


☐

I'm not robot


reCAPTCHA

Continue

Paraview developer guide

This marks the beginning of a series of blog posts that I'm going to write (and I hope others also write) to give a... Read more Kitware is releasing the VTK Textbook and VTK User's Guide for PDF download. Follow the links below to a landing page where you can buy... Read More ParaView Documentation ParaView is a source and cross-platform scientific data analysis and visualization tool that enables the analysis and visualization of large datasets. ParaView is both a general-use end-user application with a distributed architecture that can be seamlessly leveraged from the desktop or other remote parallel computing resources and an extensible framework with a collection of tools and libraries for various applications, including scripts (using Python), Web view (via ParaViewWeb), or in-situ analysis (with Catalyst). ParaView leverages parallel data processing and rendering to enable interactive viewing for large datasets. It also includes support for large screens, including tiled displays and immersive 3D displays with head tracking and wand control features. ParaView also supports script and batch processing via Python. Using the included Python modules, you can write scripts that can run almost all the features exposed by the interactive application and much more. ParaView is open-source (BSD licensed, friendly commercial software). As with any successful open source project, ParaView is supported by a community of active users and developers. Contributions to both the code and the user's manual that help make the tool and documentation better are always welcome. Did you know? The ParaView project began in 2000 as a collaborative effort between Kitware Inc. Initial funding was provided by a three-year contract with the U.S. Department of Energy's ASCI Views program. The first public release, ParaView 0.6, was announced in October 2002. Independently of ParaView, Kitware began developing a web-based display system in December 2001. This project was funded by PHASE I and THE US ARL's Phase I and II SBIRs and eventually became the PVEE. PVEE has contributed significantly to the development of paraview's client/server architecture. PVEE pioneered ParaViewWeb, a modern ParaView-based web viewing solution. Since the beginning of the project, Kitware has successfully collaborated with Sandia, LANL, ARL and various other academic and government institutions to continue development. Today, the project is still strong! In September 2005, Kitware, Sandia National Labs and CSimSoft development of ParaView 3.0. This was a great effort focused on rewriting the user interface to be more user friendly and developing a quantitative analysis framework. ParaView 3.0 was released in May 2007. The user's manual has been designed as a guide for using ParaView application. It is oriented towards users who have a general understanding of common data visualization techniques. For scripting, a working knowledge of the Python language is assumed. If you don't have new to Python, there are several tutorials and guides to get you started that are available on the Internet. Did you know? In this guide, we will periodically use these Did you know? to provide additional information about the topic in question. Common Errors Common Errors blocks are used to highlight some of the common problems or complications that can occur when it comes to the topic of discussion. This guide can be divided into two volumes. Section 1 in Section 8 can be considered the user's guide, where various aspects of data analysis and visualization with ParaView are covered. Section 1 in section 11 in the manual reference part provides detailed information about the various components of the scripting UI and API. This guide tries to cover most of the features commonly used in ParaView. ParaView's flexible, pipeline-based architecture offers plenty of possibilities. If you find yourself looking for some untreated feature in this guide, refer to wiki pages (and/or feel free to ask about them on the mailing list (. ParaView is open source. The complete source code for all features described in The ParaView Guide can be downloaded from the ParaView . We also provide binary files for major platforms: Linux, Mac OS X and Windows. You can get the source and binary files for the official versions, as well as follow the active development of ParaView, downloading the night builds. Provide details about how to compile ParaView using source files outside the scope of this guide. Refer to the ParaView Wiki (for more information. View model 1.1: Input process objects A, B, and C, and/or the output of one or more data objects. Data objects represent and provide access to data; process objects operate on data. Objects A, B, and C are source, filter, and mapper objects, respectively. [SML96] visualization is the process of converting raw data into images and renderings to gain better cognitive understanding of data. ParaView uses the VTK, Visualization Toolkit, to provide the backbone for data visualization and processing. The VTK model is based on the data flow paradigm. In this paradigm, data passes through the system transformed into every step by modules known as algorithms. Algorithms can be common operations such as cutting or generating contours from the data, or they could be derived amounts of calculation, etc. Algorithms have input ports through which they take the data and output ports through which they produce the output. You need ingesting manufacturers ingesting in the system. These are simply algorithms that do not have an input port but one or more output ports. They're called sources. Readers who read data from files are examples of such sources. In addition, there are algorithms that turn data into primitive graphics so that it can be rendered on a computer screen or saved to disk in another file. These algorithms, which have one or more input ports but not output ports, are called sinks. Intermediate algorithms with input ports and output ports are called filters. Together, sources, filters, and sinks provide a flexible infrastructure where you can create complex processing pipelines by simply connecting algorithms to perform arbitrarily complex tasks. For more information about the VTK programming model, refer to [SML96]. This way of looking at the view pipeline is the basis of the ParaView workflow: you can insert data into the system by creating a reader, which is the source. You can then apply filters to extract information (for example, iso-boundaries) and render the results in a view, or to save data to disk using writers(sinks). ParaView includes readers for a multitude of file formats typically used in the world of computational science. To efficiently represent data from various fields with different characteristics, VTK provides an advanced data model used by ParaView. The data model can simply be considered as a way to represent data in memory. The different types of data will be covered in more detail in Section 3.1. Readers produce a suitable data type to represent the information contained in the files. Based on the data type, ParaView allows you to create and apply filters to transform data. You can also view data in a view to produce images or renderings. Just as there are different types of filters, each of which allows you to perform different operations and processing types, there are different types of views for generating various types of renderings, including 3D surface views, bar views and 2D lines, parallel coordinate views, and so on. Did you know? Visualization Toolkit (VTK) is an open source software system freely available for 3D computer graphics, modeling, image processing, volume rendering, scientific visualization, and information visualization. VTK also includes auxiliary support for 3D interaction widgets, two-dimensional and three-dimensional annotation, and parallel computing. Inside, VTK is implemented as a C toolkit, which requires users to build applications by combining various objects an application. The system also supports automatic wrapping of the core of C, C++, in Python, Java, and Tcl so that VTK applications can also be written using these interpreted programming languages. VTK is used worldwide in commercial applications, research and development, and as the basis of many advanced display applications such as ParaView, VisIt, VisTrails, VisTrails, MayaVi and OsiriX. ParaView comes with several executables that serve different purposes. This is paraview's primary graphical user interface (GUI). In most cases, when we refer to ParaView, we are actually talking about this application. It is a cross-platform Qt-based user interface that provides access to God processing capabilities. The main parts of this guide are dedicated to understanding and using this application. pypython is the Python interpreter that runs ParaView Python scripts. You can think of this as the paraview equivalent for scripting. Similar to pypython, pvbatch is also a Python interpreter running Python scripts for ParaView. The only difference is that while pypython is used to run interactive scripts, pvbatch is designed for batch processing. In addition, when running on processing resources with MPI functionality, pvbatch can run in parallel. We will cover this in more detail in Section 6.9. For remote viewing, this executable represents the server that performs all data processing and, potentially, rendering. You can have paraview connect to pvservers running remotely on an HPC resource. This allows you to create and control the visualization and analysis of the HPC resource from the desktop as if it were simply processing it locally on the desktop! These can be thought of as the pvsrver divided into two separate executables: one for the data processing part, pvdataserver, and one for the rendering part, pvrenderserver. Splitting these into separate processes allows you to process and render data on separate sets of nodes with appropriate processing capabilities suitable for the two tasks. As with pvsrver, paraview can connect to a pvdataserver- pvrenderserver pair for remote viewing. Unless otherwise noted, all discussions about remote view or client-server view in this guide apply to pvsrver and pvdataserver - pvrenderserver configurations. paraview is the graphical front end for the ParaView application. The user interface is designed to allow you to easily create pipelines for processing data with arbitrary complexity. The UI provides panels to examine and modify pipelines, to modify parameters that in turn affect processing pipelines, to perform various data selection and inspection actions for data introspection, and to generate renderings. Various aspects of the user interface will be covered for the best part of this guide. Let's start by looking at the various components of the user interface. If you run paraview for the first time, you will see something fig. 1.2. The UI consists of menus, anchorable panels, toolbars, and a window, which is the middle of the application window. Fig. 1.2 paraview application window. application, (File menu), undo/redo (Edit menu), to toggle panel, and toolbars (View menu). In addition, menus provide ways to create sources that generate test datasets of various types (Sources menu), as well as new slicers for data processing (Filters menu). The Tools menu gives you access to some of paraview's advanced features, such as managing plugins and favorites. The panels offer the ability to view the status of the application. For example, you can examine the set display pipeline (Pipeline Browser), as well as the memory used (Memory Control) and the parameters or properties for a processing module (Properties panel). Many of the panels also allow you to change the displayed values, such as the Properties panel not only shows the parameters of the processing module, but also allows you to edit them. Many panels are sensitive to context. For example, the Properties panel changes to show the parameters of the selected form or view. The viewport or the middle of the paraview window is the area where ParaView renders data-generated results. Containers in which data can be rendered or displayed are called views. You can create different types of views, all arranged in this window area. By default, a 3D view is created, which is one of the most commonly used views in ParaView. To better understand how to use the application interface, consider a simple example: creating a data source and applying a filter. The display process in ParaView begins by bringing the data into the application. Section 2 explains how to read data from various file formats. In addition to reading files to insert data into your application, ParaView also provides a collection of data sources that can produce sample datasets. These are available in the Sources menu. To create a source, simply click any item on the Sources menu. Did you know? when you move the cursor over items on any menu on most platforms (except mac os x) a brief description of the item appears in the status bar in the lower-left corner of the window For example, clicking > Sphere will create a producer algorithm that generates a spherical surface, as shown in Figure 1.3. Figure 1.3 View in paravista: Step 1. - Some things to note: A pipeline module is added to the Pipeline Browser panel with a name derived from the menu item, as it is The Properties panel fills with text to indicate that it shows the properties of the highlighted item (which, in this case, is Sphere1), as well as to display some widgets for parameters such as Center , Radius , and so on. In the Properties panel, the Apply button is activated and highlighted. The 3D view remains unchanged because nothing new is yet displayed or rendered in this view. Let's take a closer look at what happened. When we clicked Sources > Sphere , referring to Section 1.2, we instantiated a source that could produce a spherical surface mesh, which is reflected in the Piping Browser. This instance receives a name, used by Sphere1 and the Browser pipeline , as well as other UI components, to reference this instance of the source. Pipeline modules, such as sources and filters, have parameters that you can modify that affect form behavior. We call them properties. The Properties panel shows these properties and allows you to edit them. Because inserting data into the system can be a time-consuming process, paraview allows you to change properties before the form runs or actually processed to insert data. As a result, the Apply button is highlighted to indicate that you must accept the properties before the application proceeds. Since no data has yet entered the system, there is nothing to show. Therefore, the 3D view remains unchanged. Suppose the default values for all Sphere1 properties are correct. Next, click the Apply button. Let's say we agree with the default values for all properties on Sphere1. Next, click the Apply button. Figure 1.4 View in paraview: Step 2. A spherical surface is rendered in the 3D view. The View section of the Properties panel now shows new parameters or properties. Some toolbars are updated, and you can see that toolbars with text, such as Solid Color and Surface , are now turned on. When we click Apply , we told Paraview to apply the properties displayed in the Properties panel. When a new source (or filter) is first applied, paraview will automatically show the data that the pipeline module produces in the current view, if possible. In this case, the sphere source produces a surface mesh, which is then displayed or displayed in the 3D view. Properties that control how data is displayed in the view are now displayed in the Properties panel in the View section. Things like color surface, rendering type or representation, shadow parameters, etc., are shown in this newly updated section. We'll look at the display properties in more detail in Section 4. Some of the commonly used properties are also on the toolbar. These properties include the data array with which the surface is colored and the type of representation. These are the changes on the toolbar that allow you to quickly change some display properties. If you change one of the properties of the sphere source, such as properties in the Properties section of the Properties panel, including radius for the spherical mesh or its center, the Apply button is re-highlighted. When you have completed all property changes, you can press Apply to apply the changes. Once the changes are applied, paraview will re-run the source of the sphere to produce a new mesh, as required. The view will then be automatically updated and the new result will be rendered. If you change one of the display properties for the sphere source, such as properties in the View section of the Properties panel (including Representation or Opacity), the Apply button does not change, changes are applied immediately, and the view is updated. The logic behind this is that, typically, running the source (or filter) is more computationally intensive than rendering. Changing the (or filter) properties causes the algorithm to run again, while changing the display properties, in most cases, triggers only a new rendering with an updated graphic state. Did you know? For some workflows with smaller data sizes, it can be more convenient if the Apply button has been applied automatically even after you make changes to the pipeline form properties. You can change this setting from the application settings dialog box, which is accessible from the > settings menu. The setting is named Auto App status by using the button from the toolbar. Based on the data flow paradigm, you create pipelines with filters to transform data. Similar to the Sources menu, which allows you to create new data sources, there is a Filters menu that provides access to the large set of filters available in ParaView. If you will perch the items in this menu, some of them will be enabled and some of them will be disabled. Filters that can work with the data type produced by the sphere source are enabled, while others are disabled. You can click one of the filters enabled to create a new instance of that filter type. Did you know? To understand why a particular filter doesn't work with the current source, simply move your mouse over the disabled item on the Filters menu. On Linux and Windows (not OS X, the status bar will provide a brief explanation of why such a filter is not available. For example, if you click Filters > Reduce , a filter is created that reduces each of the mesh cells by a fixed factor. Just like before, when we created the origin of the sphere, we see that the newly created filter is given a new name, Shrink1 , and is highlighted in the Browser. The Properties panel is also updated to show the properties of the new filter, and the Apply button is highlighted to require you to accept properties for the filter so that it can run and the result can be rendered. Clicking back and forth between Sphere1 and Shrink1 in the Pipeline Browser displays the Properties panel and updates the toolbars, reflecting the status of the selected pipeline module. This is an important concept in ParaView. There is a notion of an active pipeline module, called an active source. Several panels, toolbars, and menus will be updated based on the active source. If you click Apply , as before, the reduction filter runs and the resulting dataset is generated and displayed in the 3D view. Paraview will also automatically hide the result from Sphere1 so that it is not displayed in the view. Otherwise, the two datasets will overlap. This is reflected in the state change for eyeball icons in the pipeline browser next to each of the pipeline modules. You can show or hide results from any pipeline module by clicking the eyeballs. This simple workflow forms the basis of all data analysis and visualization in ParaView. The process involves creating sources and filters, changing their parameters, and displaying the result generated in one or more views. In the rest of this guide, you'll cover various types of filters and data processing that you can run. You will also cover different types of views that allow you to produce a wide range of 2D and 3D views, as well as inspect data and drill down. Common errors Beginners often forget to press the Apply button after creating sources or filters or changing properties. This is one of the most common pitfalls for new users to the ParaView workflow. While this section refers to pypython, everything we discuss here is also applicable to pvbatch. Until we start searching for parallel processing, the only difference between the two executables is that pypython provides an interactive shell where you can type commands, while pvbatch expects the Python script to be specified in the command-line argument. ParaView provides a scripting interface to write scripts to perform tasks that can be performed using the GUI. The scripting interface is accessible via Python, which is an interpreted programming language popular with the scientific community for its simplicity and capabilities. While a working knowledge of Python will be useful for writing scripts with advanced features, should be able to follow most of the discussion in this book about ParaView scripting even without much Python exposure. ParaView provides a paraview package with several Python modules that expose various features. The main scripting interface is provided by the simple module. When you pypython, you should see a prompt in a terminal window as follows (with some platform-specific differences). Python 2.7.5 (default, Sep 2, 2013, 05:24:04) [GCC 4.2.1 Apple LLVM 5.0 Compatible (clang-500.6.68)] on darwin Type help, copyright, credits, or license for more information >>> You can now type commands at this prompt and ParaView will execute them. To import the ParaView scripting API, you must first import the simple form from the paraview package as follows: >>> from paraview.simple import - Common errors Remember to press the Enter or Return key after each command to run it. Any Python interpreter will not execute the command until Enter is hit. If the module loads successfully, pypython will have a prompt for the next command. >>> from paraview.simple import - >>> You can treat this as in the same state as when paraview started (with some differences that we can ignore for now). The application is ready to enter data and start processing. Let's try to understand the workflow by looking at the same use case that we did in Section 1.4.2. In paraview, you created the data source by using the Sources menu. In the scripting environment, this is mapped by simply typing the name of the source to create. The sphere source will be created with a default property set. Just like with paraview, as soon as a new pipeline module is created, it becomes the active source. Now, to display the active source in a view, try: the Show call will prepare the view, while the Render call will cause rendering. In addition, a new window will be displayed, showing the result (Fig. 1.6). This is similar to the state after tapping Apply in ui. Window 1.6 showing the python code result. Set a single property in the active source. >>> SetProperties(Radius=1.0)—You can also set multiple properties. >>> SetProperties(Center,[1, 0, 0], StartTheta=100) Similar to the Properties panel, SetProperties affects the active source. To query the current value of any property in the active source, use GetProperty . >>> >>> radius - GetProperty(Radius) >>> print(radius) 1.0 >>> >>> center - GetProperty(Center) >>> print(Center) [1.0, 0.0, 0.0] The SetProperties and GetProperty functions work the same function as the Properties panel properties section, which allow you to set and introspect pipeline module properties for the active source. Similarly, for the Display section of the panel or display properties, we have the functions and GetDisplayProperty . >>> >>> SetDisplayProperties(Opacity=0.5) >>> GetDisplayProperty(Opacity) 0.5 Common errors Note that property names for the SetProperties and SetDisplayProperties functions are not enclosed in double quotation marks, and property names for GetProperty and methods are. In paraview, each time you press Apply or Edit a display property, the UI automatically renders the view. In the scripting environment, you must do this manually by calling the Render function each time you want to render again and examine the updated result. fime: we lack blurb on camera reset. Similar to creating a source, to apply a filter, simply create the filter by name. Create the 'Shrink' filter and link it to the active source, which is the 'Sphere' instance. >>> Shrink() - As soon as the Shrink filter is created, it will now become the new active source. All methods that act on the active source now act on this filter instance, and not on the Previously Created Sphere instance. : Displays the resulting data and renders it. >>> Show() >>> Show() >>> Render() If you tried the script above, you'll notice that the result isn't exactly what we expected. For some reason, the cells have retracted are not visible. This is because we lost a phase: in paraview, the UI was smart enough to automatically hide the input dataset for the newly created filter after reaching the application. In the scripting interface, such operations are the responsibility of the user. We should have hidden the source of the sphere from view. You can use the Hide method, the counter part of Show , to hide the active source. But now we have a problem : when we created the reduction filter, we modified the active source to be the compaction instance. Fortunately, all the features we've discussed so far may require a first optional topic, which is the source or instance of the filter to operate on. If specified, this instance is used in place of the active source. The solution is as follows: - Get the input property for the active source, such as input for reduction. >>> shrinkInput - GetProperty(Input) - This is in fact the instance of the sphere that we created earlier. >>> print(shrinkInput)—Hide the instance of the sphere in an explicit <<paraview.servermanager.Sphere object= at= 0x11d731e90=>>explicit. >>> Hide(shrinkInput)—Render the result again. >>> Render() Alternatively, you can also get/set the active source by using the GetActiveSource and SetActiveSource functions. >>> shrinkInstance - GetActiveSource() >>> print(shrinkInstance) - Get the input property for the active source, which is input. >>> sphereInstance - GetProperty(Input) - This is actually the instance previously created sphere. >>> print(sphereInstance) - Change the active source to sphere and hide it. >>> SetActiveSource(sphereInstance) >>> Hide() - Now restores the active source <<paraview.servermanager.Sphere object= at= 0x11d731e90=>>at= 0x11d731e90=>>on the compaction instance. >>> SetActiveSource(shrinkInstance) - Render the result >>> UpdatePipeline(proxy=sphere)—Let's check the limits again. >>> sphere. GetDataInformation(). GetBounds() (-0.48746395111083984, 0.48746395111083984, -0.48746395111083984, 0.48746395111083984, -0.5, 0.5) - If we call UpdatePipeline() again, this will have no effect since the pipeline has not been modified, so there is no need to re-run. >>> UpdatePipeline() <&>>> sphere. GetDataInformation(). GetBounds() (-0.48746395111083984, 0.48746395111083984, -0.48746395111083984, 0.48746395111083984, -0.5, 0.5) - Now let's change a property. >>> sphere. Radius = 10 - Limits do not change because the pipeline has not been run back. >>> sphere. GetDataInformation(). GetBounds() (-0.48746395111083984, 0.48746395111083984, -0.48746395111083984, 0.48746395111083984, -0.5, 0.5) - Let's update and see. >>> UpdatePipeline() >>> sphere. GetDataInformation(). GetBounds() (-9.74927902216797, 9.74927902216797, -9.74927902216797, 9.74927902216797, -10.0, 10.0) We will examine the sphere. GetDataInformation API in Section 3.3 in more detail. For time datasets, UpdatePipeline accepts a time argument, which is the time for the pipeline to be updated. To upgrade to 10.0: >>> UpdatePipeline(10.0) - Alternate way to do the same: >>> UpdatePipeline(time=10.0) - If you are not using the active source: >>> UpdatePipeline(10.0, source) >>> UpdatePipeline(time=10.0, proxy=source) The paraview application also provides access to an internal shell, where python commands and scripts can be entered just as with pypython. To access the Python shell in the GUI, use the > Python Shell menu option. A dialog box appears with a prompt exactly like pypython. You can try inserting commands from the previous section into this shell. As you type each of the commands, you will see the UI update after each command, such as when you create the source instance of the sphere, it will appear in the Pipeline Browser . If you change the active source, the pipeline browser and other UI components will be updated to reflect the change. If you change the display properties or properties, the Properties panel is updated to reflect the change as well. Fig. 1.8 Python Shell in paraview provides access to scripting. Python shell in paraview supports completion for instance functions and methods. Try pressing the Tab key after partially typing any commands (as shown in Fig. 1.8). This guide provides a fair overview of the ParaView Python API. However, there will be cases where you just want to know how to complete a particular action or sequence of actions that you can perform with the GUI using a Python script instead. To do this, paraview supports actions in the UI such as Python scripts. Simply start the track by clicking Tools > Start Track. paraview now enters a mode where all your actions (or at least those relevant to scripting) are monitored. Each time you create a source or filter, open data files, change properties, and hit Apply , interact with the 3D scene, or save screens, and so on, actions will be monitored. Once you're done with the series of actions you want to create a script, click Tools > Stop Track. paraview then displays an editor window with the generated track. This will be the equivalent of the Python script for the actions performed. You can now save it as a script to use for batch processing. © Copyright 2020, ParaView Developers Revision eb7609ae. Built with Sphinx using a theme provided by Read Documents. Documents.

