15
Light Table: An Interface To Visual Information Systems

Patrick Purcell
Dan Applebaum

The Media Laboratory
Massachusetts Institute of Technology

Introduction

A primary aim of the Light Table project was to see if a combination of the optical laser disc, local area networks, and interactive videographic workstation technology could bring a major visual collection, (such as the Rotch Visual Collections of the Massachusetts Institute of Technology), to a campus-wide population of undergraduate users.

VIS (Visual Information System) is the name being given to the new genre of information technology. Much research and development effort is currently being applied to areas where the image has a special significance, for example in architecture and planning, in graphic and fine arts, in biology, in medicine, and in photography.

One particular advance in the technology of VIS has been the facility to access visual information across a distributed computer system via LAN (Local Area Networks) and video delivery systems, (such as campus TV cable). This advance allows users to retrieve images from both local and remote sources, dispatching the image search through the LAN, and receiving the images back at their workstation via dedicated channels on the campus TV cable.

Light Table is the title of a system that acts as a computer-based interactive videographic interface to a variety of visual information systems described in the body of this paper. It takes its name from the traditional, back-lit, translucent light table that lecturers use to assemble and view collections of slides for talks and seminars. The component of Light Table which is being reported in greatest detail here, a software outcome called Galatea, is a versatile and robust system capable of controlling video devices in a networked environment.

Context for Light Table: MIT Project Athena

The context for these developments has been Project Athena, a comprehensive educational program, launched at MIT in 1983, with the collaboration and support of DEC (Digital Equipment Corp) and IBM (International Business Machines), to develop a new technical infrastructure for under-graduate use (recently extended to graduates) across campus in the coming decade Norman 1989).

While Athena's primary mission is educational enhancement, its technical goals have been several. One goal is coherence, namely to achieve a uniform operating system at the Athena user interface. Unix (AT&T) was the choice as the operating environment (Leffler et.al. 1989). Another goal has been to develop a digital network which would accommodate the campus-wide computing activities of hundreds of workstations, together with file servers, mail servers and printers. Athena has also sponsored new courseware and teaching strategies to take advantage of this computing context of unprecedented scale and complexity. The standard Athena Workstation is a 32 bit, 1-3 MIPS CPU with 4-8 MB of RAM, a hard disc, a monochrome megapixel display, a mouse and an Ethernet interface. The present deployment of approximately 800 machines, is a mix of DEC VAXstations and IBM PC/RTs. Selected workstations are configured with color displays and have multimedia capability.

In the technical area, there have been many interesting spin-offs, with potential well beyond the original Athena agenda. Among the most prominent has been a window system called "X" (Scheifler and Gettys 1987), a powerful network-transparent windowing program and the environment for Light Table.

Athena's educational mission has yielded a number of distinctive and interesting results, including computer-aided design tools, and interfaces to laboratory experiments. Multimedia applications having agendas in education, entertainment and communications, promise to be among the most significant Athena outcomes (Mackay and Davenport 1989).

The Visual Information System

The focus of this paper is the electronic Light Table as a highly interactive videographic interface connected to different visual information systems. So far, the VIS source images have almost entirely come from the MIT Retch Visual Collections.

Rotch Visual Collection

The Retch Visual Collection is a library of approximately 300,000 images in the form of slides, photographs, drawings, and microfiche. Covering the areas of architecture, urbanism, art and photography, it serves as a research and educational resource to faculty and graduate students. The progressive transfer of visual images from this important source to electronic media will make a radical change in the role of the Collection at MIT by making it accessible to undergraduates using Athena workstations across the campus.

The image banks accessed by Light Table include Archfile, Picassofile and the Boston Project. Archfile is an architectural image library that is comprised of approximately 5,500 records of architectural design. Each record in Archfile is divided into nine fields in the following manner. Fields 1 and 2 define the location of the building or project. Field 3 assigns the building as one of eight architectural types, (including religious, residential, medical, public, educational and commercial). Picassofile (Purcell and Okun 1983) is the name given to a small, but detailed image bank and data base built around the works of Pablo Picasso. The Boston collection (Smith 1987) is the most recent of the databased image collections to have been compiled. It consists of approximately 7,000 examples of significant architectural design or planning projects in the metropolitan Boston area. Its parameters of search and the organization include fields such as architect, location, frame number, year of construction etc. For example, a search to show the work of architect I.M. Pei in Boston generates 32 references and associated images. To constrain the search further the user may wish to specify a year of construction, or specify a building type (for example residential), or specify a particular precinct in metropolitan Boston, or combine all of the above.

Light Table

Light Table, begun in 1978, presents multiple images on the screen, as independent entities, each with its own border. Goals of the project have included portability, distributed use, and efficient use of resources.

To enhance portability for Light Table, (especially for different resolutions and for both DEC and IBM platforms), the X Window System was chosen to control the graphics output. Since X-Windows was designed to provide a machine- and device-independent method of creating interactive graphics, it was on its way to being adopted as an industry standard. Its use made it possible to develop a single program that would run both on the variety of hardware utilized by Project Athena and later in outside environments as well. The first version of Light Table was created using the X Window System, version 10. This first version only used digital representation of stored images. These images could be stored in a one bit, eight bit, or twenty four bit format. Images from any of the three storage formats were converted for display on either one bit or eight bit display depths.

Figure 1 General view of Athena video workstation showing the electronic 'light table.

Figure 2 MIT Rotch Visual Collection, with conventional light table.

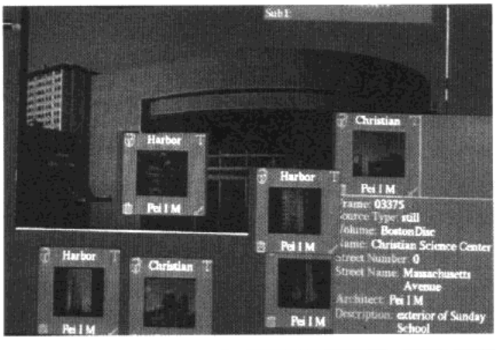

Figure 3 Images being retrieved on light table interface. Collages of mages at various scales are shown.

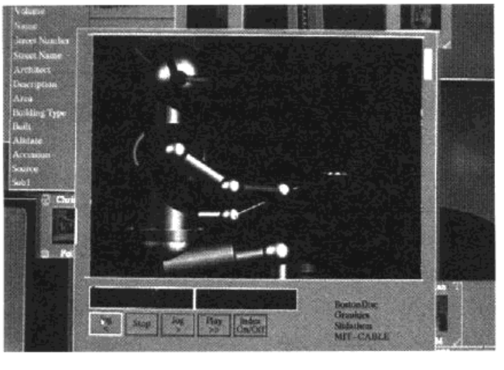

Figure 4 Light Table window running a digital videographic animation "beat dedication" (Sabiston-Visible Language Workshop)

The user interface for this system contained individual slides for each image plus a database control window. Each slide was comprised of several components. The primary component was a miniature version of the image. Around this was a wide border with several icons, a title, and the artist's name. The icons, located in the corners, could display the slide's image at full size, delete the slide, or protect the slide against accidental deletion. The database search window, (outside of the slides), controlled which images would be loaded into slides.

Although providing a good user interface for a slide library, the prohibitive cost of digital image storage required that an alternate storage method be found. The early version of Light Table was modified to take advantage of videodisc storage technology. In addition to using images stored on the file system, the program could search to a single image frame on a videodisc stored in a videodisc player connected to the computer with an R5232 serial line. Displaying an image found in this manner on a video monitor next to the main graphics display was not a very good substitute for the slide metaphor used earlier. A method of digitizing the analog videodisc images was required. The solution entailed using the Parallax Graphics model 1280 videographic subsystem to digitize the analog video directly into the frame buffer, so that the images were displayed at the same location and size as the digitally stored images.

To provide a distributed operating environment, the videodisc control method had to change from a local control system to a network-based control system. An experimental program was designed to listen to access requests from other programs and to control the videodisc players accordingly. Since this program program served the needs of other programs, it was called a "server.' This program remained running at all times, making it a "daemon" process. Light Table was redesigned so that it sent commands to this videodisc server, instead of sending commands directly to the videodisc players, hence, making it a "client" for the videodisc "server.' The commands issued to the videodisc server were independent of the brand of videodisc player being used; the server was responsible for adapting to the particular player being used. As new players were added, only the server required modification.

As the server could respond to requests from multiple clients, and manage switching between clients, users could share the videodisc players. Since the video digitizing process takes only about one thirtieth of a second, a client only needed a frame kept on the player for a short time. By allowing users to share the videodiscs and players, a more efficient use of relatively scarce resources was accomplished. Two videodisc players, controlled by a workstation and feeding an audio/video switch, were installed in the Rotch Visual Collection. The output of the switch was modulated to be sent out over the MIT campus cable television system. Users taking advantage of Light Table could access 108,000 frames of visual information from many locations on the campus. Once the experimental server was completed, work began on a more powerful and more flexible videodisc control system which would be available for other projects around the campus. This system would later be called Galatea.

The other piece to be changed from local access to network access was the database retrieval system. In the original system, the program simply read the information about the available images out of a file in the computers file system. To provide remote access to this database and to prevent duplicate copies of what could now be a very large file, a database search server was created. This server responded to requests over the network for information about the images that were available in databases pertaining to architecture, art and documentary video.

At the same time as the database server and Galatea were developed, the decision was reached to rewrite Light Table from scratch to utilize the new version of the X Window System, version 11, which was established as an industry standard, and which would permit work with the latest equipment from many vendors. This new Light Table required virtually no local resource, but used network servers almost exclusively. The graphics output was through the X Window System, using network capability to retrieve information from the custom database server. The videodiscs were controlled through the Galatea Network Video Device Control System.

Galatea

Galatea (Applebaum 1989) was designed not only to provide remote access to centralized video devices, but also to handle a distributed arrangement of video resources. Shared devices did not all need to be located on a "master' server, but could be located at several sites. This distributed flexibility was essential in the Project Athena environment, since each videographic workstation usually had its own local video resources, in addition to its access to the remote central server. In order to control the local resources, usually a videodisc player and an audio/video switch, a Galatea server was run on the videographic workstation. This server was also capable of forwarding requests for video resources to other Galatea servers on the network, the most common of which was the primary campus video server, located centrally, in the Retch Visual Collection.

Naturally, keeping track of how all of the video devices were wired together was one of the main tasks of a Galatea server. Each server managed a collection of videodisc players and audio/video switchers. A Galatea server was capable of providing several outputs. If the audio/video switch in a system could handle such routing, a single videodisc server could be configured

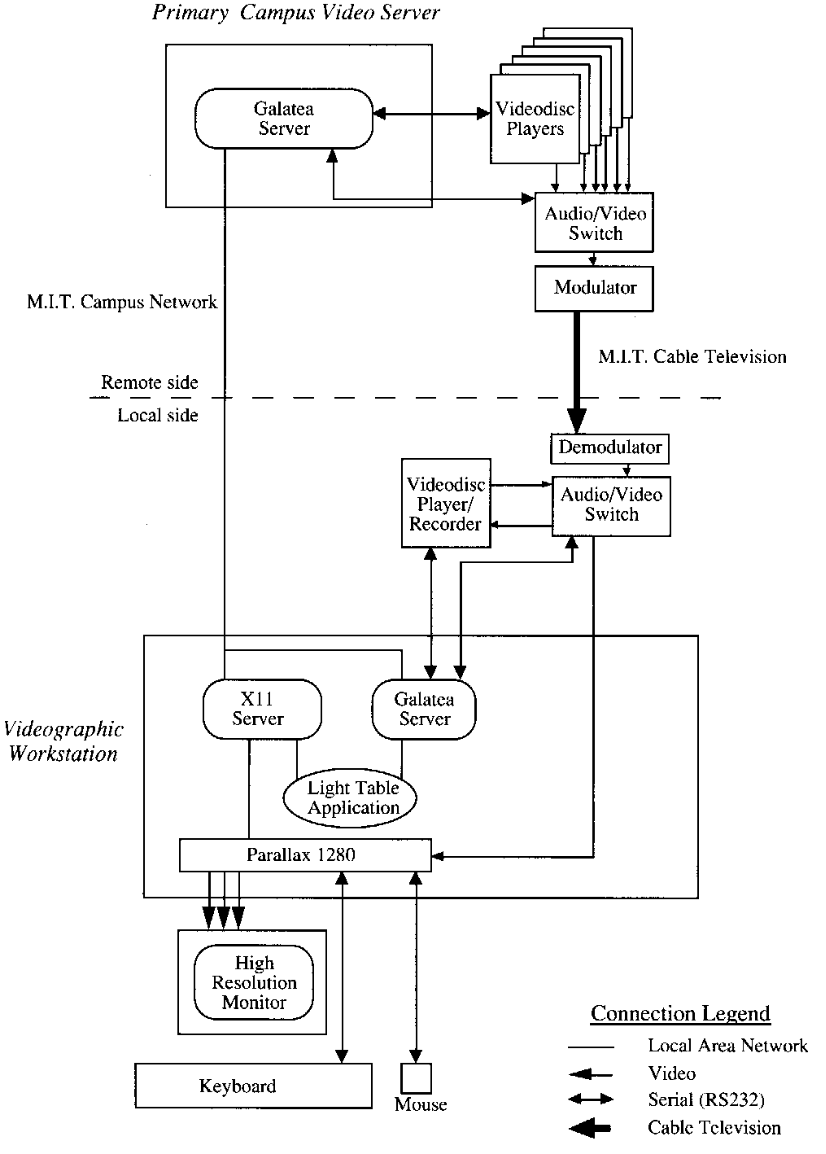

Figure 5 Diagram of Light Table system including Galatea

to send the output of any of several videodisc players to any of several different destinations. By supporting multiple 'virtual server outputs, a single server could support several channels on a cable television system, or several displays on a single videographic workstation. In order to make the configuration even more flexible, switchers could be chained. An output of one switcher could feed an input to another switcher. For example, two four input, one output" switchers could be combined into one "seven input, one output' switcher by passing the output of one of the switchers into an input on the second switcher.

All of the details of the routing system were hidden from client programs, such as Light Table. When a client first connected to a Galatea server, it established which of the server's "virtual" outputs it wished to use. The server then passed the client a list of disc volumes that could be routed to the specified output. Once the client had this information, it needed only to indicate to the server which volume was desired, and the server would handle all of the specific audio and video routing.

The clients task was also simplified because a Galatea server could actually combine several copies of a single disc into one volume. If two copies of the Archfile disc were present, the server would provide the clients with the information that the Archfile volume was available. Normally, the server would guarantee that the multiple copies of the disc would appear to all the clients as just one. For a frame search, however, the Galatea server would first query all of the copies of a volume as to which player containing a copy of the disc was closest to the requested frame. To reduce the time needed to perform a search, the server would then use the closest player.

A major effort has been made to enable the Galatea servers to compensate for device failures. Network crashes, workstation crashes, failure of videodisc player and failure of audio/video switch were all handled automatically, with no intervention necessary from an operator. The usual method of dealing with the indication of a video device failure was for a server to clear all its internal device tables and rebuild its device information from scratch. It then proceeded to notify any other servers and clients that relied on its volume information that a volume table rebuild took place. Failures of networks and workstations required much more complicated compensating techniques. When a server served other Galatea servers, it maintained a file containing the network addresses of those other servers. If its workstation crashed and was rebooted, the Galatea server would notify each of the servers listed in that file that it was back on line. This technique could compensate for workstation crashes. Network crashes were slightly more complex. When a local server detected that a remote server became unavailable, it would undertake a volume table rebuild. If the remote server did not respond during the rebuild, it would not be recorded in the volume table. However, the local server would periodically check to see if the remote server was available by sending out small network packets requesting a response. If a response was received, another volume table rebuild would be performed so that the remote server, with its associated video devices, could be used again.

Although originally designed for the Light Table project, Galatea has been extended to perform well in many other contexts. Galatea can handle live video control requests, seamless splicing, and recording onto write-once video media. The live video requests can be used to view discs which contain motion video segments. The seamless splicing code can sequence several segments and, through a series of pre-rolls and video switches, can provide the seamless playback of the segments. The recording functions can be used to do frame by frame recording of computer generated animations. Alternately, the record functions can be used to take snapshots of the workstation user to send to other persons, or to copy a personal video tape into the main storage for use by others.

Galatea is an attempt to standardize access to video resources, so that various video client programs can work together on a single workstation, and so client programs are portable from site to site. Currently there are about one dozen universities and corporations using Galatea as the video device control system for their research. Galatea's built-in management of the various types of failures is a significant improvement over previous server designs. In summary then, it can be said that the Light Table interface, combined with Galatea and the database server provide reliable image delivery and effective portability. Galatea, which has grown to become an entire project, was the key to creating a system with these characteristics. The power of Galatea to provide disparate forms of video access was also essential in its deployment, so that many projects can share video resources for long periods without human intervention.

Acknowledgements

The concepts of the visual information system and the development of a versatile user interface to such systems owe much to our colleagues at MIT and elsewhere. Osnat Ron was the developer of the prototype Archfile system. Morissa Miller, Carter Pfaelzer and Paul Paternoster developed the second generation of Archfile. Thanks to Henry Okun, author of Picassofile, whose art historical erudition was the basis of the Picassofile data base. A special mention for Merrill Smith, originator of the Boston Project, and erstwhile Head of the MIT Retch Visual Collections. Thanks to Fang Pin Lee for sustained contribution to the Boston Project. Christopher Thorman made many contributions including the "Image Delivery System". Several persons have helped in the development of Galatea. Hal Birkeland has made considerable improvements and corrections to the server implementation. With Ben Rubin, he wrote the code for seamless splicing in the system. The input of Ben Rubin and Paul Boutin has led to the implementation of many features in the current version. Russ Sassnett provided the model and most of the underlying device control system. The special database server for networked environments was a creation of Jenifer Tidwell. Overall, Glorianna Davenport's constant encouragement, ideas and suggestions were essential to the proper development of the system. Light Table and its server Galatea have been sponsored by MIT Project Athena. The facilities provided to Light Table project by Muriel Cooper and colleagues in the Visible Language Workshop and by other colleagues in the MIT Media Laboratory were greatly appreciated.

References

Applebaum, DI. 1989. 'The Galatea Network Video Device Control System". Cambridge: MIT Media Laboratory.

Leffler, SH, MK McKusick, MJ Karels, and JS Quarterman. 1989. The Design and Implementation of the 4.3BSD UNIX Operating System. Reading MA: AddisonWesley.

Lippman, A. 1980. "Movie Maps: An Application of the Optical Video Disc to Computer Graphics". In SIGGRAPH '80 Proceedings. New York: ACM.

Mackay, WE, and C Davenport. 1989. "Virtual Video Editing in Interactive Multimedia Applications". Comm. ACM Vol 32/7. New York: ACM.

Murman, EM. 1989. 'The Athena Project: Goals. Philosophy Status and Experiences". MIT Athena Report. Cambridge.

Purcell, PA, MW Smith, and CP Thorman. Sept. 1986. 'The Image Delivery System". Report to the Council for Library Resources. Washington.

Purcell, PA, and H. Okun. 1983. "Information Technology and Visual Images: Some Trends and Developments". In proc coot International Federation of Arts Libraries. Munich.

Scheifler, RW, and J. Gettys. April 1987. "The X Window System". ACM Transactions on Graphics 5(2) pp 79-109. New York: ACM.

Smith, MW. 1987. "The Boston Project". Internal Paper. Cambridge: MIT Dept of Architecture.