



13th Summer School on **SCIENTIFIC VISUALIZATION**

Paraview for large data visualization

Ivan Spisso - [i.spisso@cineca.it](mailto:i.spisso@ Cineca.it)
SuperComputing Applications and Innovation Department



OUTLINE

- Paraview architecture details
- Large data visualization
- Filters and data explosion
- Rendering of large data
- Client-server mode, example of cluster use



Paraview architecture details

http://www.paraview.org/Wiki/Users_Guide_Client-Server_Visualization

ParaView is designed as a three-tier client-server architecture.

The three logical units of ParaView are as follows:

Data Server: is responsible for data reading/filtering/writing;

All of the pipeline objects seen in the pipeline browser are contained in the data server.

The data server can be parallel.

Render Server: is the unit responsible for rendering.

The render server can also be parallel, in which case built in parallel rendering is also enabled.

Client: The unit responsible for establishing visualization.

The client controls the object creation, execution, and destruction in the servers, but does not contain any of the data.

If there is a GUI, that is also in the client.

The client is always a serial application.



Parallel architecture

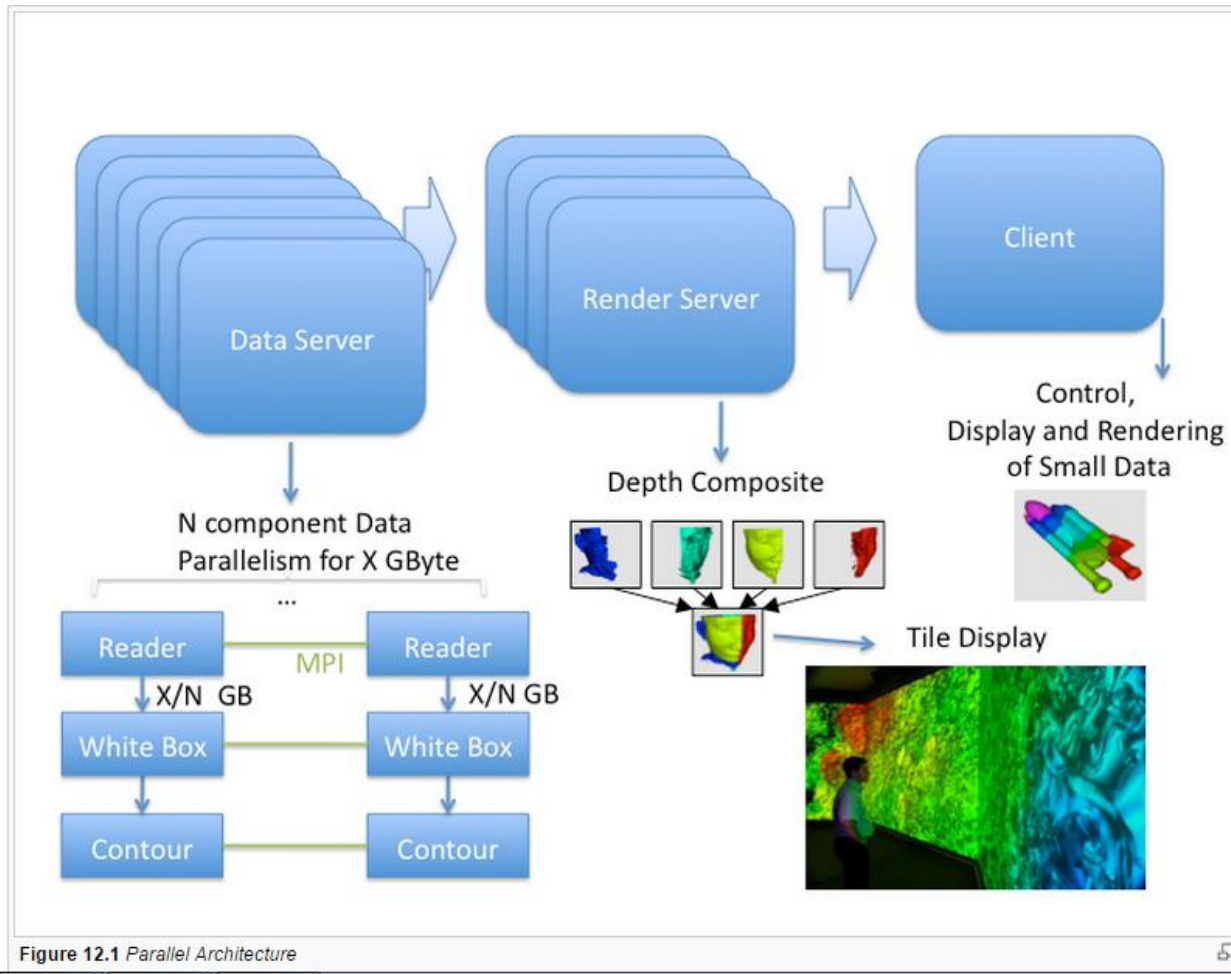
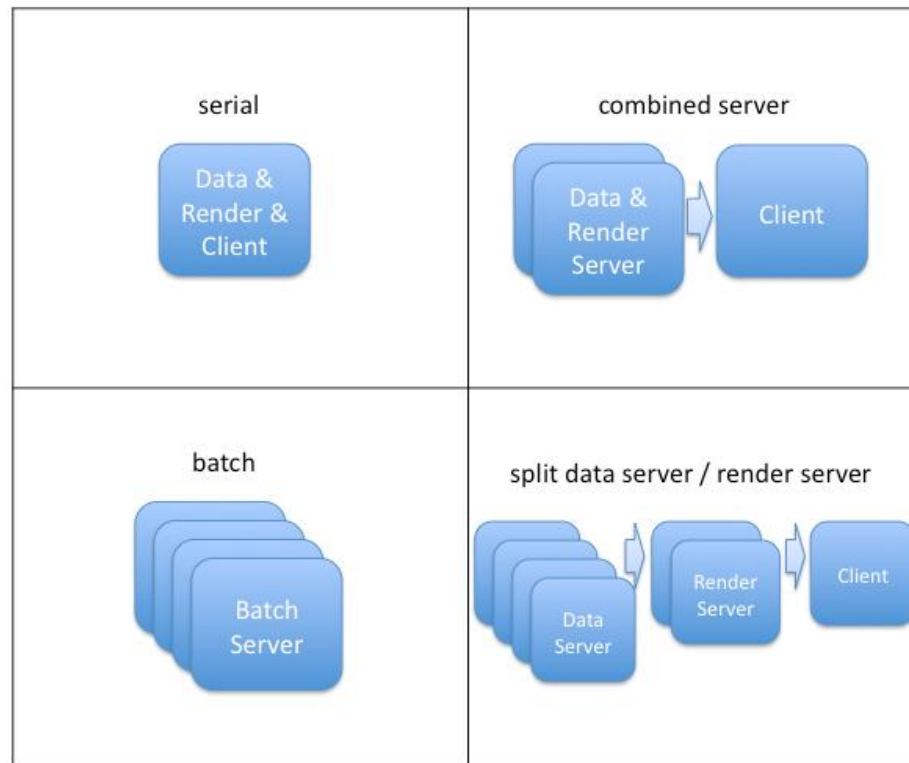


Figure 12.1 Parallel Architecture



Common Configurations of the logical components

The three logical components can be combined in various different configurations. When ParaView is started, the client is connected to what is called the built-in server; in this case, all three components exist within the same process. Alternatively, you can run the server as an independent program and connect a remote client to it, or run the server as a standalone parallel batch program without a GUI. In this case, the server process contains both the data and render server components. The server can also be started as two separate programs: one for the data server and one for the render server.



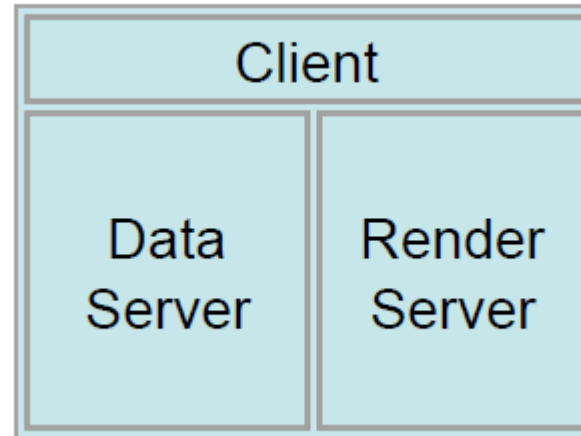


Paraview architecture details

Standalone mode

In standalone mode, the client, data server, and render server are all combined into a single serial application.

When you run the ParaView application, you are automatically connected to a built-in server so that you are ready to use the full features of ParaView.



OUR CASE UP TO NOW




Paraview architecture details

Client-server mode. In client-server mode, you execute the *pvserver* program on a parallel machine and connect to it with the ParaView client application.

The *pvserver* program has both the data server and render server embedded in it, so both data processing and rendering take place there.

The client and server are connected via a socket, which is assumed to be a relatively slow mode of communication, so data transfer over this socket is minimized.

**Example of
implementation on PLX
cluster, later on**






Paraview architecture details

Client-render server-data server mode. In this mode, all three logical units are running in separate programs. The client is connected to the render server via a single socket connection.

The render server and data server are connected by many socket connections, one for each process in the render server. Data transfer over the sockets is minimized.

Not recommend since it born to take advantage of heterogeneous environments where one might have a large, powerful computational platform and a second smaller parallel machine with graphics hardware in it.

In practice benefits are almost always outstripped by the time it takes to move geometry from the data server to the render server.






Large data visualization

- Loose coupling between components → flexible framework
- Drawback → larger memory footprint
- Why? Because each stage of this pipeline maintains its own copy of the data. Whenever possible, ParaView performs **shallow copies** of the data so that different stages of the pipeline point to the same block of data in memory. However, any filter that creates new data or changes the values or topology of the data must allocate new memory for the result. If ParaView is filtering a very large mesh, inappropriate use of filters can quickly fill in all available memory.
- **When visualizing large data sets, it is important to understand the memory requirements of filters.**






Filters and data explosion

- Append Datasets
- Append Geometry
- Clean
- Clean to Grid
- Connectivity
- D3
- Delaunay 2D/3D
- Extract Edges
- Linear Extrusion
- Loop Subdivision
- Reflect
- Rotational Extrusion
- Shrink
- Smooth
- Subdivide
- Tessellate
- Tetrahedralize
- Triangle Strips
- Triangulate





Filters and data explosion

- Clip 
- Decimate
- Extract Cells by Region
- Extract Selection 
- Quadric Clustering
- Threshold 

Input data



filter

Output: unstructured
data of the reduced
magnitude respect to
input data



Filters and data explosion

- Cell Centers
- Contour 
- Extract CTH Fragments
- Extract CTH Parts
- Extract Surface
- Feature Edges
- Mask Points
- Outline (curvilinear)
- Slice 
- Stream Tracer 



Input data


filter

Output: unstructured
data of the reduced
magnitude respect to
input data










Filters and data explosion

- Block Scalars
- Calculator 
- Cell Data to Point Data
- Curvature
- Elevation
- Generate Surface Normals
- Gradient
- Level Scalars
- Median
- Mesh Quality
- Octree Depth Limit
- Octree Depth Scalars
- Point Data to Cell Data
- Process Id Scalars
- Python Calculator
- Random Vectors
- Resample with dataset
- Surface Flow
- Surface Vectors
- Texture Map to...
- Transform
- Warp (scalar)
- Warp (vector) 





Filters and data explosion

- Annotate Time
- Append Attributes
- Extract Block
- Extract Datasets
- Extract Level 
- Glyph 
- Group Datasets 
- Histogram 
- Integrate Variables
- Normal Glyphs
- Outline
- Outline Corners
- Plot Global Variables Over Time
- Plot Over Line 
- Plot Selection Over Time 
- Probe Location 
- Temporal Shift Scale
- Temporal Snap-to-Time-Steps
- Temporal Statistics

Input data

filter

Output: no data or
shallow copy of input
data



Cutting out large data

how to and recommendations

- Extract iso-surfaces from a volume using the Contour filter;
- Perform a Slicing over the data
- Clipping
- Threshold
- Extract Selection
- Extract Subset

NOTE: all of these filters with the exception of Extract Subset will convert structured data types to unstructured grids. Therefore, they should not be used unless the extracted cells are of at least an order of magnitude less than the source data.

- When possible, replace the use of a filter that extracts 3D data with one that will extract 2D surfaces. For example, if you are interested in a plane through the data, use the Slice filter rather than the Clip filter.
- If you are interested in knowing the location of a region of cells containing a particular range of values, consider using the Contour filter to generate surfaces at the ends of the range rather than extract all of the cells with the Threshold filter.
- Be aware that substituting filters can have an effect on downstream filters. For example, running the Histogram filter after Threshold will have an entirely different effect than running it after the roughly equivalent Contour filter.



Rendering of large data

Rendering is the process of synthesizing the images that you see based on the dataset.

The ability to effectively interact with your data depends highly on the speed of the rendering. The speed of rendering is proportional to the amount of data being rendered.

The interactive render is a compromise between speed and accuracy.

ParaView supports two modes of rendering that are automatically clipped as necessary.

Mode-1: **still render**, the data is rendered at the highest level of detail, when the user is not interacting with the data. This rendering mode ensures that all of the data is represented accurately. At any time when interaction of the 3D view is not taking place, ParaView uses a still render so that the full detail of the data is available as you study it.

Mode-2: **interactive render**, speed takes precedence over accuracy. This rendering mode endeavors to provide a quick rendering rate regardless of data size.

As you drag your mouse in a 3D view to move the data, you may see an approximate rendering while you are moving the mouse, but the full detail will be presented as soon as you release the mouse button.



Client-server mode, configuration server side

1. Inspect the `pw_pvserver.sh` script
2. `qsub pw_pvserver.sh`
3. `qstat -n`
4. Take note of the node(s) where your pvserver job is running

```
ispisso0@node097:~ - Shell - Konsole
Session Edit View Bookmarks Settings Help

1 #!/bin/bash
2 #PBS -l walltime=1:00:00
3 #PBS -l select=2:ncpus=12:mpiprocs=12:ngpus=2
4 #PBS -q R1909052
5 #PBS -A cin_staff
6 module load profile/advanced
7 module load autoload paraview/4.1.0-mpiserver
8 cd $PBS_O_WORKDIR
9
10 echo "Running on " `hostname`
11 echo Available resource are `cat $PBS_NODEFILE`
12 echo "Job started at `date` on nodes: `cat $PBS_NODEFILE` "
13
14 NPROCS=`wc -l < $PBS_NODEFILE`
15 export DISPLAY=:0.1
16 echo 'NPROCS is' $NPROCS
17 ##mpirun -bynode -np 12 pvserver -display :0.1 -np 12 pvserver -display :0.2
18 mpirun -bynode -np 12 pvserver : -np 12 pvserver
19 echo "Job finished at `date`"

[ispisso0@node097 ~]$ qstat -n
node351.plx.cineca.it:
Job ID          Username Queue   Jobname   SessID  NDS  TSK  Req'd  Req'd  Elap
-----
1909601.node351 ispisso0 rvn_visu pw_pvserve  --    1   2    --   00:20 Q   --
--
1912565.node351 ispisso0 visual   ispisso0-p 23376  1   1    --   12:00 R 03:35
node097ib0/6
1912910.node351 ispisso0 R1909052 pw_pvserve 13028  2  24    --   01:00 R 00:04
node103ib0/0*12+node102ib0/0*12
```

X Connections

One of the most common problems people have with setting up the ParaView server is allowing the server processes to open windows on the graphics card on each process's node.

Multiple GPUs Per Node

It is becoming commonplace to put multiple GPUs on each node in a cluster. Taking advantage of these multiple GPUs can be tricky.

Typically, each of these GPUs will have its own display. For example, if you have two GPUs on a node, they are probably referenced by the displays localhost:0.0 and localhost:0.1.



Connecting to a server configuration client side

1. Open Paraview with RCM on the visual queue
2. Click the “connect” button
3. Edit Server Configuration. Insert Host from the client side (The first one from the list)
4. Click configure
5. Click Connect

The screenshot shows the Paraview interface with the 'Connect' button in the top toolbar circled in red. Two dialog boxes are overlaid on the interface:

Choose Server Configuration

Configuration	Server
100	cs://node088ib0:11111
pippo	cs://node364ib0:11111
rvn	cs://node364ib0:11111

Edit Server Configuration

Name: 100
Server Type: Client / Server
Host: node088ib0
Port: 11111

Client-server mode



Summer School on SCIENTIFIC VISUALIZATION

TurboVNC: node097:6 (ispisso0) [Tight + JPEG 1X Q95]

Layout #1 x +

cs://node103ib0:11111
motorBike_500.vtk

Information Properties

Properties

Apply Reset Delete ?

Search ... (use Esc to clear text)

Properties (motorBike_500.vtk)

Display (UnstructuredGridRepres

Representation Surface

Coloring

U Magnitude

Show Edit Rescale

Styling

Opacity 0.31

Cube Axes

Show Axis Edit

3D

client

node097	System Total	14.895 GiB	31.74%
	paraview	200.37 MiB	0.42%

server

node103	System Total	5.51 GiB	11.69%
	pvsrver	4.17 GiB	8.86%
node102	System Total	4.19 GiB	11.02%
	pvsrver	4.07 GiB	8.63%

Selection Display Inspector

Cell Labels

Point Labels

Selection Color

common_tools - Konqueror ParaView 4.1.0 64-bit

ispisso0@node097: ~ - She

16:34
Tuesda
2014-10



Client-server mode

Display by the *vtkProcessID*

The screenshot shows the ParaView 4.10.0 client interface. The main 3D view displays a motorbike model colored by the *vtkProcessID* variable. The color scale ranges from 0 (blue) to 23 (red). The Pipeline Browser shows the data source *motorBike_500.vtk* and the *ProcessIdScalars1* filter. The Properties panel shows the *Display (UnstructuredGridRepresentation)* and *Coloring* options. The Memory Inspector shows the system memory usage for the client and server nodes.

Node	System Total	paraview
node097	4.99 GiB 31.81%	239.65 MiB 0.50%
node103	5.98 GiB 12.68%	4.66 GiB 9.88%
node102	5.67 GiB 12.02%	4.55 GiB 9.65%

The screenshot shows the *Find Data* dialog box in ParaView. The *Create Selection* tab is active, showing search criteria for *Process ID* 1. The *Current Selection* table lists the selected cells.

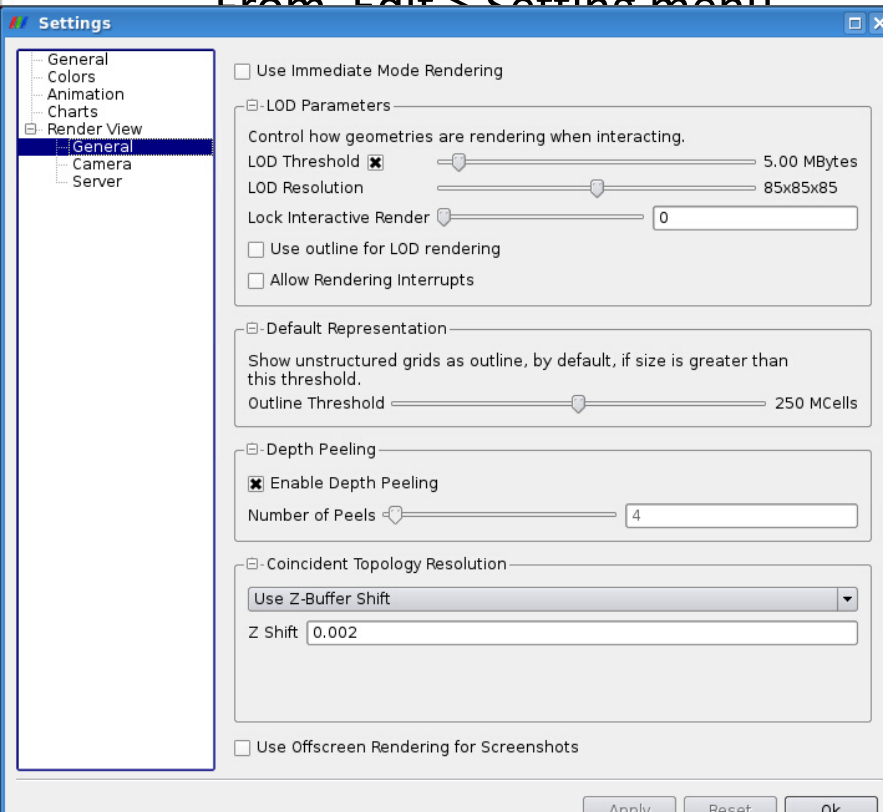
Cell Type	Process ID
2042390 Pyramid	23
2042391 Pyramid	23
2042392 Pyramid	23
2042393 Tetrahedron	23
2042394 Tetrahedron	23
2042395 Pyramid	23
2042396 Tetrahedron	23
2042397 Tetrahedron	23



Bonus Slides

Rendering basic parametric settings

From Edit > Setting menu





Basic parameters settings

Use Immediate Mode Rendering When checked, geometry is sent to the graphics card for immediate rendering. When unchecked, the geometry is first compiled into display lists for more efficient rendering. The display lists usually render faster, but require initial time to compile during the first frame and extra memory to store.

LOD Threshold Controls when to replace the geometry with a decimated version of the geometry during interactive rendering. The checkbox turns the feature on or off. When on, the slider gives a threshold for the feature. If the geometry size is below the threshold, it is considered small enough to render. When the geometry size is above the threshold, the decimated form is used during rendering. When unchecked, the full geometry is always rendered.

LOD Resolution Controls the size of geometry to create for geometric level of detail. Moving the slider to the right results in a more coarse representation.

Lock Interactive Render Specify a pause between an interactive render and a still render. This pause allows you to let go of the mouse and drag again (to perhaps change from rotate to pan) without having to wait for a full still render.

Use outline for LOD rendering Normally ParaView uses a decimated version of the geometry for its interactive render. When this box is checked, ParaView instead uses a simple bounding box outline. This saves the time to make the decimated geometry and any rendering time.

Allow Rendering Interrupts When checked, a still render may be interrupted by a request to perform an interactive render. Not all rendering modes support interrupts.



Basic parameters settings

Outline Threshold When an unstructured dataset exceeds this threshold, the default representation mode is set to outline instead of surface.

Enable Depth Peeling ParaView uses an algorithm called depth peeling to properly render translucent surfaces. With it, the top surface is rendered and then “peeled away” so that the next lower surface can be rendered and so on. Can be expensive: try shutting off depth peeling or adjust the number of depth peels used. Using more peels → more depth complexity → slower (less peels → faster).

Coincident Topology Resolution It is sometimes the case you wish to render lines that are coincident with a surface (for example, rendering a wireframe of cells on top of their surface). Rendering systems generally can have problems resolving hidden surfaces in these circumstances. These options allow you to change the “tricks” ParaView's rendering uses to resolve these issues. Generally you do not need to change these options unless there is a problem with the rendering.

Use Offscreen Rendering for Screenshots ParaView usually uses offscreen rendering when creating screenshots to avoid having other windows on your desktop affect the images. However, some hardware places restrictions on offscreen rendering that can cause artifacts. If you find that your saved images look different than the ones on screen (for example, more dark), try disabling this feature.