Report Number: WUCS-91-56

1991-12-01

Remote Visualization: Challenges and Opportunities Panel Statement

Gurudatta M. Parulkar

Remote Visualization: Challenges and Opportunities
Panel Statement

Guru M. Parulkar

WUCS-91-56

December 1991

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

Appeared in the Proceedings of IEEE Visualization '91, October 1991

Remote Visualization: Challenges and Opportunities

Organizer: Guru M. Parulkar, Washington University in St. Louis

Panelists: Jack Bowie, Digital Equipment Corporation
Hans-Werner Braun, San Diego Supercomputer Center
Roch Guerin, IBM T.J. Watson Research Center
Guru M. Parulkar, Washington University in St. Louis
Daniel Stevenson, MCNC

Motivation and Scope

Scientific visualization has emerged as a major computer-based field of study. It is becoming increasingly evident that the visualization is indeed a critical tool for discovery and understanding as well as a tool for communication and teaching. The use of visualization for many scientific applications has been well documented. New visualization applications have been rapidly emerging with developments in visualization methodology and underlying technologies. In short, visualization is beginning to revolutionize the way research is done in various disciplines of science and engineering.

Remote visualization is visualization that utilizes data and computing resources that are physically distributed. There are three components in the visualization process, namely data, computation, and user interaction. As long as the locations of these components are not all the same, the need for remote visualization arises. Furthermore, the scientist may wish to do parts of the visualization computation on separate machines in order to distribute the computation load and achieve better performance. Therefore, we believe a significant fraction of practical visualization will be remote visualization.

It is also important to note that the proposed national program for the High-Performance Computing and Communications (HPCC) and National Research and Education Network (NREN) essentially requires as one of its important components development of the infrastructure to support remote interactive visualization.

Efficient remote visualization requires support far beyond what is needed for visualization on a single computer. It requires networks with high bandwidth and low latency, an efficient interstage communication mechanism on the networks, and proper adaptation and partitioning of the visualization computation. Trends in high speed networking suggest that the next generation networks will provide enough bandwidth to allow remote visualization. However, it is important to note that host-to-host connections with high bandwidth and performance guarantees are necessary but not sufficient conditions for successful deployment of demanding distributed applications. As a result of decreased queueing delay in newer networks, the speed of light propagation delay will increasingly dominate the communication latency in wide area networks. The networks will continue to have more packet errors and losses than a local environment. Moreover, the new networks have much higher bandwidth-delay products which affect flow and error control strategies at the application and transport levels.

Recent developments in visualization and collaboration tehcnologies suggest that computer based collaboration applications will need to include support for multimedia and visualizations. This leads to a new set of interesting issues that have to do with synchronization and concurrency control among various data streams and among various end points of the conversation.

It is the responsibility of applications and the transport protocols to cope with these conditions and to convert the high bandwidth into high performance for the applications. It has been recognized that suitable solutions can be found by developing deep understanding of the communication requirements of various classes of applications. The interactive remote visualization presents a unique set of interprocess communication (IPC) requirements that have not been adequately addressed by existing research on networking or visualization.

The purpose of this panel is to emphasize the need for and importance of remote visualization; to predict the impact of remote visualization on application algorithms, communication protocols and underlying networks; and outline opportunities for research and development to support remote visualization in the context of the NREN.

End-user Requirements

Jack Bowie

I will represent the end-user point of view on remote visualization. What are the key application issues in remote visualization? How do users access these capabilities? What kind of visualization environments can support remote visualization? What is the role of industry standards? How will the concepts of remote visualization be adopted and implemented by the vendor community?

Network Infrastructure

Hans-Werner Braun

Over the past several years, the US research and education data network infrastructure has experienced dramatic changes in scope, both in terms of available bandwidth as well as a move towards more ubiquitous access. About five years ago the National Science Foundation started to build up infrastructure available to the scientific research community as an outgrowth of technology previously developed by DARPA researchers. The result was a much more widespread use of data communications, and therefore, a need for even more and higher-performance infrastructure. The national data communication fabric has since evolved from a 56kbps network to a 1.5Mbps network in 1988 and to the beginnings of a 45Mbps network by the end of 1990. Efforts are underway to research the issues surrounding networking at gigabit speed and eventually to upgrade the generally available infrastructure to such levels. The latter is part of multiagency efforts towards a National Research and Education Network (NREN) that will be part of the larger international Internet.

Researchers today can access the national network by means of regional networks that campuses connect to. Those regional networks then connect to national interconnection networks like the backbone network of the National Science Foundation, NSFNET, networks of other agencies, or the networks of commercial service providers.

The national networking infrastructure, including the concept of the National Research and Education Network, is instrumental to research throughout the country. It enables researchers to exchange large amounts of data with remote computers and, in many cases, to create images of data sets, which then can be sent across the network. It also allows for communication between colleagues, both domestically and abroad via attached international connections.

Today's high-performance networks are used for both aggregations of many data flows as well as an enabling factor for applications that demand high bandwidth by themselves. One such high-demand application is remote visualization. While just a few years ago remote visualization was not reasonably possible, given the available bandwidth, today's bandwidths allow people to design and implement advanced applications that require remote interactions with serious traffic flows between the two or more end points of the resulting conversations. As we move towards gigabit per second speeds, such applications will become more and more interactive and oriented towards real time, possibly continuous data flows, which would enable the development of even more advanced applications. Those data flows could be used for remote visualization. Another example is high-speed intercomputer communication such as the distributed applications that are being researched between supercomputers in the context of gigabit testbeds.

At the same time, a plain availability of high bandwidth for the aggregation of demands of many users is insufficient to accommodate the requirements of more demanding applications, like remote visualization. As the network evolves, a need for service guarantees also evolve, as the network has to strive towards predictability. Unless service guarantees, like specific amounts of bandwidth or network latencies, are at acceptable levels for end users, continued evolution of high-demand applications will be hampered.

Public vs. Private Network Solutions

Daniel Stevenson

In the state of North Carolina we have an excellent infrastructure for data communications use by our academic research community. The MCNC owned and operated data network consists of a private 45 Mbps backbone interconnecting Ethernets and FDDI rings. The network (which has been in place for five years) connects 13 major sites and over 4000 workstations/hosts.

The recent creation of a supercomputing division at MCNC and the deployment of a Cray Y-MP has created a discontinuity in the usage of our data network. As the community of applications researchers has gained access to MCNC's supercomputing facilities over the past 18 months, their knowledge of supercomputing techniques has grown dramatically over a rather short time. Our visualization group maintains a busy schedule supporting the needs of various end users doing computational science on the Cray. One of the key thrusts we see for the visualization and communications groups is the development of a technology base to support remote visualization by our user community. But while our network capacity is sufficient to support one or two visualization processes simultaneously, we anticipate that the growing need for high-speed networks to support visual output from computational models will begin to overwhelm the existing data capabilities provided by our network.

As we chart the paths for our future development of network capability we expect to increasingly turn to use of public network solutions rather than continue to rely exclusively on private network solutions. We see this process as being consistent with the probable direction of the NREN. The telecommunications standards we expect to provide the performance and reach characteristics we foresee needing to support our user community include SONET and ATM. Other technologies that are being pursued by various organizations suffer from a variety of problems, insufficient capacity to support multiple visualization users on a single network, limited reach, availability only in private networks.

We are finding that our early interactive visualization applications require more computational power than can be delivered by the Cray. To achieve interactive image rates the models are partitioned across multiple compute engines with each processor assigned a task best suited for its particular architecture. In the VISTAnet project the Cray is coupled with the Pixel-Planes 5 (a high performance graphics processor) in addition to other hosts. While the resulting image stream requires bandwidth of less than 100 Mbps the data flow between Pixel-Planes 5 and the Cray is expected to reach several hundred megabits per second. We expect the model of functional partitioning between compute servers and visualization servers to continue as the demand for computational steering increases.

While ATM promises plenty of reach and reasonable bandwidth for visualization it is not without its technical shortcomings. These include difficulties associated with small cells in high performance links and the relative complexity and rigidity of the adaptation process. Furthermore the characteristics of ATM (BER, cell loss statistics) seem significantly different than the models that the current internet protocols are based on. As part of the VISTAnet project we are involved with North Carolina State University in efforts to understand (from a performance perspective) the interactions between ATM networks and transport layer protocols.

Interprocess Communication Support

Guru M. Parulkar

As part of our research on end-end communication in high speed networks, we have recently started an effort to study a set of distributed visualization applications. Our aim is to understand the real communication requirements of these applications and appropriately tailor the protocols to provide efficient support.

We came to understand that a distributed asynchronous pipeline is an appropriate model for most remote visualization applications. This is also evident from emerging visualization tools such as apE and AVS. In these distributed pipelines, interstage communication incurs significant delay due to the physical distribution of pipeline stages. The slowest stage becomes the bottleneck and limits pipeline speed. In order to achieve efficient pipelining, the interstage communication has to satisfy a number of conditions:

- A visualization application typically involves a large number of data segments of considerable size. These segments have to be streamed through the pipeline with minimum delay to allow overlapped processing.
- Computation and communication speed should be well balanced for optimal utilization of both resources.
- An application should be able to specify its error tolerance requirements to the IPC mechanism and the IPC mechanism should enforce it only to the degree that is necessary to satisfy the requirements, thus avoiding any unnecessary error control overhead.
- Buffer overflow in the pipeline should be avoided because loss of partially processed data due to overflow may require restarting the pipeline from

an earlier stage, thus wasting computing cycles and introducing unnecessary (and unacceptable) delay.

Existing IPC mechanisms and underlying communication protocols do not satisfy these requirements. We have proposed an IPC solution with appropriate transport mechanisms that to address the problems and provide efficient IPC support. There are three important aspects to our solution:

- We propose a novel interprocess communication paradigm called segment streaming to facilitate implementation of asynchronous pipelines across networks. A segment is a logical unit of data that is independently processed by application processes. By segment streaming, an application process only needs to make a single system call to perform the request for all segments; each segment will be transmitted when ready across a connection without the latency of request or setup. Therefore, it supports segment prefetching and allows overlapping between intrastage computations and interstage communications. This provides a necessary condition for efficient operation of the pipeline.

- Asynchronous pipelines across networks also require carefully engineered flow control to avoid frequent pipeline overflows and flushings. We propose a two stage flow control. We use connections with rate control between application and the underlying network. The rate is set at the time of connection set up and is rarely adjusted during the life of the connection. On top of the rate control, we propose to use a simple window control. The purpose of this window control is to do end-to-end flow control between pipeline stages. It is important to note that the window control is not responsible for any congestion control within networks or for error control as is the case with the existing transport protocols.

- Error control for visualization applications is challenging because their error control requirements lie on a spectrum between those of data and video applications. A typical data application requires error free data exchange which means every packet loss and corruption has to be recovered. A typical video application on the other hand does not require any error control (as long as the error rate is reasonably low) because the subsequent frames overwrite the previous frame. In the case of visualization and other imaging applications, the error tolerance varies depending on the application and pipeline stage. We propose an application-oriented error control scheme which does error control as dictated by the error tolerance of a particular pipeline stage.

Impact on Networks and Protocols

Roch Guerin

The PLAnet architecture is a new architecture for fast packet-switched networks, which is to be tested and deployed in a number of forthcoming testbeds, e.g., the Aurora testbed. The architecture extends across local, metropolitan, and wide area environments, and supports a variety high-speed connection services across the network. In particular, it allows for both datagram and connection oriented services, and provides sophisticated bandwidth management functions. This allows the network to not only ensure guaranteed resources to connections, but also to allocate bandwidth as a function of connection needs and requirements. A fast setup and disconnect mechanism further enhances the flexibility of the network.

Such a flexibility is key to efficiently supporting both existing and emerging high-speed applications with heterogeneous and often highly variable bandwidth requirements. This is particularly true for connections between high performance graphic workstations and supercomputers, which can have very high throughput requirements but not necessarily sustained over long periods of time. In order to provide both an adequate Grade-Of-Service (GOS) to these connections and efficiently use the network resources, it is critical to properly characterize these requirements and map them onto the appropriate network connection types.

I will briefly review the features of the PLAnet architecture and the associated transport protocol, and then turn to the issue of supporting high-speed applications such as remote visualization. This is done based on some experience with a high-speed graphic engine developed at IBM. The feature of the traffic generated by this graphic engine will be reviewed and their implications in terms of bandwidth requirements and connection characteristics will be discussed. For example, the relation between the peak rate, long term rate, and burstiness of a connection, and the amount of network bandwidth it requires will be detailed. The focus will be on the impact of the traffic generated by such high speed workstations on the network, and on the best ways of supporting the associated connections. In particular, the various possible mapping onto the bandwidth allocation procedures of PLAnet will be discussed, and the trade-offs associated with the different alternatives available (i.e., burst vs connection reservation, lower GOS vs lower latency, etc.) will be identified. Although the PLAnet architecture is used as a base for these discussions, most results should be applicable to other fast packet-switched networks, e.g., ATM.