PathForward Project

Procurement & Materiel
University of California /DOE
Contract W-7405-ENG-48

STATEMENT OF WORK

The ASCI Challenge

The challenge undertaken by the Accelerated Strategic Computing Initiative (ASCI) to ensure the integrity of the U.S. nuclear weapons stockpile, without actual testing, is a formidable one. In order to meet this challenge, a computational modeling, simulation and analysis environment must be provided which enables unprecedented confidence in simulation results. The goal is to develop the computational equivalent of an actual experiment that simulates operation of and effects on a complete weapon system, including all subsystem components, at extremely high fidelity. The task of the Visual Interactive Environment for Weapons Simulation (VIEWS) program is to provide "see and understand" capabilities for the enormous datasets generated by these digital experiments. The results are fundamental requirements for the VIEWS program that are unprecedented in problem/analysis complexity and magnitude.

Since ASCI capability-needs exceed both current and projected technology levels associated with natural evolution, ASCI is consciously making efforts to accelerate technology development and commercialization where possible and appropriate. One way in which ASCI is encouraging such acceleration is by providing funding to key industrial research and development activities that show promise for delivering commercially viable technologies that would be useful to the needs of ASCI and VIEWS.

Increasing Visualization Needs

A critical need identified within the VIEWS program is to develop scalable visualization technologies that are matched to the massively scalable computation systems being constructed by ASCI. As dataset sizes continue to grow, this problem will become increasingly acute. Current computational platforms deliver 10TF of computational horsepower with 30TF and 100TF platforms planned within the timeframe of this subcontract. These computational platforms consist of thousands of nodes ranging from 2 to 128 processors per node. Each platform has its own system software. ASCI's scalable rendering systems need to be tightly coupled to the compute platforms which are producing the data sets.

While parallelism is being used at many levels in today's highest performance graphics systems, when scaled to their design limits these systems fail to deliver the performance necessary to provide desired levels of interactivity, even for current ASCI data sets.

ASCI visualization requirements have been identified that project a need for systems capable of delivering one thousand times the performance of the highest performing, hardware-accelerated graphics pipeline on a

single dataset in the year 2004. In particular the ASCI VIEWS program has identified two major milestones that it hopes to achieve in the next few years.

- 2002 – Demonstrate an interactive 64 Mpixel display driven by a scalable rendering system.
- 2004 – Demonstrate the ability to do real-time analysis on a 200TB ASCI dataset.

Both of these milestones depend critically on a scalable rendering and visualization infrastructure.

Desired Performance Features

1. Rendering extracts of large, time-varying datasets at interactive speeds

The table below provides some projected raw rendering and display requirements over the next few years.

Technology Timeline	1999 (Current)	2001	2004
Simulation Output			
# unstructured zones	100M	400M	1G
# vis variables (8 bytes)/zone	10	20	20
# vis dumps	200	500	1000
Structured Volumes			
Texture Memory	1GB	8GB	32GB
Fill Rate/frame	4 Gpixel	18 Gpixel	200 Gpixel
Unstructured Volumes			
Generated triangles/frame	400M	1.6G	4.0G
Surface Rendering			
Rendered triangles/frame	20M	100M	640M
Depth complexity	4	6	10
Display Resolution	2000 x 2000	4000 x 4000	8000 x 8000

This table is given in terms of data or pixels per frame, where a frame can correspond to one timestep of a large-scale simulation. The numbers in the table are not hard, in that the numbers are intended to represent the relative order of magnitude of the rendering problem. To be used interactively, the technology must be capable of generating images at frame rates between 5 and 60 Hz, depending on dataset sizes and display modality. It is important that frame rates not drop below 5 Hz even on the very largest datasets. Moreover, achieving interactive speed should not induce severe latency penalties. Image generation latency should not exceed 2-3 frames in any case.

As noted, datasets may contain hundreds or thousands of time steps. While most datasets typically contain spatial and temporal coherency, there are many datasets within which little (particularly temporal) coherency can be exploited. In arriving at the frame rate requirements, it is assumed that each frame can be from a different timestep. Additionally, rendered surfaces are often dynamically

generated, so technologies that rely on large geometry and texture caches must provide enough bandwidth and computational capacity to handle the case where the visualization primitives are recomputed every frame.

2. Support for both multiple offices and shared theater facilities

Scientists spend most of their time working in their offices alone or with a few colleagues. At each of the DOE weapons labs, support is needed for high-end graphics in at least 100 offices, with perhaps 20-30 concurrent users. Supporting a large number of offices from a single shared rendering facility is important. At first blush, it would seem that traditional graphics workstations might satisfy this need, but the size of ASCI datasets prohibits this. In virtually all cases, the datasets involved will be multi-terabyte aggregations of thousands of files that must reside on large RAID-based servers.

Collaborative theatre facilities using large multi-screen tiled displays and multiple screens are also becoming common within ASCI VIEWS. Scientists find utility in collaborative interactions supported by such facilities and the increased resolution available in such facilities is also of real utility. Current theatres use up to 15 tiles and provide display resolutions up to 6000x3000 pixels, which will continue to increase, as noted.

The proposed technology must be capable of supporting both desktops and theatre facilities. The ideal technology would be one capable of channeling all its rendering power to either kind of environment with little or no switching overhead. In the first case, it would need to support dozens of concurrent users each exploring a different, possibly multi-terabyte dataset on a desktop or wall display with a few megapixels. Dataset sizes are not necessarily reduced when lower pixel count displays are being utilized. In the second case, the technology must support a seamless high-performance display on a large tiled wall display. Thus the technology must be capable of dynamic resource allocation and resource aggregation to match dataset and display size requirements.

3. Multiple, distributed, remote displays over digital networks

The systems will support users located throughout a large campus. Thus it needs to deliver interactive imagery to users geographically distributed throughout this campus. These users will have heterogeneous display devices as noted previously, ranging from standard desktop monitors with stereo support and/or flat panels to tiled multi-screen displays. The system must couple to appropriate digital networks capable of supporting these various types of displays. Technologies that deliver to displays over standard computer networks are preferred to custom solutions. The laboratories' LANs consist of ATM OC-3 and 100 Base-T switched networks today with switched gigabit networks beginning to be deployed. There is also interest in the remote displays being supported across the WAN with multiple gigabit links available next year.

4. Support for rich graphical functions and primitives

The system must support a rich palette of graphics functions and primitives including volume rendering (both structured and unstructured), and rendering of opaque and translucent polygonal surfaces. It must be possible to combine these primitives in a single generated image.

Additionally, preliminary research into the use of programmable pixel and vertex pipelines within the VIEWS program suggests that the use of such features can potentially reduce the magnitude of the visualization data flow problem. As such, support for programmable pixel combiners, multi-texturing and programmable vertex submission environments are very desirable features.

5. Supporting software infrastructure

- Application transparent support for complex displays (e.g., remote network displays, tiling, etc)
- Legacy applications support (e.g., support for single context/pipeline OpenGL based applications)

The labs have a large body of existing software, much of which will be supported on these systems. The majority of this software is written under Unix utilizing OpenGL. While it is understood that applications will need to be modified to effectively utilize these systems, it is desirable if the systems provide some level of legacy support and that they provide largely transparent mechanisms to simplify the necessary porting effort.

6. Scalable data and visualization services integrated into rendering platform

- Parallel data paths for bringing data into the rendering system
- Ability of the platform to support data and visualization services in addition to rendering
- Seamless data integration with the largest ASCI computational engines. An example of this is a shared file system or a data server that feeds data to a rendering system.

The rendering systems will need to support an end-to-end visualization architecture. This includes the complete data path from disk or remote computational engine to screen. A proposed technology that might be capable of offloading some level of the visualization computation in addition to supporting scalable rendering is very desirable, especially in many of the remote visualization scenarios. The data being rendered will be generated and, perhaps, pre-processed on a separate platform. The capacity for, and the ease with which the proposed engine may be integrated with, this platform must also be taken into account.

7. Long-term architectural viability

- The use of commercial off-the-shelf components to leverage parallel technology development efforts
- Cost effective solution, support and timely upgrade paths
- Independent component upgrades, allowing the system to be customized, scaled and balanced to specific problems as well as allowing for individual components to be upgraded along different technology paths

These systems will form the backbone of the Tri-Labs capability visualization services for some time to come. Solutions that offer cost effective scalability and upgradability are required to achieve these goals. To that end, solutions capable of leveraging asynchronous emerging market technology

trends and development are desired as they potentially allow a system to avoid premature obsolescence and extend the lifespan of substantial infrastructure investments.

8. Packaging/management issues

- Reasonable power/space requirements (e.g., rack mounted vs. standalone components)
- Complete system administration package
- Systems stability - features, such as redundant power supplies, that increase uptime
- Security - features that enable proper encryption/protection of classified information
- Integration into existing laboratory computational and display infrastructures

The proposed system must be realizable and maintainable within the peculiarities of the DOE environments. The ideal system would provide a complete administration package, scaled to the desired system size. This would include the capacity for dynamic alternative routing to help address component failure issues common in extreme systems. Finally, the system must be capable of supporting a secure computing environment as is present in the labs.