Using Remote Visualization Capabilities at CINES

Nicolas Monnier, Yannick Sommerer
CERFACS

Originally Published in La Gazette du CINES, Number 15-16, April – July 2003
(Translated and reproduced with the approval of CINES)

1. Background

The continual increase in processing power available at national computing centers allows the creation of more complex numerical simulations and more accurate results. However, increases in complexity and accuracy can lead to huge result files that are difficult to exploit. The management of many large pre and post processing files has become a major inhibitor that slows the development of more realistic numerical simulations. Not only has it become costly to transfer files between a laboratory and the computing facility, but the visualization systems required to study such large volumes of data are too expensive for most research laboratories.

The CERFACS research project, lead by Yannick Sommerer, “Extreme computing for turbulent combustion” illustrates these facts. Coupling multi-physics is one of the current areas of research for turbulent combustion where thermal-acoustic phenomena and fluid-structure interactions play a fundamental role in flow stability within combustion chambers. In order to correctly predict the level of noise and pollution generated by turbines, we must simulate large-scale problems on high-resolution meshes and take into account as much as possible the flow physics both before and after the combustion chamber.

Though the CINES computers can process a numerical simulation with a 10^7 point mesh in a reasonable time (which is already a step beyond traditional capabilities in this domain), the post-processing of the generated data is a computing challenge on its own.

- Such a large simulation cannot be done in one step. The computation must be regularly checked to make certain that it is converging properly. Intermediate results are regularly post-processed to validate the results with about ten 80MB files required for each iteration. This results in a transfer of over one half a gigabyte of data for each intermediate check.
- At the end of the simulation, over one hundred result files of 80MB each (the result of ten intermediate phases) files have to be processed through a powerful visualization server.

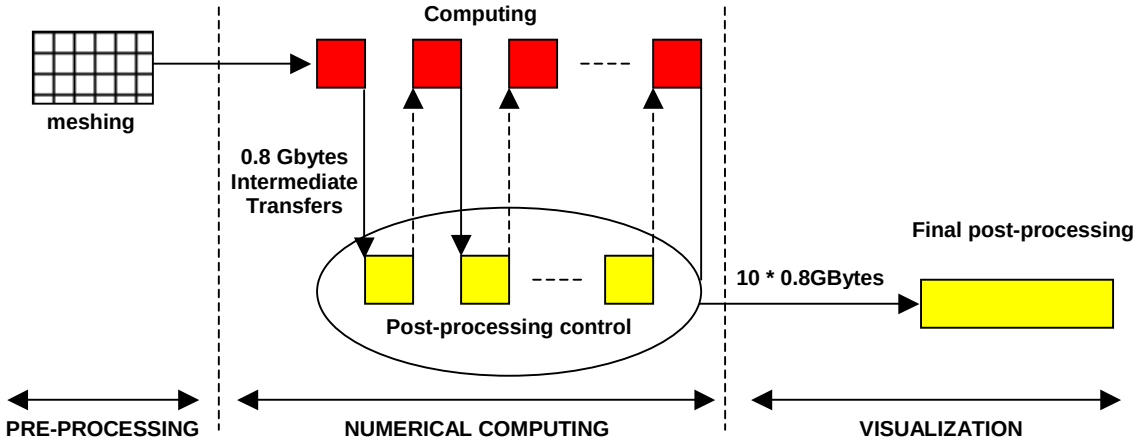

With a sustained 400Kb/s[1] network bandwidth between CERFACS and CINES, data download was

[1] 400 Kbits per second is the average rate between CINES and CERFACS at the time of the tests. This rate can be higher or lower according to the use of RENATER and REMIP bandwidth. Under ideal circumstances, a transfer rate of 3.9 Mbits/s can be observed.

the main bottleneck for this numerical simulation. The download time would have been about 4.5 hours for each post -processing check, and about 45 hours for the final data set to be exploited!

Following the event “Visualization and Large Scale Computing” organized by CINES on October 21st 2002, we were interested in evaluating the features of OpenGL Vizserver™ computing solution and its ability to solve our problem. In summary, visualization was done on the Silicon Graphics® Onyx2® visualization system located at CINES and the graphical results were sent across the network and using OpenGL Vizserver and remotely displayed on a client located at CERFACS, more than 200 Km away. In this way, post-processing can be done through the network without transferring any data sets but only images which can be compressed using different techniques to fit the available network bandwidth.

2. Technical Approach

The tests previously carried out by CINES (OpenGL Vizserver test between Montpellier and Nancy – 800 Km) and by IDRIS (OpenGL Vizserver tests between Paris and Poitiers – 400Km), were both successful. They demonstrated the effectiveness of this type of remote visualization which was more powerful than visualization on a local PC due to the power of the Onyx2 hardware. In both cases though, our common network provider RENATER dedicated a private virtual circuit to the operation (tunneling for bandwidth reservation).

In our case, we wanted to see the performance of OpenGL Vizserver with our regular network connection, and without any special installation or particular procedure. Therefore, our network connection and loads during the project were typical of our daily environment.

CERFACS and CINES are connected through regional network REMIP and national network RENATER with a 4 Mb/s peak bandwidth.

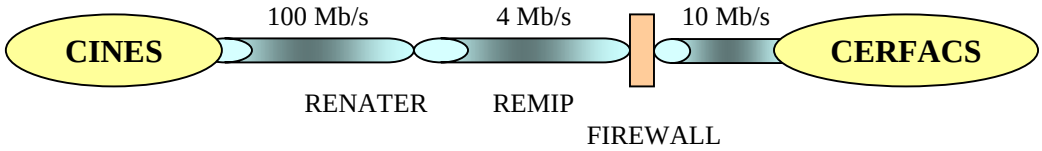

3. Experiments and Results

In order to estimate the time saving brought by remote visualization for results analysis, we compared the two following procedures:

1. Without remote visualization:

    - Downloading the data: each file size was about 8 MBytes for this test, but as stated above each file will be about 80 MBytes in a real environment.
    - Processing the data on our local visualization system.
    - Visualization of the results with EnSight® Gold by running a script that creates a movie from a set of JPEG images.

2. With remote visualization using OpenGL Vizserver:

    - Login to the Onyx2 system at CINES via OpenGL Vizserver.
    - Processing the data on the Onyx2 at CINES.
    - Visualization of the results with EnSight Gold through OpenGL Vizserver using the same script as was used in the case without remote visualization.
    - Download these JPEG images used to create the film.

Ease of use evaluation:

CERFACS’s connection to the RENATER network (4 Mb/s) enables having a good refresh rate in the EnSight Gold windows as long as we use one the available compressors of OpenGL Vizserver. The best compromise between images fluidity and quality is obtained through the "Interpolated Color Cell

Delta (4:1)" compression type. Zoom and image translations (the most refresh rate consuming operations) are usable. The ability to change compression type on the fly with OpenGL Vizserver is particularly useful in adapting the compression to the type of visualization used.

The impact of network saturation did exist, but fortunately their duration were relatively regular and limited.

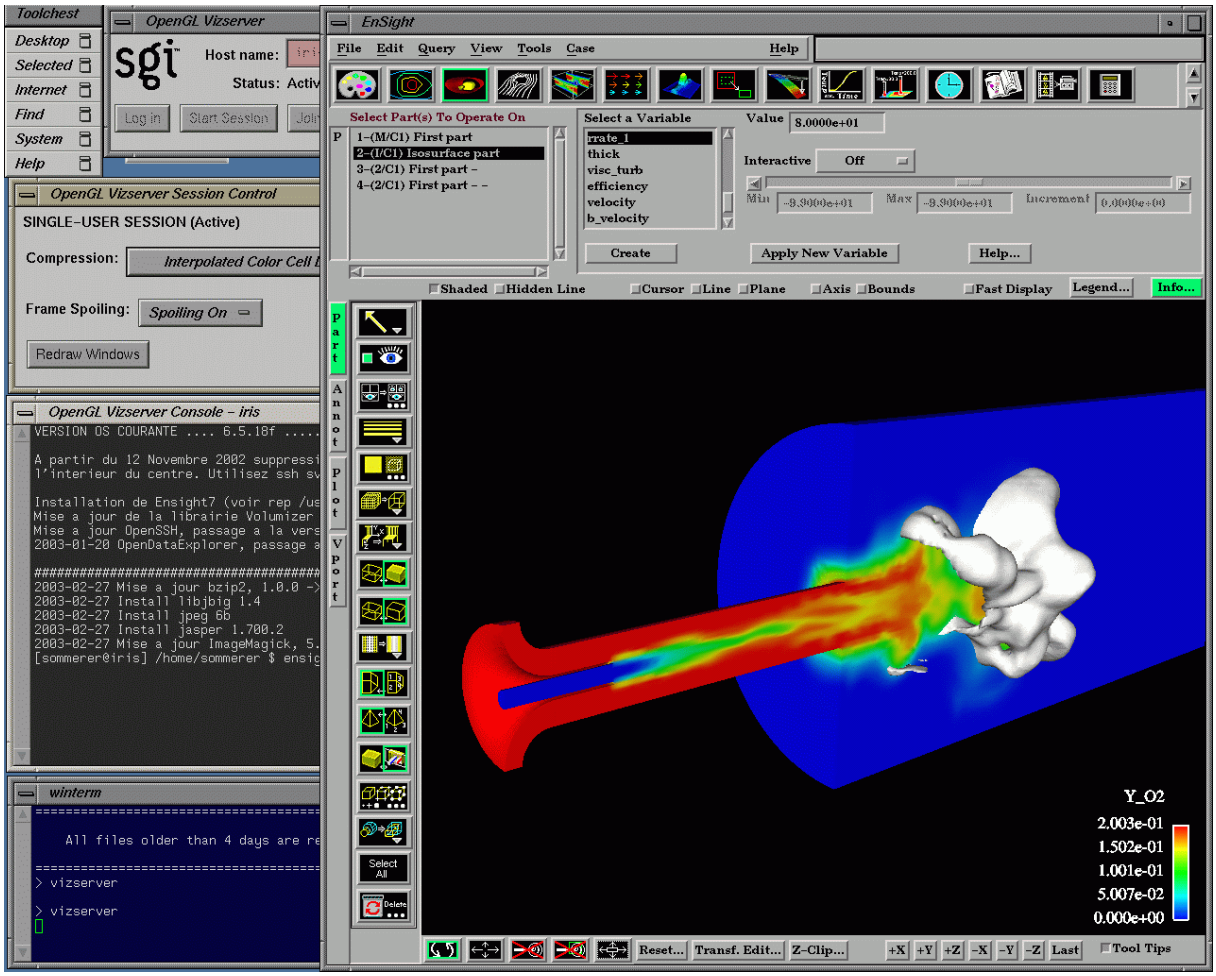

CINES post-processing and remote visualization from CERFACS during a combustion simulation

Comparative evaluation of post-processing times:

The following table allows comparing the execution times of the procedures described above:

	With remote visualization	Without remote visualization
File transfer time	11 s	1015 s
Data processing time	42 s	33 s
EnSight Gold session time	446 s	441 s
Total visualization time	499 s	1489 s

Post-processing via OpenGL Vizserver divides by 3 the time devoted to the total operation. Moreover, we must mention that this test case is small sized (10 MB per solution file, 7 solutions loaded). For target data sizes (usually higher than 80 MB per solution file, 10 solutions loaded) the effectiveness of remote visualization will increase:

- The processing capability of Silicon Graphics Onyx2 is markedly greater than that available to researchers at CERFACS. Visualization of large volumes of data on the CINES systems is much more interactive through OpenGL Vizserver than on our local machines.

- Only some small Postscript files were downloaded (small images making a movie) compared to data sets on the order of 0.8 GB or more! For example, only a few tens of seconds are enough to receive the movie files generated at CINES via OpenGL Vizserver whereas more than 4.5 hours would have been necessary to transfer the 10 intermediate solution files of 80MB each that would allow us to carry out this animation locally. The post-processing time using the Onyx2 at CINES and OpenGL Vizserver will be divided by a factor much greater than 10 in the real environment.

4. Conclusion

Very positive, these tests argue for the use of the Onyx2 together with the computing power of CINES. Whereas the transfer times between CINES and CERFACS made post-processing and control almost impossible, remote visualization provided by OpenGL Vizserver divided by more than 10 the time necessary for post-processing the data.

Usable between CINES and CERFACS on a WAN with a reasonable peak bandwidth (4Mb/s in our case) in real operating conditions (no dedicated VLAN), OpenGL Vizserver allows the control the results, their analysis and the building of remote movies. OpenGL Vizserver provided an acceptable loss of precision and delivered visualization that is more interactive through RENATER than Silicon Graphics® Octane® workstations on our local area network.

This service provides an entire, geographically spread community with a high-capacity post-processing server, centralizes visualization software and consolidates of our laboratory's needs. The proximity between the computing systems and the post-processing facilities allows efficient monitoring during long running, complex numerical simulations. Having this system available is an essential part of the growth of CINES's computing capabilities.

We insist on thanking Mr. Philippe Falandry and Mr. Jean-Chistophe Penalva for their support throughout this study, as well as the Simulog Company for the providing an EnSight Gold license during this test period.

What is CINES ?

CINES is the French National Computer Center for Higher Education (the Centre Informatique National de l'Enseignement Supérieur). Sponsored by the French Ministry of Research and New Technologies, CINES is one of the three main computer centers in France. It is located in Monpellier, near the Mediterranean Sea.

The center has 3 missions :

- To offer leading edge platforms and services for large scale computing within the French academic and scientific research community. User from universities and public entities have access to a 768 CPU SGI® Origin® 3800 server (0.7TFLOPS) and an IBM® cluster (1 TFLOPS) and through a 1 Gbit/second network operated by RENATER, the French Research Network. User support, training and help to analyze and visualize output from simulation applications are provided to enable the community to efficently use these systems.
- To provide infrastructure for large data bases and web services. A very large spectrum of applications are hosted, ranging from a reference catalog for documents found in university libraries (http://sudoc.abes.fr/), to medieval manuscripts (http://liberfloridus.cines.fr/); and university broadcasting channels (http://www.canalu.fr).
- To offer network expertise and training. Organized in common with RENATER, courses are regularaly offered on Ipv6, BGP, and metrology.

For more information , see http://www.cines.fr/ or send e-mail to doc@cines.fr.

What is CERFACS ?

CERFACS is the European Centre for Research and Advanced Training in Scientific Computation. With 80 researchers, it is one of the world's leading research institutes working on efficient algorithms for solving large scale scientific problems in various application fields : Climate Modeling and Global Change, Computational Fluid Dynamics, Electromagnetism and Control. This involves the evaluation

of existing tools and the development of new ones that exploit high-performance serial, vector, parallel and distributed computing resources.

What is SGI ?

Founded in 1982 and based in Mountain View, California, SGI (also known as Silicon Graphics), is a leading developer of high performance computing, data management, and visualization hardware and software solutions. These solutions help accelerate workflows and remove data management bottlenecks for government, manufacturing, energy, entertainment and sciences industries.

What is OpenGL Vizserver?

OpenGL Vizserver is a key part of the SGI® Visual Area Networking (VAN) strategy to accelerate large data and visualization work flows. OpenGL Vizserver provides an application transparent environment that enables remote users and distributed teams to access SGI visualization systems over the network. Users are able to take advantage of the full visualization, computing and high bandwidth data management capabilities without transferring data to their local system or travelling to the site of their data.

OpenGL Vizserver takes a stream of images that are generated in real time and distributes them in a compressed form to remote and mobile cliens running their choice of Microsoft® Windows®, Linux®, Solaris™ and SGI® IRIX® operating system software, and returns application control from those clients to the remote visualization server. Because the amount of compressed image data shipped over a network is independent of the size of the underlying application data , remote users are able to achieve a significant acceleration of their visualization work flow.

OpenGL Vizserver also enables multiple users in different locations to see the same results and share control of the remote application. When combined with it's application transparent capabilities, OpenGL Vizserver is able to take existing stand-alone applications and immediately turn them into collaborative ones – without waiting for software vendors to add specific collaborative funcationality.

3605 [11.14.2003] J14448