NSF/ARPA Science and Technology Center for Computer Graphics and Scientific Visualization

GV-STC Digital Media Network
October 1993

Abstract

This document describes the design and functionality of the data component of the low-latency, dedicated, digital video and data network the GV-STC has developed as part of its digital media infrastructure. The network supports simultaneous multi-point transmission of compressed video and digital data at aggregate rates of up to 1.5 Mb/s (T1). The network is designed to support Center-wide access to specialized instrumentation, computational resources, and visualization systems at individual Center sites. The network's combination of low-latency, guaranteed bandwidth data transmission and high-quality video feedback, provides a unique resource for studying a number of important research problems including: time-critical visualization; bandwidth/latency tradeoffs in remote interactive systems; image compression technologies for synthetic images; and collaborative research tools.

1 Introduction

The original proposal for the Science and Technology Center for Computer Graphics and Scientific Visualization (GV-STC) acknowledged the need for allocating substantial effort and resources to the goal of collaboration. Graduate student and faculty exchanges and airport meetings were originally envisioned as the primary method to achieve collaboration between the geographically separated sites of the Center.

However the first year of Center operation demonstrated that ongoing collaborative research requires more frequent contact than travel schedules permit. Electronic mail and phone connections do well enough to support routine correspondence, but face-to-face discussions and shared graphics environments are critical to group collaborations.

In September 1991, at the NSF's request, the GV-STC submitted its Criteria for Evaluating Center Effectiveness, outlining research goals and standards against which the Center proposed to be judged. That report lists more specific goals for collaboration, stating that, "In three years we expect to have developed an effective working methodology for sharing software, models and images in a controlled fashion amongst ourselves, our industrial sponsors and other collaborators ... The [GV-STC] will eventually serve as a model collaboratory through a high bandwidth, video teleconferencing communication network. We will attempt to install a high-speed communication system as soon as the technology will affordably allow. Our ten year goal is to have permanent real-time video windows into each laboratory, enabling continuous conferences from each of the five laboratories to each of the others, in a sense creating 'electronic neighbors'."

In May of 1992, the GV-STC proposed a five-way videoteleconferencing system to the National Science Foundation, and in September of 1992 received supplemental funding for hardware acquisition and carrier line charges. The document "STC Proposal for A Widely Distributed Videoteleconferencing Environment" (May 1992) describes the design and technical specifications of the system.

Full-time, dedicated, five-way audio and video communication has now been operational for nearly one year. We are rapidly accumulating technical expertise and pedagogical experience with the system. As one significant outcome, we have developed specialized multi-site software to support scheduling and control of conferencing and video production functions from workstations anywhere within the Center. Our videoconference control software is described in the document "GV-STC Videoconference Controller (VCC)" (October 1993). Our networked video production software is described in the document "GV-STC Video Widgets" (October 1993).

New datalink components coming on line in November 1993 will enable use of the GV-STC T1 network as a low-latency, guaranteed bandwidth data network in concert with its videoconferencing functions. This will provide real-time multimedia connections between our sites. The datalink components will improve remote access to specialized resources at individual Center sites and will open new avenues for research into time-critical computing, interaction-at-a-distance, data compression and transmission algorithms, and collaborative research tools.

The remainder of this document describes this new datalink portion of our network.

2 Project Overview

The GV-STC T1 network was activated in January 1993. For the past ten months we have been using the network primarily for videoconferencing; conducting frequent research and administrative meetings, colloquia, and since September 1993, teaching a twice-weekly, five-university, graduate level seminar on advanced topics in computer graphics and scientific visualization. The frequent face-to-face contact afforded by the video medium has allowed us to develop our collaborative research and technical agendas to the the point where we are ready to begin the second phase of development of our digital media infrastructure.

The goal of this second phase of development is to provide a low-latency, assured bandwidth digital channel connecting the five sites of the GV-STC. The digital channel will share the T1 links with the existing video channel, and will be coordinated with it, providing the Center with a unified audio/visual and data communications medium. This multimedia system will be used for several purposes. First, it will provide Center-wide access to site-specific instrumentation such as Caltech's MRI imaging laboratory, Cornell's light measurement laboratory, or Utah's NC machining facility. Next, it will provide a dedicated channel to specialized computational resources such as Cornell's FDDI-coupled workstation cluster, or UNC's PixelPlanes machine. Finally it will provide a testbed for fundamental research on issues such as time-critical/bandwidth-limited visualization, visualization/interaction at a distance, and synthetic image compression/transmission algorithms.

3 Design Issues

3.1 Network topology

Figure 1 shows the hardware configuration of the GV-STC T1 network. The network has a spoke and hub configuration with the hub located at UNC-CH. At each site the T1 line from the long distance carrier is broken out into 24 64Kb/s channels by a Channel Service Unit (CSU). The CSU has three ports which can be configured to connect to any contiguous bank of 64Kb channels. In this way fractional T1 service can be allocated to devices connected to the CSU.

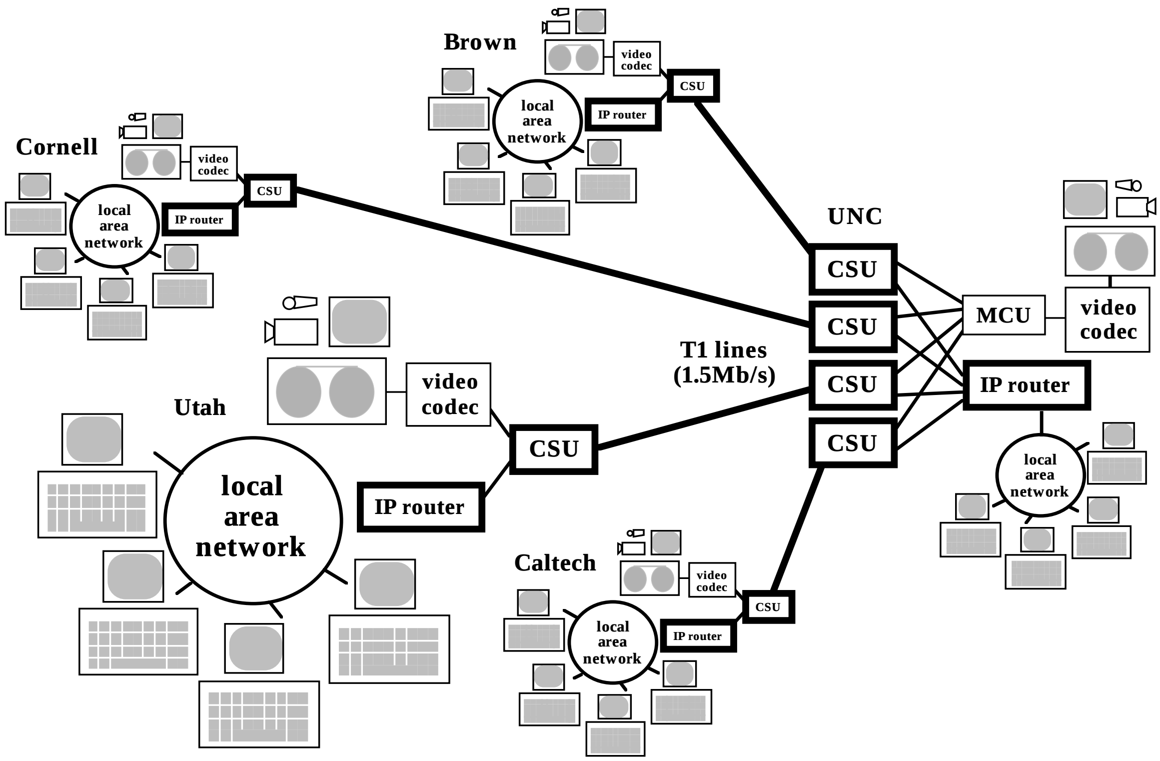

Figure 1: GV-STC T1 network

3.2 Bandwidth allocation

At each of our sites, one port on the CSU is connected to our videoconferencing codecs. The codecs can transmit compressed video and audio at rates from 64Kb/s up to 1.5Mb/s (full T1). Another port on the CSU is connected to a data packet router running the TCP/IP protocol. The router is also capable of passing data at rates from 64Kb/s up to full T1. When both video and data are being transmitted on the T1 line the codec and router must share the available bandwidth.

A router is connected to each of our local area networks. Data packets originating at one site travel from the local area network through the router, get merged at the CSU with digital video and audio coming from the site's videoconferencing codec and are sent cross-country on the T1 lines to the network's hub at UNC-Chapel Hill.

3.3 Hub functions

At the hub are four CSU's which breakout the T1 lines originating at the other sites. The ports on each of these CSU's are connected to specialized equipment which handles switching functions. The videoconferencing codec channels are connected to a device called the Multipoint Control Unit (MCU). The MCU handles voice-activated switching during videoconferences as well as performing many other codec related functions. Another port on each of the hub CSU's is connected to a four-port data packet router. This serves as the switching point for all data traffic carried on the T1 network.

3.4 Traffic control

One issue that required serious design consideration was the choice of whether to use static or dynamic data packet routing. On the Internet, routers communicate with each other periodically to make their connectivity known. In this way the network "knows" which of the many network paths from computer A to computer B are available and which ones are offering the best throughput. This is known as dynamic routing and it serves the Internet community well by providing a robust and adaptive scheme for assuring uninterrupted traffic flow. Unfortunately, while packet delivery is virtually assured, there are no latency guarantees since the load on any particular link in the network is a continuously varying function of the number of other available links, and their aggregate bandwidth.

The GV-STC's decision to use static routing in its T1 network is based on its need for low-latency, assured bandwidth, data transmission between sites in support of its research goals. With static routing network-wide "knowledge" of the availability of the data path over the GV-STC's T1 network can be restricted to local area networks, sub-networks, and even individual workstations within the Center's laboratories providing a much greater level of control over traffic load and packet latency.

4 Support for Collaboration and Research

Our multimedia T1 network is designed to support collaboration on continuing research topics as well as provide a foundation for new research on networked digital media. In the area of research collaboration, researchers at Brown will now be able to use the interface tools they have developed for Utah's Alpha-1 solids modeler to design an object, and they will then be able to watch and interactively refine the the model as the object is manfactured at Utah's NC machining facility. Researchers at Cornell will be able to "walk" through complex radiosity simulations displayed in real time on UNC's PixelPlanes machine. MRI data from Caltech can be volume rendered at UNC or Cornell and viewed and manipulated from Caltech. The potential for significant interaction between researchers at our geographically dispersed sites has been increased enormously.

While the immediate use of the network will be to support collaborative research and teaching, the medium itself suggests many worthwhile research projects. As networks grow and remote, interactive access to instrumentation, computational resources, and visualization systems becomes commonplace, issues of network bandwidth and latency become central to making remote access workable. Strategies for time-critical visualization need to be developed where interactive responsiveness and image quality can be optimized within the contstraints placed on the communication by the network.

A related topic is the development of compression and transmission algorithms for graphical data and synthetic images. In general, visualization produces an information explosion as we go from a model to images. If we are doing real-world modeling such as in medical visualization, and we add data acquisition into the process, then we also have a vast amount of data being generated prior to visualization. Where to place the network in the stream from instrumentation to visualization, and how to effectively manage the vast amounts of information which must be processed are open questions which our multimedia network will allow us to begin to address.

5 Conclusions

The GV-STC originally requested funding for its T1 network because we felt that a geographically dispersed center would only succeed if a means was provided by which researchers at the individual sites could get into each other's laboratories, offices, and meeting rooms on a regular basis. The system we've developed provides five-way, full-time, simultaneous, bi-directional, audio/visual and data communication between our Center sites. This multimedia connection gives our researchers much greater access to each other's laboratories, each other's tools, and each other.

To develop this system we have had to work together to understand the technology and adapt it to our needs. Our research support environments are now more tightly integrated and at the same time we now have immediate access to the specialized and highly developed resources each center site has to offer. We now have a unique and vital tool which we can use to pursue our research goals.

6 Acknowledgements

This work was supported in part by the NSF and ARPA Science and Technology Center for Computer Graphics and Scientific Visualization (ASC-89-20219). All opinions, findings, conclusions or recommendations expressed in this document are those of the author and do not necessarily reflect the views of the sponsoring agencies.