

HPC for Remote Visualization and Interaction with Scientific Applications

D. Riaño Nuñez¹, C.J. Barrios Hernández¹, J. T. Hernández Peñaloza²

1. High Performance and Scientific Computing Center (SC3-UIS) www.sc3.uis.edu.co
 2. IMAGINE (UniAndes) imagine.uniandes.edu.co
 Universidad Industrial de Santander www.uis.edu.co
 Bucaramanga, Colombia



Universidad Industrial de Santander



Abstract

Most of the scientific applications required visualization tools to evaluate their results. Also the scientific data are increasing exponentially, therefore the computational resources and visualization hardware must be more complex and expensive. NVIDIA GRID(TM) [1] is one of the latest remote graphics technology solution, designed for cloud gaming, allows hardware virtualization and improves the delay problems with remote applications. This technology can be used to develop scientific applications, integrating high performance computing and taking advantage of remote graphics. In this poster we present a work in progress, integrating remote visualization, interaction and high performance computing for the development of scientific applications, and also taking advantage of ultra high resolution display walls as a collaborative environment.

Background and motivation

Graphics as a Service (GaaS) is a new trend in computer science. This approach removes the need to purchase expensive GPU for rendering tasks. Working as a service, end-users can access to remote desktop environments with high-end configurations for a low cost. However, bandwidth is one of the common issues in cloud gaming, affecting the gaming experience with delay. GaaS requires high-end PC, this requirement is also for scientific applications, making this architecture ideal for researchers with high performance visualization needs. Therefore this visualization is the result of a long research process that generated a very large amount of data; the sheer size of data generated made high performance computing a necessity to obtain results.

When developing HPC applications visualization is often considered an independent step in the process, hence creating an opportunity to take advantage of this GaaS technology including HPC architectures like NVIDIA Tesla(TM) GPU.

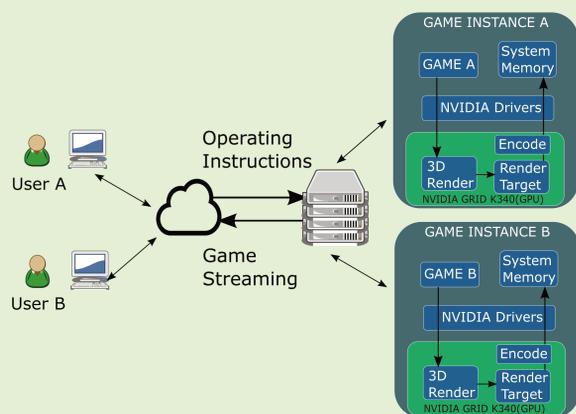


Fig 1. Architecture diagram of cloud gaming framework [2]

NVIDIA(r) GRID

NVIDIA GRID(TM) is a technology for desktop virtualization and video game streaming that offers remote graphics processing using GPU virtualization, remote processing and session management libraries. Furthermore it allows access to GPU for multiple users running their applications simultaneously with high speed. NVIDIA GRID(TM) SDK provides two component software APIs [2] :

- NvFBC: captures the entire desktop, including windows applications.
- NvIFR: captures from a specific render target.

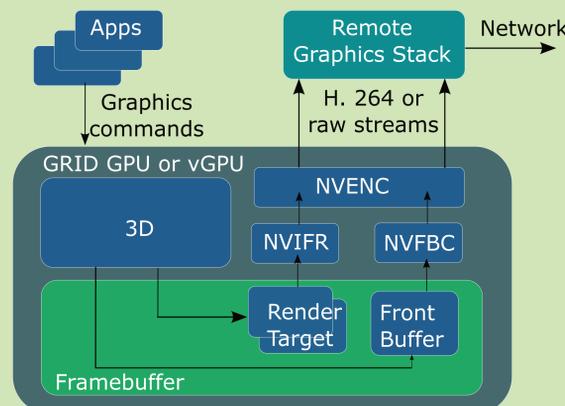


Fig 2. NVIDIA GRID(TM) Remote Graphics SDK [3]

Related Work

GPGPU computing is part of high performance and scientific computing, but GPGPU platforms are not available in all research sites. We proposed visualization as a service for scientific computing using GUANE, a hybrid supercomputing platform based on NVIDIA TESLA(TM) GPU. [4]

Software as Paraview [5] which provides a Scientific visualization tool, was developed to analyse extremely large datasets. It runs in various modes of operation: Client-Server Mode, Distributed Server Mode and Tiled-Display Mode.

SAGE2 [6] is a graphics middleware that enables real-time streaming of ultra-high resolution images and supports collaboration.

HPC Architecture SC3-UIS

UIS High performance and scientific computing center (SC3-UIS) provides the fastest supercomputer in Colombia (GUANE) to researchers in many fields of science.

Our infrastructure has 16 computing nodes with: 2 Xeon processors, 8 NVIDIA Tesla(TM) Fermi GPUs, 104 GB of RAM connecting with an infiniband network.

The NVIDIA GRID(TM) Server is a HP ML350R with: 6 Core Xeon processor, 16 GB of RAM and NVIDIA GRID(TM) K2. The NVIDIA GRID(TM) K2 has 2 NVIDIA Kepler(TM) GPUs with 8GB of DDR5 memory, 1536 cores per GPU and allows multiple user connections.

Further Work

Owing to the quantity of information that can be processed and visualized, the experts and developers must acquire high-end hardware and perform complex setup processes. Because of those processes, often cumbersome and time-consuming, the research itself gets less effective work time from the experts and developers. Besides, once the setup of the development environment is finished, bandwidth limitations are found.

There are similarities between hardware requirements for cloud gaming platforms and remote scientific application visualization. In this work we will propose a workflow for the development of applications that require HPC and visualization, using NVIDIA GRID(TM) to decrease delay that can affect the user-experience while interacting with a remote scientific application. In addition, the proposed workflow will aim to eliminate the necessity of acquiring high-cost hardware for the visualization of scientific application output.

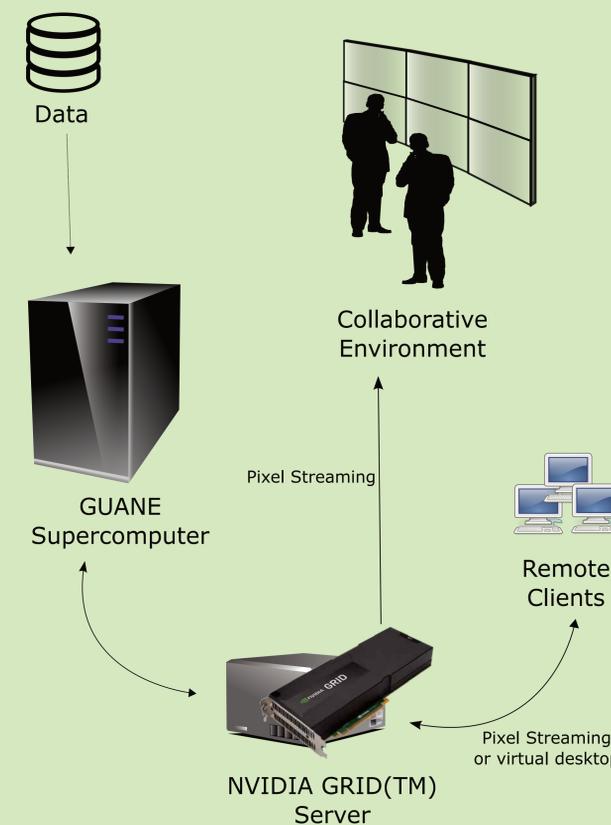


Fig 3. Proposed connection model of HPC-Visualization

References

1. NVIDIA GRID, Website: <http://www.nvidia.com/object/nvidia-grid.html>
2. Hou, Q., Qiu, C., Mu, K., Qi, Q., & Lu, Y. (2014). A Cloud Gaming System Based on NVIDIA GRID GPU. 2014 13th International Symposium on Distributed Computing and Applications to Business, Engineering and Science, 73–77. <http://doi.org/10.1109/DCABES.2014.19>
3. Delivered High-Performance Remote Graphics with NVIDIA GRID Virtual GPU. Website: <http://on-demand.gputechconf.com/siggraph/2014/presentation/SG4119-Delivering-High-Performance-Remote-Graphics-NVIDIA-GRID-Virtual-GPU.pdf>
4. Visualization as a Service, Website: http://on-demand.gputechconf.com/gtc/2013/poster/pdf/P0253_Gelvez%20Cortes.pdf
5. Nam, S., Jeong, B., Renambot, L., Johnson, A., Gaither, K., & Leigh, J. (2009). Remote visualization of large scale data for ultra-high resolution display environments. Proceedings of the 2009 Workshop on Ultrascale Visualization - UltraVis '09, 42–44. <http://doi.org/10.1145/1838544.1838550>
6. Marrinan, T., Aurisano, J., Nishimoto, A., Bharadwaj, K., Mateevitsi, V., Renambot, L., ... Leigh, J. (2014). SAGE2: A New Approach for Data Intensive Collaboration Using Scalable Resolution Shared Displays. Proceedings of the 10th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing, (CollaborateCom), 177–186. <http://doi.org/10.4108/icst.collaboratecom.2014.25733>



Super Computación y Cálculo Científico UIS



GPU RESEARCH CENTER



GPU EDUCATION CENTER