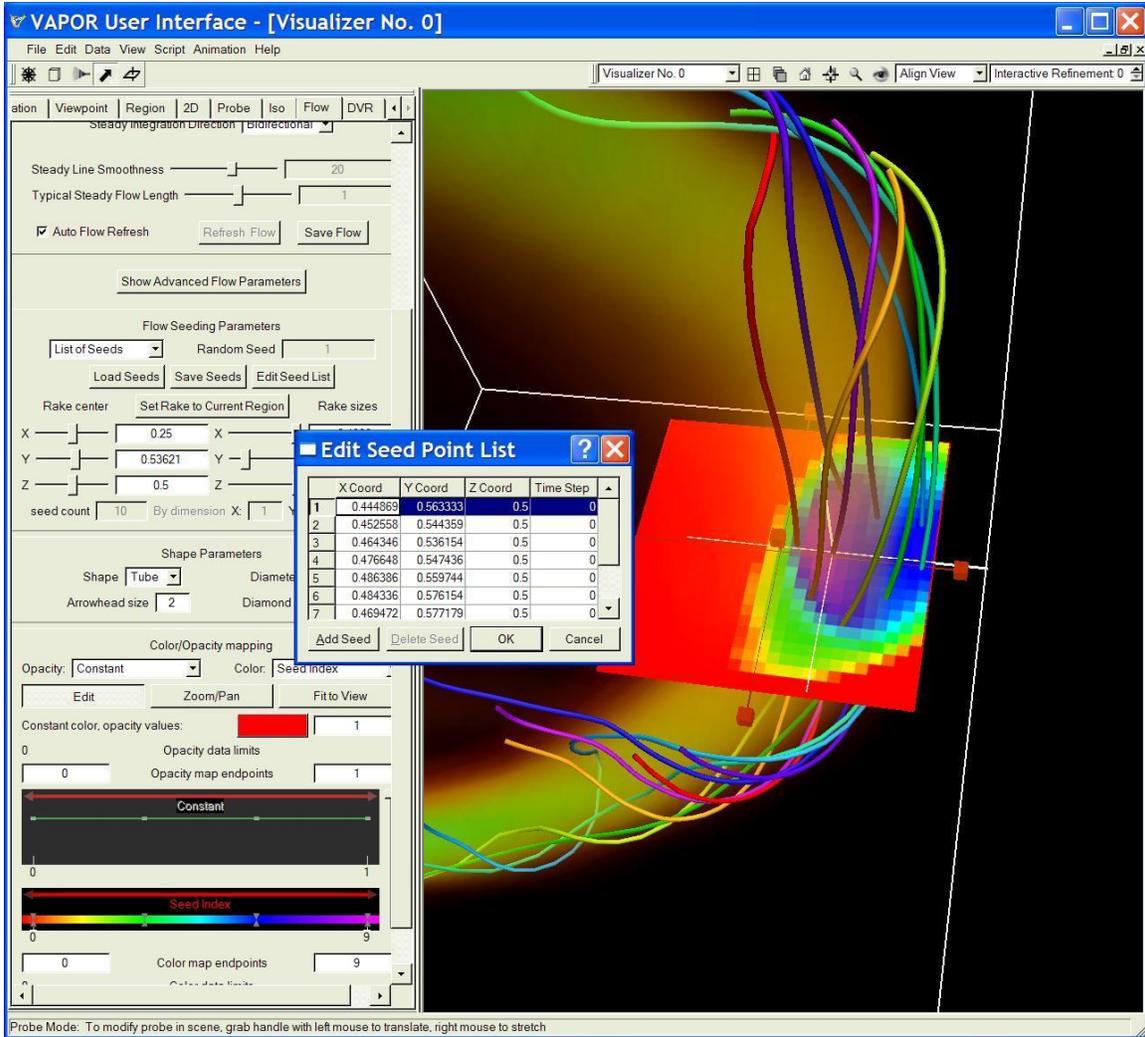


Figure 3.9: Color mapping and editing the list of seed points

In the above screen shot, each of the 9 field lines is colored differently. Note that each seed has four coordinates. The third (Z) coordinate is always 0.5 because we chose the probe to be parallel to the X-Y plane. The fourth coordinate is 0 corresponding to the time step we are currently visualizing. To specify seed points that are valid for all time steps, set the time step to "*" or "-1".

Section 4: Field Line Advection

A feature new in VAPOR 1.1 is the ability to animate field lines under the influence of a velocity field. This is useful in MHD simulation, when there is both a magnetic field and a velocity field, and one would like to see how the magnetic field lines move in the velocity field. Vapor supports a field-line advection algorithm that works as follows:

1. The user specifies a set of seed points for the steady (magnetic) field at a starting time T_0 .
2. These seed points are integrated in the steady field, resulting in a set of field lines.
3. Along each field line, the steady field magnitude is maximized, finding a point with the strongest steady field along that line.
4. The maximizing points are then advected in the unsteady (velocity) field to the next time step T_1 , to be used as seeds in the steady flow.
5. Steps 2 to 4 are repeated for each subsequent time step T_n .

VAPOR also supports a variant of the above algorithm, wherein multiple points along each field line are all advected to the next time step, and the seed at the next time step is obtained by maximizing the steady field among all the advected points from the field line.

For this example we shall examine the OT128 data set, provided by Pablo Mininni of NCAR. This dataset illustrates the phenomenon of magnetic reconstruction in 3D. The dataset has a magnetic field (b_x, b_y, b_z) and a velocity field (v_x, v_y, v_z). Early in the simulation, the velocity field forces together magnetic field lines of opposite polarity, resulting in a current sheet (surface of increased current amplitude) between the oppositely directed field lines. The rotating velocity field bends the field lines near the current sheet, causing them to twist and reconnect.

Later in the simulation, current sheets are sometimes formed between two sets of magnetic field lines at right angles to each other. Using VAPOR, we show how these magnetic field lines can be identified and animated over time, illustrating how the magnetic field lines evolve in the vicinity of the current sheet.

The OT128 data set has a low-resolution (`ot128_lowres.vdf`) and a higher resolution version (`ot128.vdf`) available for download from the VAPOR download page. Either set can be used for this analysis. The `ot128_lowres` dataset has two vector fields: The magnetic field $\mathbf{B} = (b_x, b_y, b_z)$ and the velocity field $\mathbf{V} = (v_x, v_y, v_z)$. There is also a scalar variable, “current”, which is the magnitude of the current (\mathbf{J}) field. These variables are present at refinement level 2 for timesteps 1 through 25. Between timesteps 25 and 61, the variables are only provided at every fourth timestep, because the data is evolving more slowly. The final timestep (61) is provided at higher resolution (at refinement level 3, i.e. a grid of size 128^3).

4.0 Setup for field line advection

Setting up field line advection is a somewhat complex process. VAPOR provides directions in the flow panel for setting up unsteady flow and field line advection. You can view these instructions from the flow tab (after selecting the flow Type to be Field Line Advection or Unsteady) by clicking on the button Field Line Advection Setup Help or Unsteady Flow Setup Help. In the following, we show how this set up is done in one specific case.

Start up VAPOR, and load the `ot128_lowres.vdf` dataset. Select the Flow tab, and select Field Line Advection as the flow Type. Set the Steady field variables to `bx`, `by`, and `bz`. Set the unsteady field variables to `vx`, `vy`, and `vz`. Check the three Vector fields periodic in: checkboxes, because this data is periodic in `x`, `y`, and `z`. Set the Steady Line Smoothness to 100 so we will see nice smooth field lines. Set the Refinement level to 2.

The Unsteady field scale factor is important to get correct because this will affect how fast the field lines move over time. The velocity field needs to be corrected both for the time and the spatial scaling in the original data. The original simulation was computed on a grid of size 2π on each side. The time difference between the time steps of this dataset is $1/8$ second. In VAPOR you can see (from the region tab) that the region is of size 1.0 on each side. We must scale the velocity vector field by the product of these factors, i.e. $(1/8)(1/2\pi)$, or approximately 0.01989 .

Next, we set up the flow seeding to have a planar array of 6x6 seed points. Scroll down to the Flow Seeding Parameters, select Nonrandom Rake, and, where it says By dimension, type in 1, 6, 6, as the `x`, `y`, and `z` grid sizes of the rake. The Seed time start should be "1", indicating that the 36 seeds will be inserted at time step 1. Then, enable this flow instance, and you will see the following:

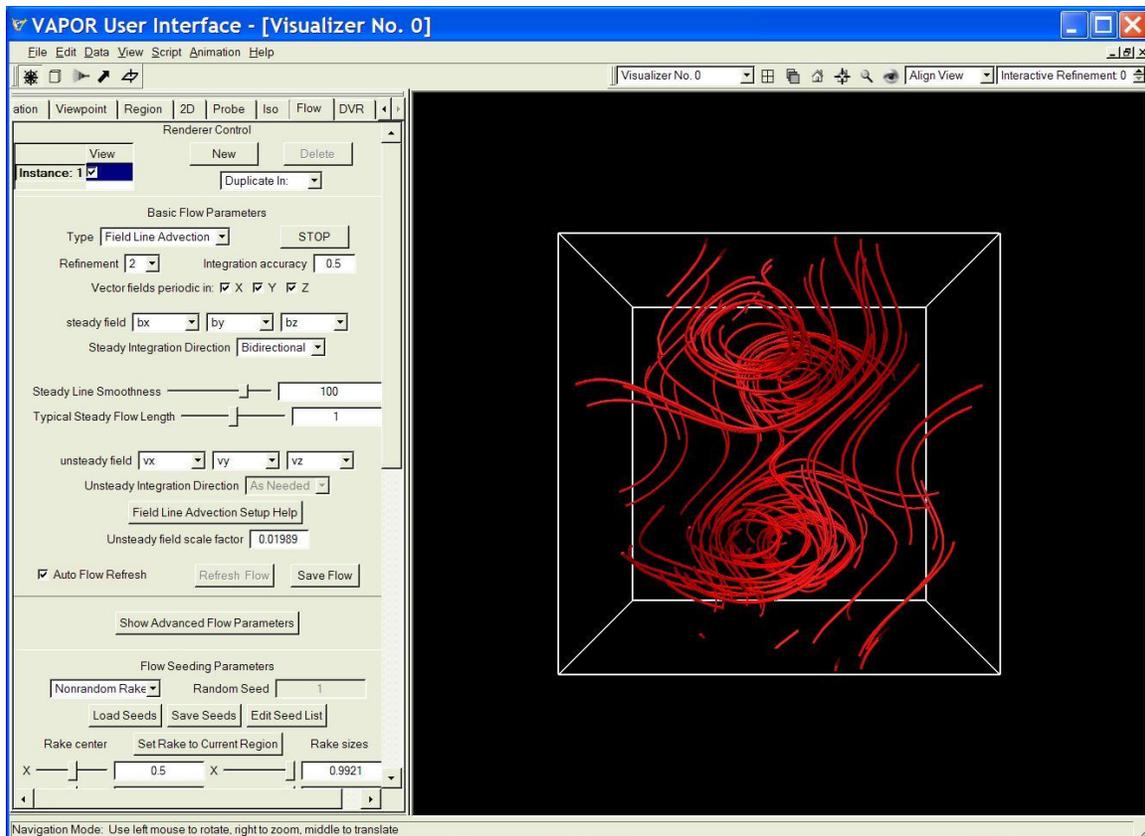


Figure 4.1: Starting magnetic field lines

4.1 Initial evolution of field lines

Note that the field lines are initially tracing more-or-less circular curves, making tubular shapes at the top and the bottom halves of the cubes. Over time these tubes flatten and rotate clockwise. Before animating them, we color-code them to make it easier to track the motion from the initial field lines. Two collections of seed points will be constructed by dividing the rake into two separate rakes and color coding them differently. Scroll down to the Flow Seeding Parameters (in the Flow tab), set the Y-value of the Rake Sizes to 0.5, and set the Y-value of the Rake Center to 0.25. The resulting rake occupies the bottom half of the volume. (You can see the rake by clicking on the rake icon at the top-left of your screen.) Click on the button labeled Save Seeds to save these seeds to a file "bottomseeds.txt". Then set the Y-center of the rake to 0.75, and save the second set of seeds to a file named "topseeds.txt". From your favorite editing program, create a text file that consists of the 36 seeds in "bottomseeds.txt" followed by the 36 seeds in "topseeds.txt". Save this file as "allseeds.txt".

Returning now to the VAPOR Flow panel, click on the Load Seeds button and load the file "allseeds.txt". You now have a collection of 72 seed points where the first 36 are on the bottom and the last 36 are on the top. Select List of Seeds (instead of Nonrandom Rake) in the Flow Seeding Parameters. Scroll down to the Color/Opacity mapping of the flow tab. Beside Color:, select Seed Index as the color-mapped entity. Then type in "71" as the right Color map

endpoint. The field lines will then be shown in 72 different colors. Edit the color map so that values 0...35 are mapped to green, and values 36...71 are mapped to maroon. Zooming in, the visualization should now appear as follows:

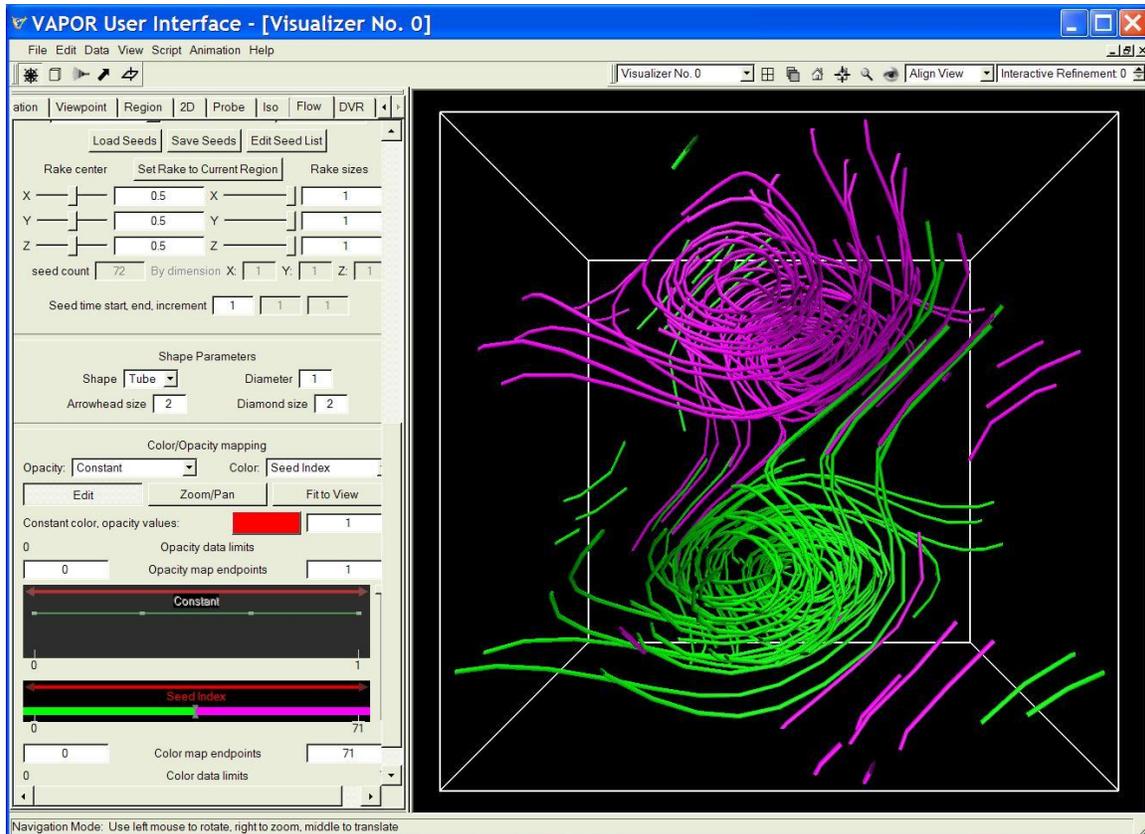


Figure 4.2: Two collections of separately colored of flow lines

To make these lines evolve over time, click on the Animation tab. Type in "25" as the End Frame Position (the data is only available for $\frac{1}{4}$ of the time steps after 25). Then advance the animation forward in time, and watch the motion of the field lines. You are seeing how the magnetic field lines are moved by the velocity field. The field lines rotate coherently for the first few time steps but eventually (say by time step 20) the lines become quite intertwined. Figure 4.3 shows this animation at time step 14:

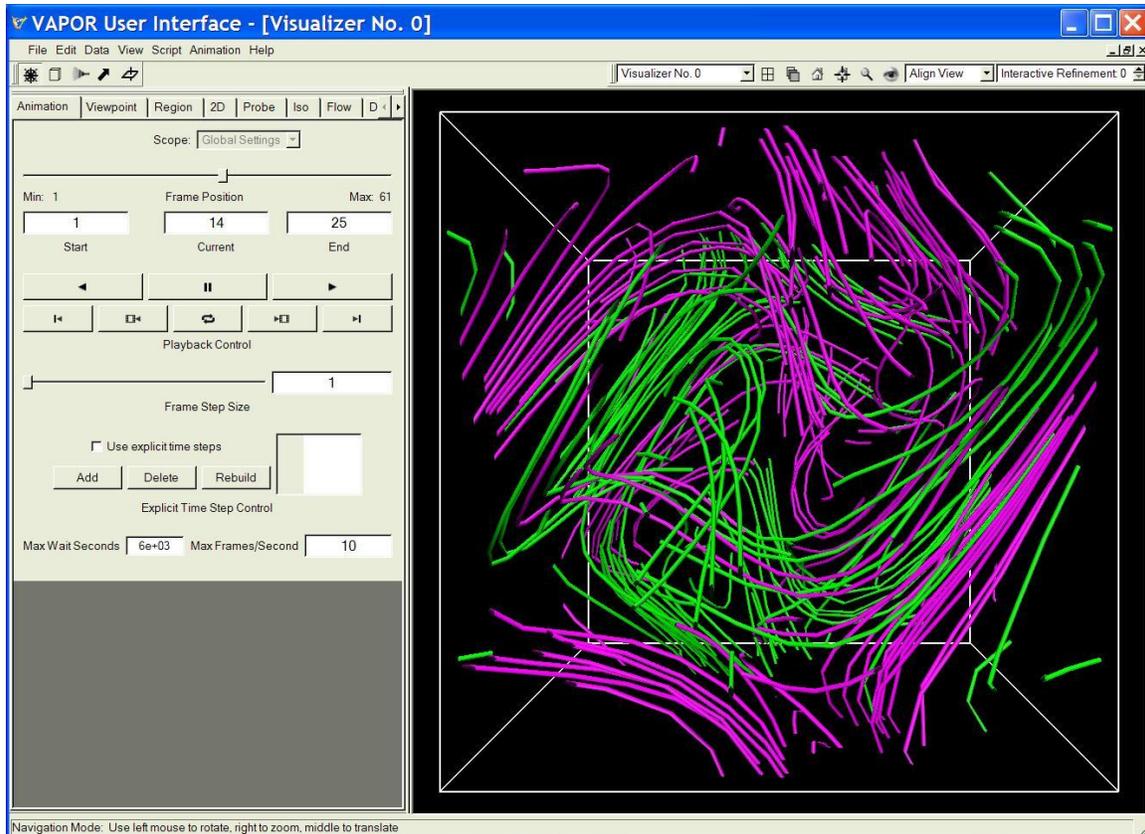


Figure 4.3: Animated magnetic field lines at time step 14

4.2 Animating field lines near a current sheet

As this simulation advances, eventually the current magnitude attains its highest magnitude on relatively flat, sheet-like shapes known as “current sheets”. We shall next observe the time-evolution of magnetic field lines near a current sheet.

We can make use of some of the parameters from the previous flow integration. Select the Flow tab, and un-check (disable) the current flow instance (Instance: 1). In the Renderer Control for this flow tab, select Duplicate In: and then select This visualizer to create another flow instance, identical to the previous. We shall set up the second flow instance to observe the flow between time steps 25 and 61.

To observe the current sheet at the later time steps, click on the Animation panel, and set both the End Frame Position and the Current Frame Position to 61. Set the Frame Step Size to 4 (so that only $\frac{1}{4}$ of the frames will be shown). Click on the DVR tab, and select “current” as the Variable. Select Refinement level 3, and check Instance: 1 to enable this DVR instance.

We shall modify the transfer function (in the DVR tab) to show the higher intensities of the current. Set the right TF Domain Bound to be 30. Press enter, then, click the Fit to View button and also the Histo button. Select the second opacity control point (the control point that is about $\frac{1}{3}$ of the way from the left to the right) and drag it down to the x-axis, at the x-coordinate of 10.0. With this

transfer function, the volume rendering exhibits several sheet-like structures, as follows:

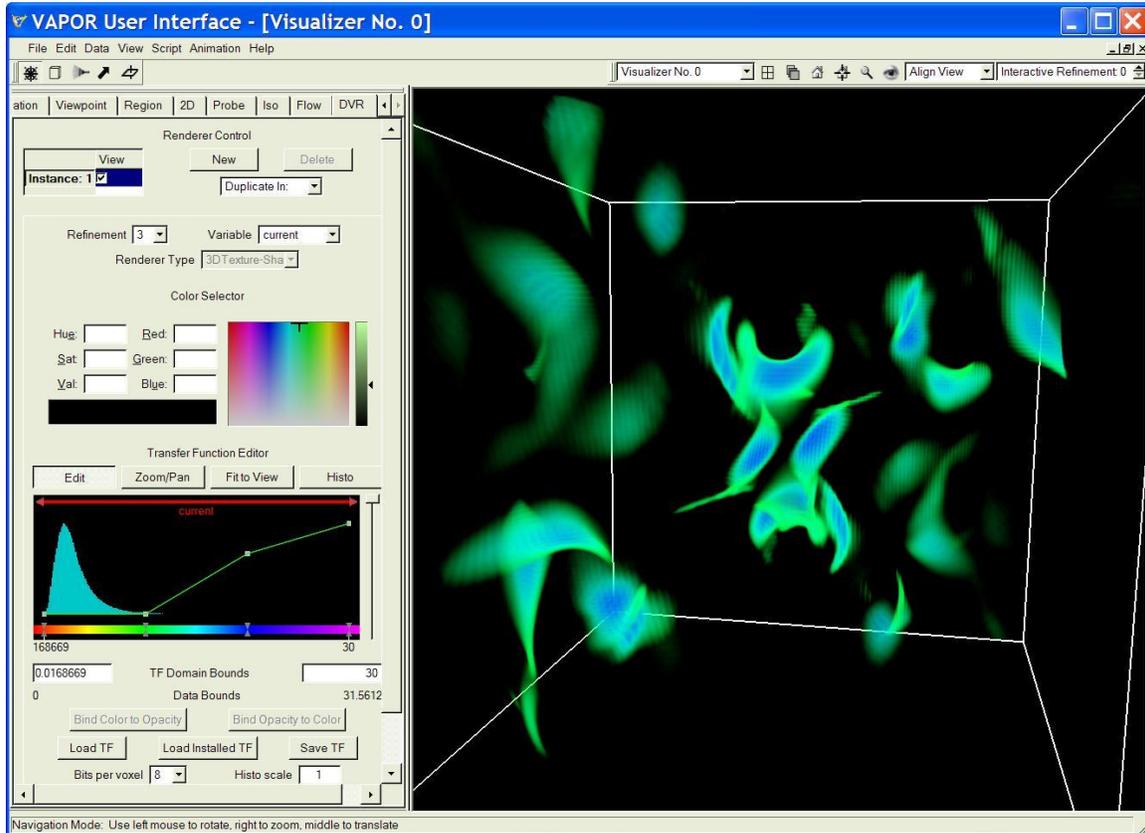


Figure 4.4: Volume rendering of current sheets at time step 61

Next we use data probes to look closely at the structure of one of these current sheets. Note the curved structure slightly above the center of figure 4.4. Click on the **Probe** tab, set the **Mapped Variable(s)** to “current”, and set the **Refinement** level to 3. Set the right **TF Domain Bound** to 30, and click on **Fit To View** as well as **Histo** in the **Transfer Function Editor** on the **Probe** tab. Enable the probe, by clicking in the **Instance: 1** checkbox. Turn on data probe mode, by clicking on the bold arrow icon below the title bar, so that you can move the probe using handles in the visualizer window. Set the **Probe/Contour Center** to (0.52, 0.65, 0.66), centering the probe in the curved (saddle-shaped) structure, as follows:

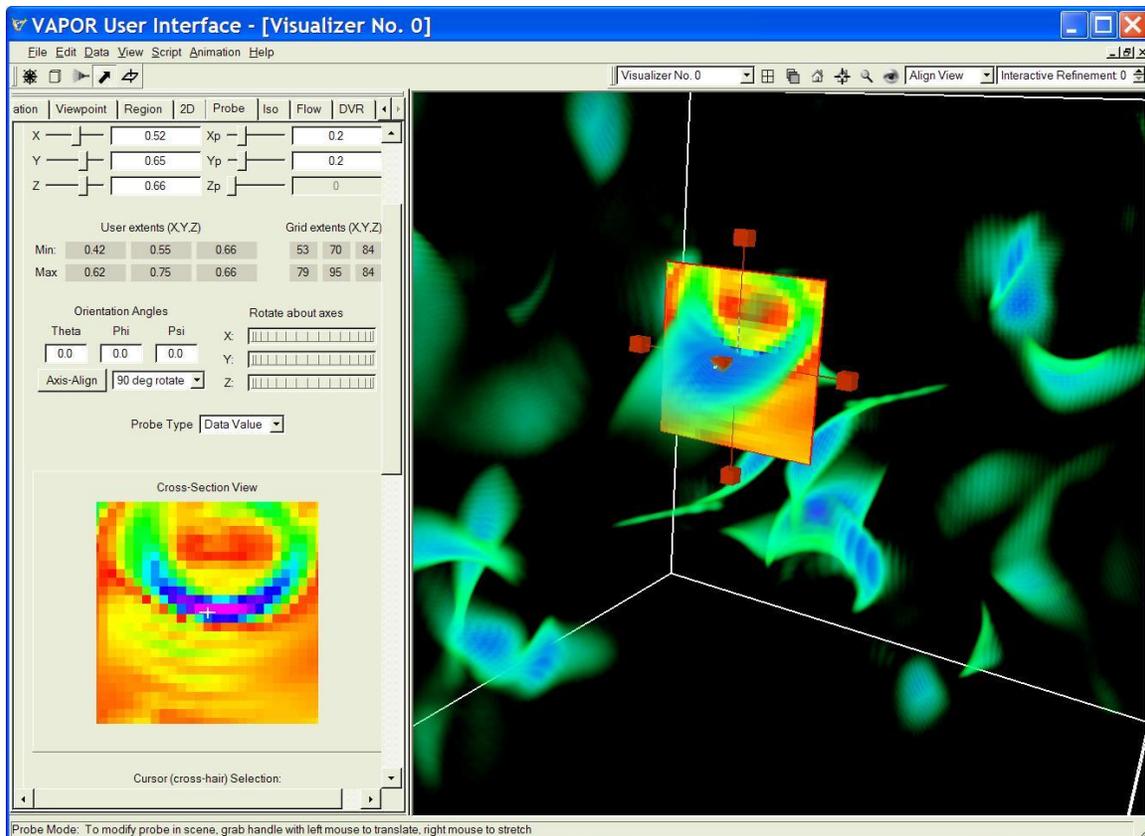


Figure 4.5: Probing a saddle-shaped current sheet

An additional probe will help us to see the 3D extent of this saddle. In the **Probe** tab **Renderer Control**, select **Duplicate in:** followed by **This visualizer**, to make a copy of the first probe instance. Then select the second Probe instance, and enable it. Using the **Y**-thumbwheel under **Rotate about axes** in the **Probe** tab, rotate this instance 90 degrees (or, just set the **Phi** angle to 90.0, or select 90 deg rotate and pick “- Y” as the rotation axis). To get a closer view, click in the **Probe** tab on the button labeled: **Use selected point to center: View**. Rotating and zooming in, you can see both probes and the saddle in the current sheet as in the following (reducing the volume opacity in the **DVR** tab helps):

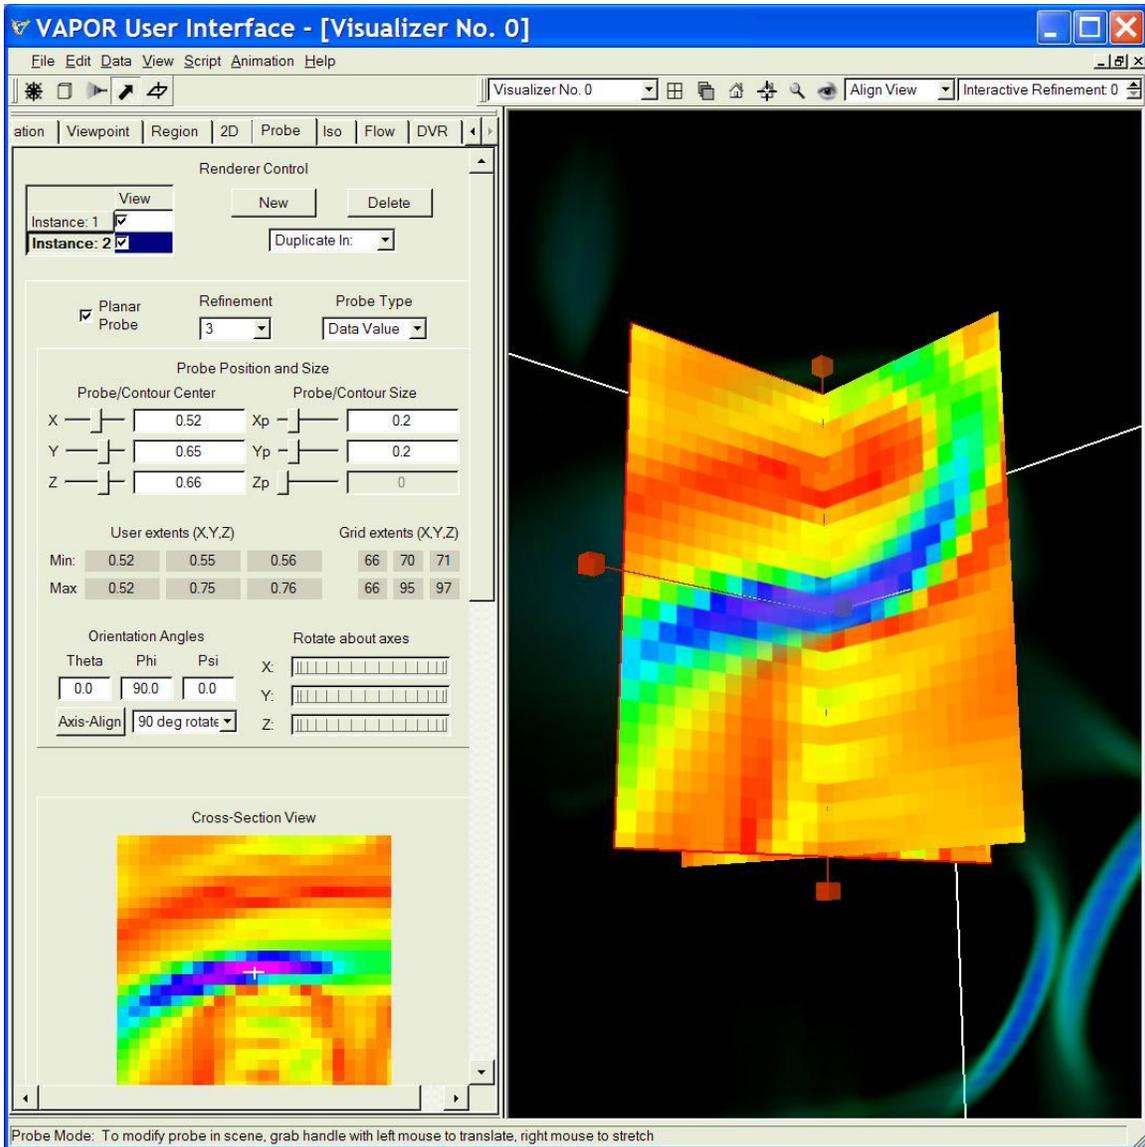


Figure 4.6: Two cross-sections of a saddle-shaped current sheet

Now we shall use these probes to construct magnetic field lines near the current sheet. Click on the Flow tab, and select Instance: 2, but don't enable it yet (i.e., don't check the checkbox). Scroll down to the Flow Seeding Parameters, make sure List of Seeds is selected, and set the Seed time start to 61. This way the seeds for the magnetic field will start at time step 61, and the magnetic field lines at other time steps will result from advecting backwards in time.

The time sampling needs to be changed, because the default values are not appropriate in this case. Click the button Show Advanced Flow Parameters (i.e., make sure the "Advanced Field Line Advection Parameters" section of the Flow panel is visible). Set Timestep Sample min, max, interval to 1, 61, 4, so that we will use only $\frac{1}{4}$ of the time steps for the field line advection.

The second flow instance has lots of seed points (72 were copied from instance 1) and we need to get rid of them. In the Flow Seeding Parameters, Click on the button Edit Seed List and you will see a list of seed points. Select one of the seeds with your left mouse button, and then repeatedly click the Delete Seed button until they are all gone. Then click OK.

Finally, enable the second Flow instance, by clicking in the Instance: 2 checkbox. Also check the Auto Flow Refresh checkbox (if it's not already checked). You will not see any flow lines, because we haven't specified any seeds for this flow.

Click on the Probe tab; we shall use it to select seed points near the current sheet. Select the first instance, and check Attach Point to Flow Seed. On the Flow tab, enable the second instance, and check Auto Flow Refresh if it is not already checked. After a few seconds you will see a flow line such as the following:

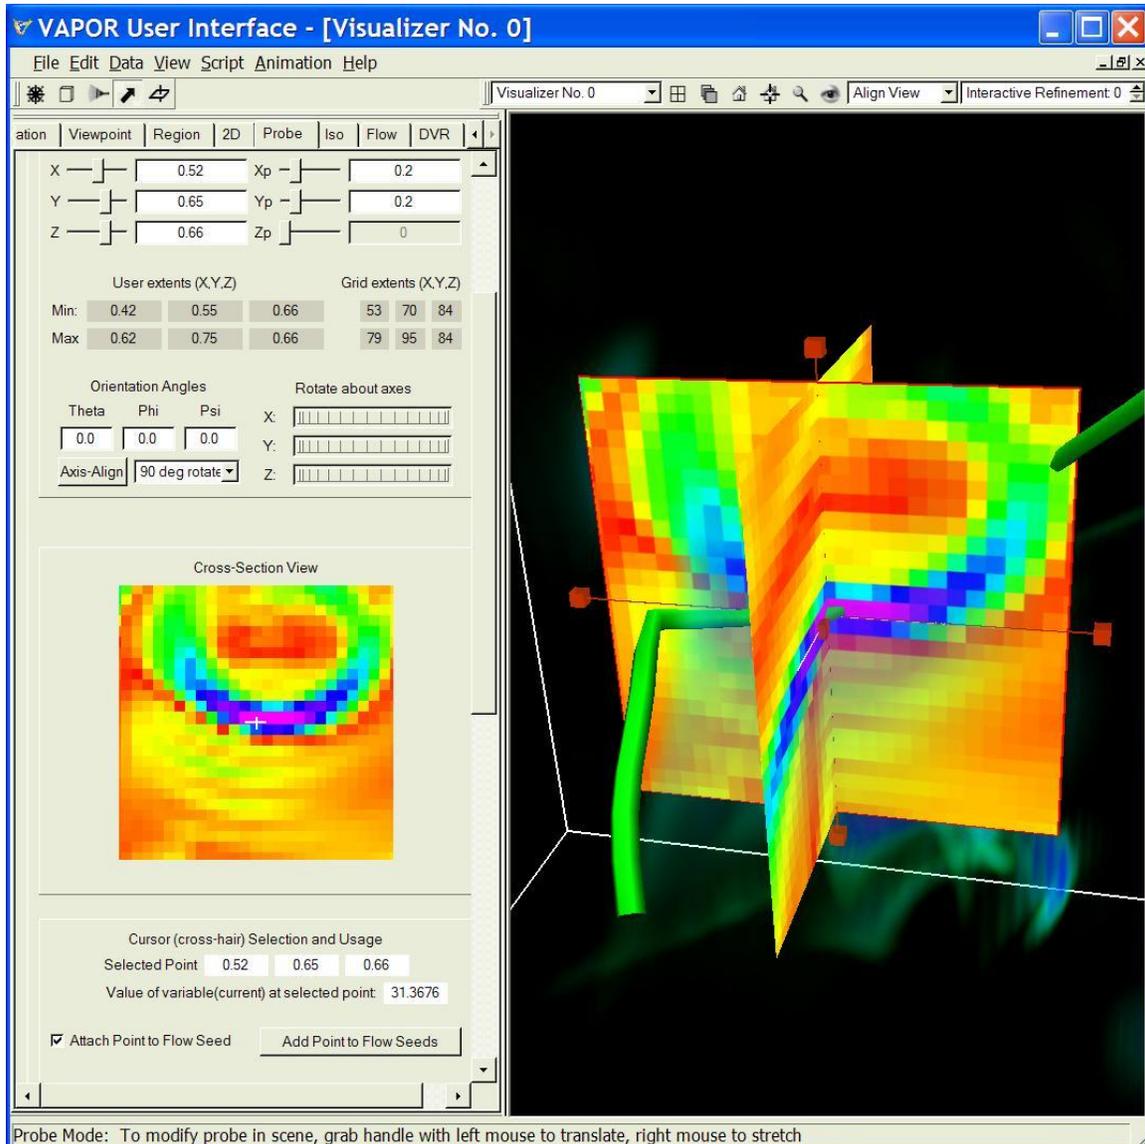


Figure 4.7: Inserting a flow line near the saddle.

We shall insert about 8 flow lines above the saddle, using the first probe instance, and then 8 flow lines below the saddle with the second probe instance. With the first probe instance selected, position the mouse at various positions slightly above the current sheet (which appears as a purple crescent in the Cross-Section View of the Probe tab). Each time you identify a seed point that results in a downward curving flow line, click Add Point to Flow Seeds, adding that point to the seed list, until you have 8 such lines. Then (using the Probe tab's Renderer Control) select the second probe instance, and choose 8 seeds that result in upward-curving flow lines, by choosing points slightly under the current sheet. Finally, returning to the Flow tab, change the color map, so that we can distinguish the two sets of 8 flow lines. Set the right Color map endpoint to 15, and you may need to modify the color mapping so that the first 8 lines are green and the last 8 lines are maroon, like the following:

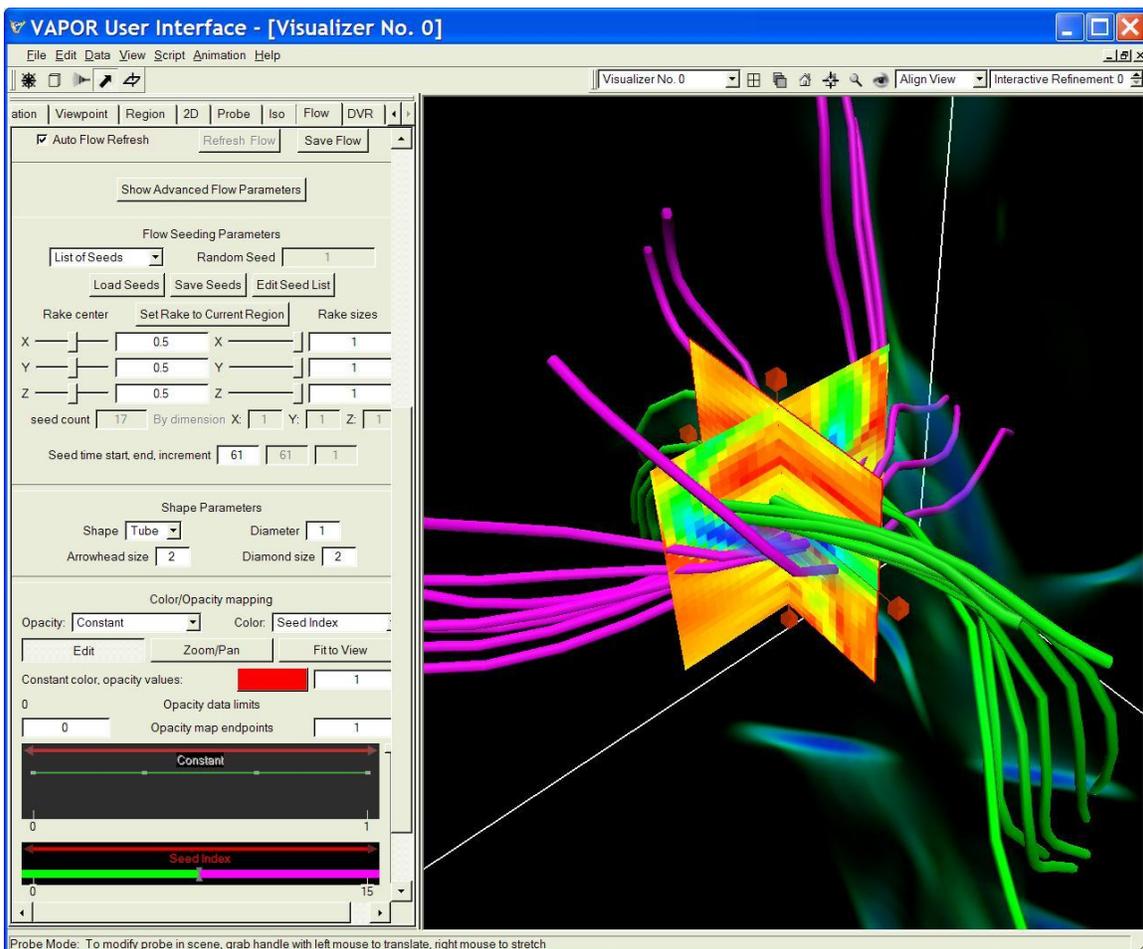


Figure 4.8: Flow lines above and below the current sheet

We can now view how the magnetic field lines move in the velocity field during the formation of the saddle shape. Disable the two probe instances by un-checking them in the Probe Renderer Control. Set the Refinement level in

the DVR tab to level 2, since the earlier time steps do not have data at refinement level 3. From the Animation panel, step backwards in time. This will cause the field lines to be advected backwards in time in the velocity field. You will note that the collection of field lines stays together for a few time steps, but gradually some of them move away. Figure 4.9 illustrates the lines at time step 33:

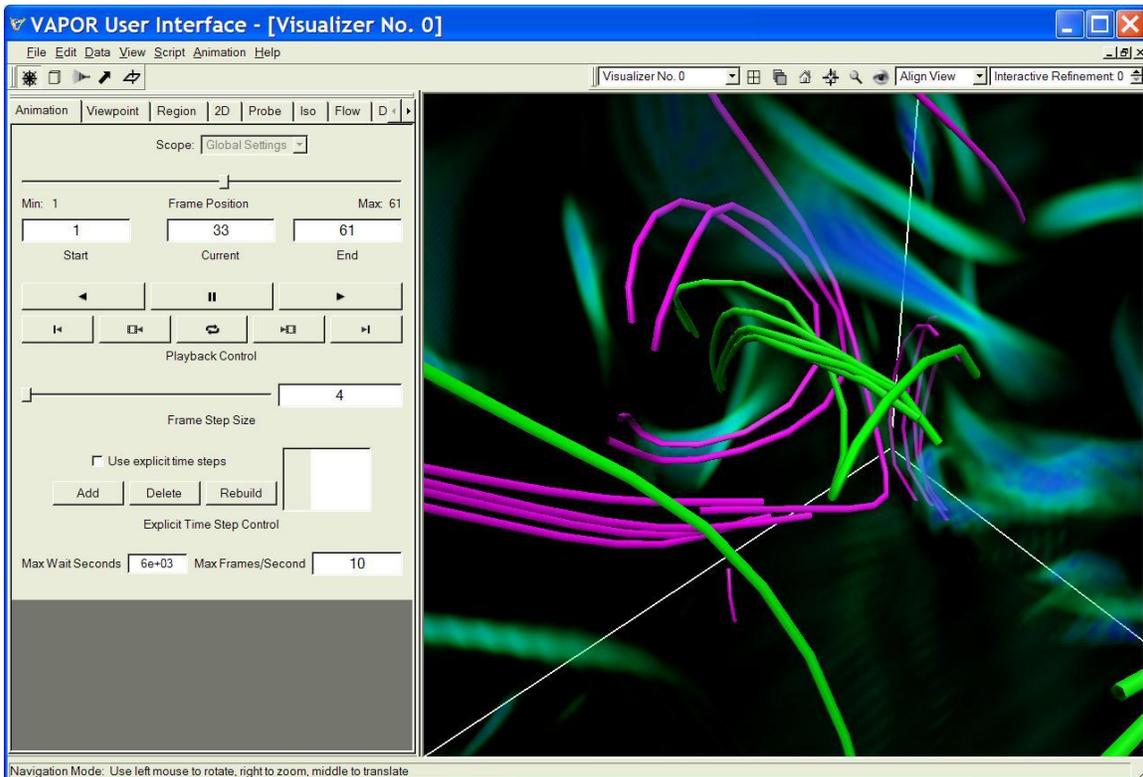


Figure 4.9: Field lines advected in the velocity field back to time step 33

We can also track these field lines using a modified algorithm that first advects multiple points along a field line and then selects the one advected point with the largest field magnitude. This alternate approach sometimes provides improved tracking of field lines. From the Advanced Field Line Advection Parameters, select Prioritize after advection (instead of the default Prioritize before advection). Type in “20” as the number of Points to advect. This will result in 20 points being advected from each field line, and the resulting advected field line is calculated from the advected point having the highest field magnitude. Using this setting, you may see that the collection of advected field lines is slightly more cohesive, showing a similar arrangement all the way back to time step 17, such as the following:

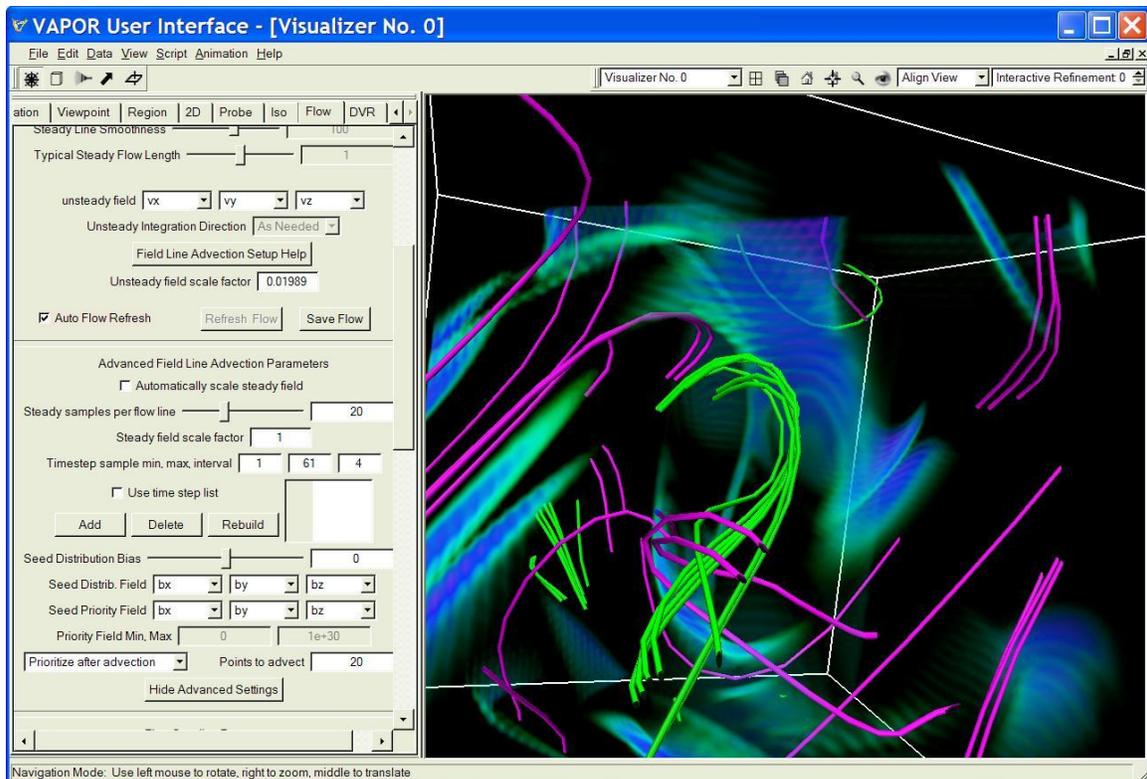


Figure 4.10: Field lines advected back to time step 17 using “prioritize after advection” option.

This data set contains numerous other current sheets. You may find it interesting to identify another current sheet and to advect the nearby field lines forward and backward in time.

Section 5: Advanced topics

5.0 IDL Import and Export

An important feature of VAPOR is the ability to perform analysis concurrently with visualization of data. Currently VAPOR supports analysis using the IDL language, from Exelis. Information about IDL is available on the Exelis website, at <http://www.exelisvis.com/ProductsServices/IDL.aspx>. In this guide we shall show how a VAPOR session can interact with an IDL session by importing and exporting data. Several examples of applying IDL code to VAPOR data sets are provided in the VAPOR VDC document, <http://www.vapor.ucar.edu/page/vapor-data-collection>. Here we shall show how those IDL routines can be used in combination with a VAPOR visualization session. You do not need to know IDL in order to run this example.

The IDL import and export features provide a means of communicating data between an IDL session and a VAPOR visualization session. From the VAPOR user interface, when you export data, you provide a description of the current data volume (in the form of an xml file) that can be read by the IDL session. From IDL, you can perform numerical operations on the data, resulting in new VAPOR data sets (described by .vdf files). Then, using the VAPOR user interface, you can import the data into your VAPOR visualization session, adding the new data to the previously existing session.

Set up for IDL analysis: These instructions are a bit different depending on whether you are using Windows or a UNIX platform. In either platform you need to write to the IDL examples directory (It's the examples/idl subdirectory of the VAPOR install directory). Copy the contents of examples/idl to a writable directory if you do not have write-access to it. Follow the setup instructions below, depending on whether you are using UNIX or Windows:

On UNIX (or Mac): To run this example you will need to run an IDL session in a shell window separate from the window in which you will run the VAPOR user interface. Initialize that shell by running vapor-setup.csh (in a c shell) or vapor-setup.sh (in a Bourne shell) then cd to your idl examples directory. Type "idl" and press enter to start running idl in that shell.

On Windows: In Windows you run IDL from the IDL IDE (I.e. the IDL desktop application). To initialize your environment, from the IDL IDE, click File->Preferences, and select the "startup" tab. Specify the Working Directory as the (writeable) directory where you have put the VAPOR idl examples. Under the IDL "path" tab add the VAPOR bin directory (which is \$(VAPOR_HOME)/bin) (in addition to the <IDL_DEFAULT> directory).

Now that your VAPOR environment has been set up, start vaporgui in its own window, and open the session file you created in section 1.6, "8spothighres.vss". If you did not save that file, you can use the session file for figure 1.11,

“guideFigure1.11.vss”. Then load the data “8spot3d-lowres.vdf” into the current session.

Exporting the data: From the VAPOR DVR tab, enable Instance: 1 to see the volume visualization of the vortex that we isolated in section 1.11. Click on the cube button near the top-left of your window, to enter region select mode. You will see a red wireframe that identifies the boundary of the current region. We would like to perform analysis of the data in that region, to better understand the flow in the vortex. From the Data menu, Click on Data→Export to IDL. This operation identifies the region around the vortex so that the IDL session can use it, as shown below:

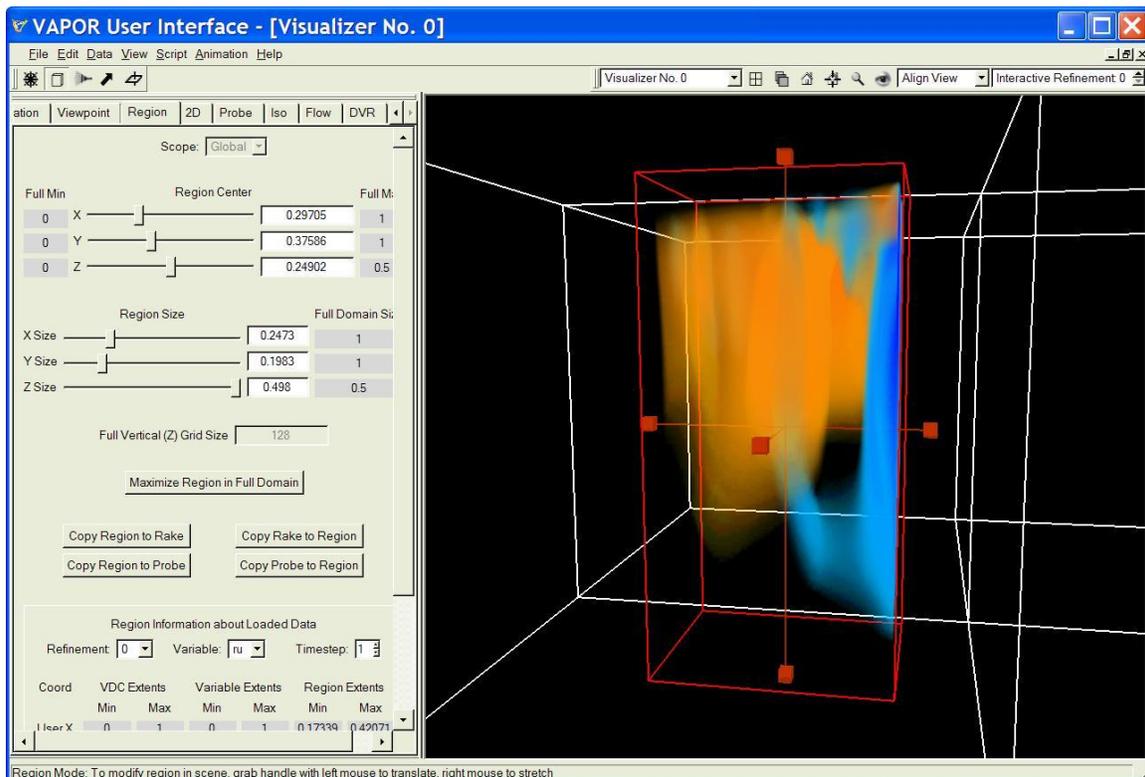


Figure 5.1: Specifying a region to export

Importing the data into the IDL session: There is an IDL script, “QuickStartEx1.pro”, provided in the VAPOR examples directory. When executed, this script will calculate the z-component of vorticity of the flow in the specified region. From the IDL command line, type “.RUN QuickStartEx1.pro” and press enter. You will get several lines indicating the progress of the idl, such as:

```
% Compiled module: $MAIN$
% Compiled module: IMPREGION (This reads the region you exported)
% Loaded DLM: LIBIDL_VDFMODULE (This loads the VAPOR library)
% Compiled module: MYDERIV (This calculates the vorticity of the flow)
```

% Compiled module: EXPREGION (This outputs the results as a .vdf file)

You may also see a message about an arithmetic error. You can ignore this. Note how long it took to execute this script. It would have taken much longer if you had calculated the vorticity for the full volume!

Merging the resulting data back into your VAPOR visualization session:

From the Data menu, click Merge (Import) a Dataset into Current Session. Then select the file “QuickStartEx1.vdf”, which will be located in your idl examples directory. A dialog will pop up, asking you to select a timestep offset. This specifies the difference in timestep numbering between your original data (8spot3d_lowres.vdf) and the new data that you created in IDL. In this case, the data was exported from timestep 0 of the original data, so the correct offset to use is 0. Click OK. VAPOR will then augment the current session by loading the new volume into the current session.

It would be a good idea to now save the new (merged) metadata in one metadata file. This will make it easier if, for instance, you would like to restore this session in the future. Click on Data→Save the Current Metadata to File and specify a vdf filename (different from 8spot3d_lowres.vdf) associated with the merged data.

To visualize the new data, click on the DVR tab, and then select the Variable “omz”. This is the new variable you have just loaded. Check Instance: 1 to enable the rendering, and you will see the data, but with a transfer function that isn’t very useful. You can edit the transfer function as we did in section 1.4 above, to create a transfer function that highlights the positive and negative values of omz:

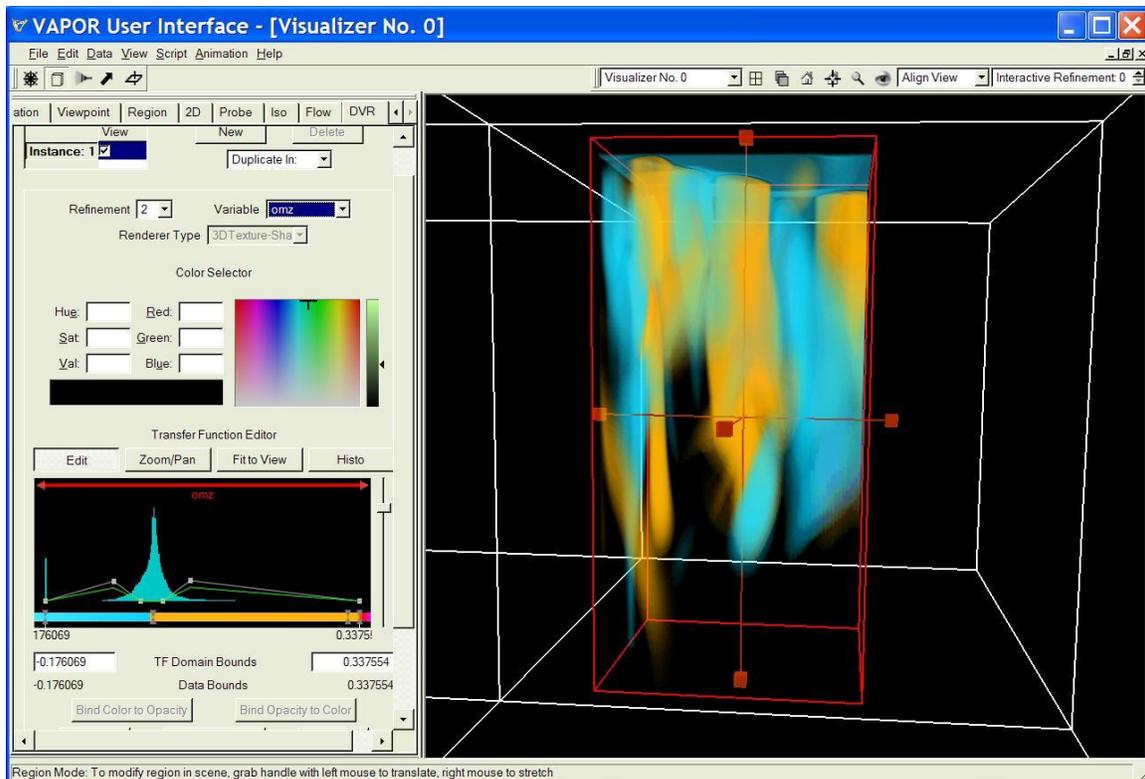


Figure 5.2: A visualization of omz with a transfer function similar to the one in section 1.

You can now use the new data (the omz variable) to position seed points in the flow. We leave as an exercise for the user to inject seed points into the flow at the points of maximal value of omz. This is done just like we showed in section 3.4 above. You will find that, when you insert seeds at a large value of omz, the resulting flow lines move upward, eventually twisting and turning back down, such as illustrated below:

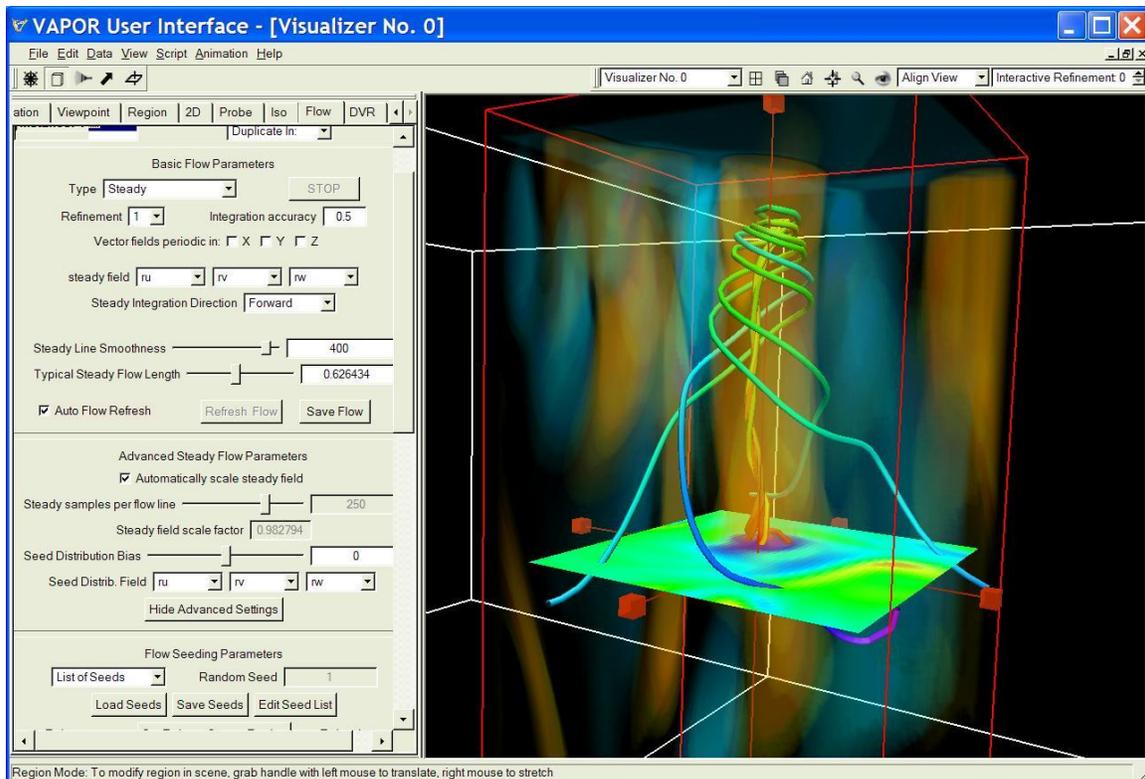


Figure 5.3: Stream lines resulting from injecting seeds at large values of omz

5.2 Using Multiple Visualizers

So far all of the visualizations we have made were performed in one window. However it is often useful to do a side-by-side comparison of different visualizations. For an example we show how to perform a comparison of the vorticity variable “omz” we just created with the “rw” variable that was visualized in section 1.

Make your VAPOR window wider by grabbing the right edge with the left mouse and dragging it to the right, if your window is not already wide enough. This will give us more room for the visualizers. Then click in the selector at the top-middle of your window, where you see the text *Visualizer No. 0*, and select Create new visualizer. You will see an empty window on the right, labeled *Visualizer No. 1*. Select the DVR tab and then click the top bar of a visualizer. You will see that the parameters in the DVR tab change, depending on which visualizer you select. Click on *Visualizer No. 1*, and then, in the DVR tab, select the Variable “rw”, select Refinement 2, and enable that instance (Instance: 1). You will see the variable rw, but with the default transfer function.

You could repeat the steps in section 1.4 to create a transfer function here. But an easier way to fix the transfer function is copy the DVR instance we previously made (in section 1.4) from Visualizer 0 over to Visualizer 1. Click on Visualizer 0 (the left visualizer), making it the active visualizer. In the renderer control for the DVR panel, click on Duplicate In, and then select Visualizer No. 1.

Then click in the visualizer window for Visualizer No. 1, and you will see (in the DVR tab) that there are two renderer instances. The first instance is the default one that was there when we created the window. Select the first one and click Delete. Now the remaining instance is the same as the DVR tab in Visualizer No. 0, and has the same transfer functions. Click the checkbox to enable Instance: 1 in visualizer 1.

Next, click on Visualizer 0, and select the variable “omz” that we previously visualized. Your VAPOR window will now show a side-by-side comparison of these two volume-rendered variables. Click on the Navigation mode button (to remove the probe’s geometric controls from the scene) and disable the probe itself, and you will see the following:

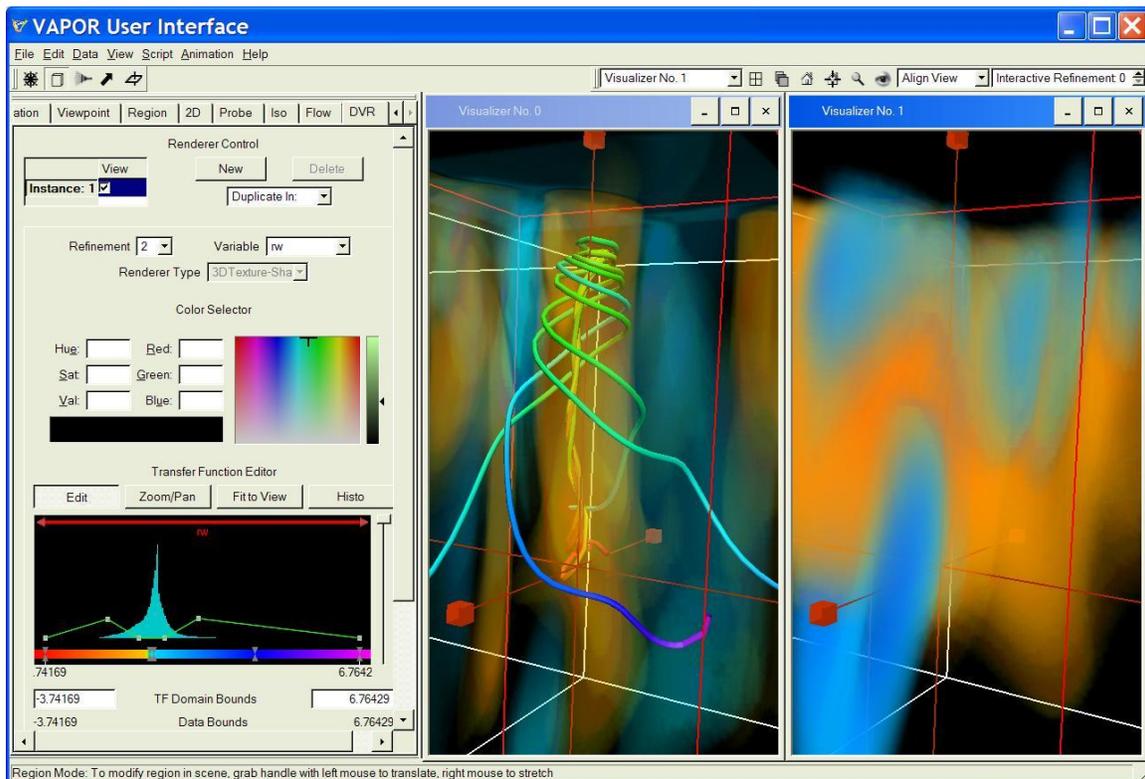


Figure 5.4: Side-by-side comparison of volume renderings of omz and rw

5.3 Using Local and Global parameters

In the side-by-side visualization of the previous section, some of the parameters in the tabbed panels are shared between the visualizer, and some are not. Animation, Region, and Viewpoint tabs each have a Scope selector indicating either Global or Local scope. By default, the Viewpoint, Animation and Region tabs are Global. If you click and drag the left mouse button in one of the visualizers, both scenes will rotate because both are using the Global viewpoint parameters. Global parameters are not available in renderer panels (i.e. DVR, Flow, and probe panels).

Select the left visualizer, click on the Viewpoint tab, and select the Scope to be Local. The viewpoint in that window will revert to the default viewpoint. Click on the “View All” button (with the eye icon) at the upper right of the main window. Then using your mouse in the left visualizer, you can position the viewpoint in the left window separately, independent of the viewpoint in the right visualizer. At any time you can select the Scope to be Global, and the left viewpoint will revert to the same (Global) viewpoint as is used in the right visualizer window.

It is sometimes useful to see the same flow lines with different volume renderings (e.g. with either omz or rw). To see the flow streamlines in both visualizers, we copy the flow instance from visualizer 0 to visualizer 1. Click on visualizer 0. On the Renderer Control at the top of the Flow tab, click Duplicate in: and select Visualizer No. 1. Then click on Visualizer 1, then select and Delete the first (preexisting) Flow instance in the Rendering Control of the flow tab. Finally click in the checkbox to enable the remaining flow instance. You will see the same flow lines in both visualizers (you may wish to reduce the opacity of the volume in Visualizer No. 1):

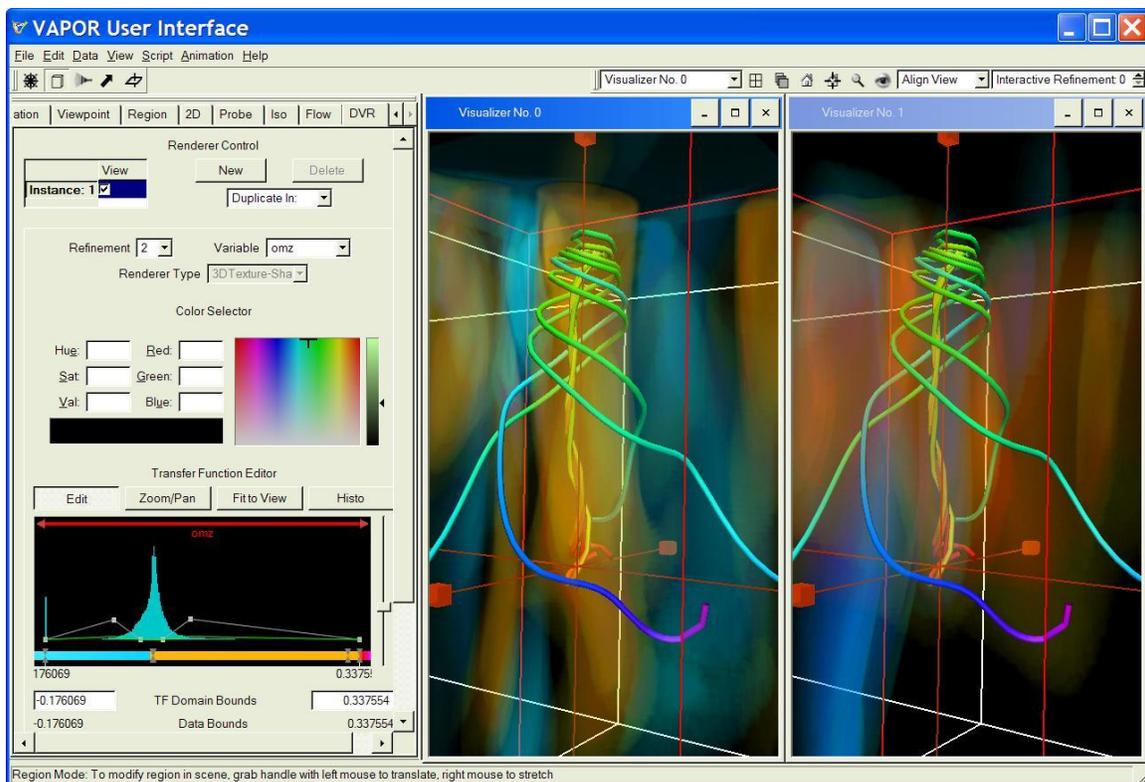


Figure 5.5: Viewing the same flow lines in different windows.

It is often useful to use Local animation control, region, or viewpoint panels, for example to view a volume from different perspectives, or to compare different time-steps side-by-side. You can always revert to the shared (Global) parameters when needed.