

Digital Display Integration Project Project Outline 2.0

J.N. Bardsley

November 1, 1999

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401
<http://apollo.osti.gov/bridge/>

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161
<http://www.ntis.gov/>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Lawrence Livermore National Laboratory
Information Science and Technology Program
Digital Display Integration Project
Project Outline 2.0
November 1, 1999

1. Introduction

The electronic display industry is changing in three important ways. First, the dominance of the cathode ray tube (CRT) is being challenged by the development of flat panel displays (FPDs). This will lead to the availability of displays of higher performance, albeit at greater cost. Secondly, the analog interfaces between displays that show data and the computers that generate the data are being replaced by digital connections. Finally, a high-resolution display is becoming the most expensive component in computer system for homes and small offices. It is therefore desirable that the useful lifetime of the display extend over several years and that the electronics allows the display to be used with many different image sources. Hopefully, the necessity of having three or four large CRTs in one office to accommodate different computer operating systems or communication protocols will soon disappear. Instead, we hope to see a set of flat panels that can be switched to show several independent images from multiple sources or a composite image from a single source.

The more rapid rate of technological improvements and the higher cost of flat panel displays raise the incentive for greater planning and guidance in the acquisition and integration of high performance displays into large organizations, such as LLNL. The goal of the Digital Display Integration Project (DDIP) is to provide such support. This will be achieved through collaboration with leading suppliers of displays, communications equipment and image-processing products, and by greater exchange of information within the Laboratory.

The project will start in October 1999. During the first two years (FY2000-1), the primary focus of the program will be upon:

- introducing displays with high information content (over 5M pixels)
- facilitating the transition from analog to digital interfaces
- enabling data transfer from key computer platforms
- incorporating optical communications to remove length restrictions on data transfer
- optimizing techniques to reduce communications bandwidth

Additional topics that will be addressed include

- obtaining greater color control
- minimization of seam widths in tiled displays
- the human factors that impact the efficient use of high resolution displays

The scope of the project will include both direct-view displays, for use by individuals and small groups, and projection systems for conference rooms and auditoriums.

The results of the project will be shared with other laboratories of the Department of Energy (DOE). Cooperation will be sought with other agencies of the U. S. Government, such as the Space Agency (NASA), the Department of Defense (DoD) and the Intelligence Agencies, both to avoid duplicative efforts and to pass on the knowledge gained during the study. In particular, we anticipate that most of the human factors work will be led by Dr. Latimer's group at the NASA Ames Laboratory.

2. Displays with High Information Content

The most important factor controlling the amount of information that can be shown on a flat panel display is the pixel count. As of mid-1999, most manufacturers offer SXGA panels, with a 1280x1024 array of multi-colored pixels. Although transitions between generations (VGA->SVGA->XGA->SXGA) have been occurring about every two years, bringing steadily increasing pixel count, the rate at which further growth will occur is uncertain. Prototype panels have been built with up to 7M pixels, but the manufacturers are hesitant to commercialize such displays. The market for such displays is not yet large and the problems of creating and transporting such complex images are daunting.

The visualization of numerical simulations, carried out on massively parallel computer systems, is one application that demands high-resolution displays. Since the ability to view several complex images at the same time can greatly increase the productivity of a programmer or analyst, display sets with multiple high-resolution panels can be extremely effective. Other important applications at LLNL include CAD/CAM, emergency response centers and the publication of technical documents.

The minimum short-term goals of the DDIP program are:

Year 1:

- acquire, integrate and test single display panels with at least 3M pixels
- install and feed display sets with 4 tiled panels in individual offices
- contract for the acquisition of prototype displays with over 5M pixels

Year 2:

- integrate and test single display panels with at least 5M pixels
- install and feed display sets with at least 20M pixels in individual offices
- acquire and integrate a single projection system with at least 20M pixels
- demonstrate a multi-purpose high-resolution display with simultaneous connections for scientific and engineering work, document preparation and retrieval, video conferencing and viewing high definition television.

It is anticipated that the direct-view displays that will be used in this project will be standard commercial liquid crystal displays (LCDs). Stimulating earlier availability of these displays is an important goal of the project. It is anticipated that during the first year, our work will concentrate on resolutions of 1280x1024 (SXGA), 1920x1200 (HDTV+) and 2048x1536 (QXGA).

We believe that it is unlikely that the projection systems that are required to meet our goals will be available from leading commercial suppliers on the desired schedule. Thus the projection display component may involve a research partnership to accelerate the development of innovative systems.

3. Digital Interfaces

Recent progress suggests that two of the major obstacles to the widespread adoption of digital interfaces may soon be removed. The first is the confusion caused by the promulgation of multiple candidate interface protocols. For example, the Japanese Electronic Industry Association (JEIDA) set up a committee to investigate Digital Interface Standards for Monitors (DISM). This group endorsed several competing approaches and their report appears to have had little influence. By virtue of diligence and good judgment, the Digital Display Working Group appears to have established de facto leadership and its work has been blessed by the Video Electronics Standards Association (VESA), whose own first attempt at standard creation was less successful.

The second barrier to the adoption of digital standards has been the small penetration of FPDs into the monitor market and the exclusive use of analog interfaces by the CRT industry. However, sales of FPD monitors have grown ten-fold in the past two years and major CRT manufacturers have introduced monitors that will accept digital signals.

The initial promoters of DDWG were Intel, Silicon Image, Compaq, Fujitsu, Hewlett-Packard, IBM and NEC, but their efforts have been endorsed subsequently by several other companies, including some who had made previous commitments to alternative standards. On April 2nd, 1999, the group issued their first detailed specification, entitled Digital Video Interface (DVI) 1.0, which incorporates many individual elements of earlier VESA standards. Using a single 3-channel link, this implementation provides input rates of up to 165 Mpixels/sec, which is sufficient to support LCD displays with up to 2M pixels, refreshed at 60Hz. At this frame rate two links will be needed for 3 Mpixel displays. It is anticipated that in version 2.0 the single-link limit may be raised to around 230 Mpix/s.

In the LLNL DDI project, we will recommend compliance with the DDWG standards and will work with vendors to accelerate the availability of more powerful implementations of these protocols. For example, although DVI 1.0 appears to be compatible with the use of optical fiber, there is no consideration of the implications of the protocols for long-distance connections. Furthermore, no implementations of the dual link architecture have yet appeared.

The initial DVI 1.0 standard seems to be unnecessarily constrained by the legacy of previous VESA standards, which do not appear to allow systems designers to fully exploit the advantages of flat panel displays. We will work with our industrial partners to explore the benefits of extensions of current standards to facilitate the use of high-resolution displays.

Specific first-year goals with respect to interfaces are:

- demonstrate a 1 km fiber optic connection capable of driving a 3 Mpixel display, preferably using no more than two fiber strands
- implement a single-link DVI system capable of driving a 3 Mpixel display
- design a dual-link DVI systems capable of driving displays with up to 8M pixels

4. Image Sources

In respect to legacy systems, the highest priority will be given to enabling data to be transferred from powerful graphics engines. These include:

- SGI Onyx 2, with Infinite Reality pipelines, running proprietary UNIX (IRIX).
- Compaq AlphaServer ES40, with PCI buses, running under Red Hat Linux

These systems are deployed in large clusters and will be used to drive multiple displays. There is also a large installed base of high-end work-stations, using both Microsoft and Unix operating systems.

Although the designers of the Onyx and AlphaServer systems have pioneered the development of scientific visualization systems, it is likely that future architectures will rely more heavily on the hardware and software being produced by the PC and video game industries. During the first year, we will collaborate with the LLNL Center for Advanced Scientific Computing (CASC) in the design of next-generation visualization systems and expect to deploy prototype systems in FY 2001.

Driving a single 3-Mpixel display at 60 Hz, with 24 bits of color data in each pixel, will require the generation and transmission of over 500MB/s. Even more would be needed if texturing information were to be passed through to the display controller. This bandwidth is beyond the capability of most single processors and communication channels. To ameliorate to problem during the first year, the creation of full screen video will be limited to 5-10 frames per second. This will require that the display controller can buffer the data and increase the refresh rate to the desired refresh rate for the display panel. The video creation rate should be increased to around 30 frames per second in subsequent years. Alternative compression schemes transport must be investigated, along with faster means of data creation and transport.

For office work, there is also a need to drive displays from workstations. Initial priority will be given to work stations using Windows-NT software. Preparations should also be made for the introduction of Windows 2000 and the Intel IA-64 chip.

We anticipate that the rate at which video information can be shown on the screen will be limited not by the optical characteristics of the display, but rather by the limited electronic bandwidth, in either the generation of the image by a computer, the transmission of the data from computer to display, or in the controlling electronics within the display. We therefore intend to seek opportunities for R&D partnerships to develop new display architectures and driving systems, so that the display controller and communications electronics can be updated without changing the optical structure of the panels.

5. Remote Visualization

The initial implementations of digital interface protocols are designed for the connection of displays to personal computers or workstations in the same room. The resulting limitations on cable lengths will be too restrictive for our needs, even to connect a cluster of graphics engines to a video wall in the same room. Thus we will give high priority to the development of digital fiber optic connections, requiring transmitters, receivers and signal boosters as well as the cable itself.

The LLNL site is one mile square and so cable lengths up to 2 km may be needed to allow point-to-point connections across the laboratory. In addition, the DOE is keen to avoid unnecessary duplication of expensive computer equipment at its laboratories in California and New Mexico and is arranging access to high bandwidth communication links.

In year 1, we will provide only dedicated point-to-point connections, preferably with simple circuit switches to optimize the use of available computer ports. In year 2 we will investigate the possible use of local and wide area networks, for example using ATM or Gigabit Ethernet systems. This should allow us to drive displays in California from computers in New Mexico, or vice-versa.

Some of the motivation of this program is provided by increased concern within DOE over information security and earthquake safety. The use of office displays to view images generated from storage media or computers in secured vaults is expected to be an important component of our new information systems. If users are able to rely on high-resolution displays to read, produce and edit documents without recourse to paper and floppy disks, they will no longer have to worry about locking up sensitive documents before they leave their offices. The need for security officers to examine brief cases for paper documents and other portable media will also be greatly reduced. Replacing heavy CRTs by LCDs and eliminating bookshelves full of thick manuals and text-books will mitigate earthquake hazards.

Although the transition to paper-less offices seems to be finally underway, the quality of electronic displays is still retarding the enthusiastic adoption of the concept. In conjunction with our colleagues at NASA Ames and our industrial partners, we will attempt to identify and remove the obstacles to complete reliance on electronic displays for access to long texts and complex graphic databases. Greater adaptability of fonts, icons, cursors and multiple windows will be investigated. Glare reduction and optimal screen placement may also be important.

6. Additional Goals

6.1 Bandwidth Management

Providing video signals at 60 frames a second for a 5M pixel display taxes the output ports of even the most expensive graphics engines as well as the communications systems. In principle, there are several obvious means to reduce this bandwidth, but the implementation of each will lead to conflicts that must be resolved.

- Reducing the bandwidth - Very few, if any, of our users will want to refresh the screen at more than 30Hz, except perhaps to produce stereo images.
- Delta (selective) transmission - Since most of our applications call for static images, it seems wasteful to send the data describing unchanged pixels all the way from the host computer. Electronics to accept only new data can easily be included in a display controller and may even be incorporated into the pixel structures. The latter solution may require poly-silicon pixel transistors and so may not be implemented within the first two years of this program
- Data compression - Effective data compression techniques will greatly extend the usefulness of the ultra-high resolution displays. However, our users often wish to examine minute features of complex images and so are very concerned about information loss during compression. Loss-less techniques may introduce unacceptable timing uncertainties into the transmission. One solution could be to introduce algorithms that warn the user when losses have occurred and provide convenient mechanisms for the user to obtain more reliable images at a slower rate.
- Adaptable display addressing - Bandwidth can be saved by taking into account the compromises made in the human vision system. In principle, the use of head-tracking systems may allow one to benefit from the different spatial resolution of the eye in the foveal region and elsewhere. However, we are skeptical about the general acceptance of these systems and they are difficult to implement for multiple viewers. Systems in which the spatial resolution and gray scale are adjusted dynamically depending on the nature of the images may be more practical. For example, lower level of microscopic detail and color content may be acceptable in rapidly changing images that in the display of stationary text, graphics or photographs.

6.2 Color Control

Lack of control over color has been a chronic problem for the electronic display industry. It is obvious whenever one looks at a set of displays in a retail showroom or in a video wall. This is a problem that is shared by the printing industry, causing difficulties for those involved in electronic publishing or advertising and leading to many returns by internet shoppers who receive goods of an unexpected shade.

LLNL currently uses an array of 15 SXGA rear-view LCD projectors to create a composite image of 20M pixels in a conference room. Although seam thickness can be minimized by proper positioning of the projectors, the color intensities drift so much that the edge of each separate image is often clearly observable. Manual adjustment of the 15 projectors is time consuming and unsatisfactory.

If funds are available and suitable partners found, we would like to design, install and test automatic color sensing and feedback mechanisms that will allow real-time adjustments to be made of the color scales in each projector. We also intend to encourage the development of systems to facilitate color control during electronic printing, so that the color of a printed document is close to that seen in the screen preview, irrespective of the technology or manufacturer of the printer and display.

6.3 Seam Minimization in Tiled Displays

We intend to deploy sets of displays in conference rooms and individual offices. Preliminary analysis suggests that seamless tiling of LCD panels will be difficult and expensive. One relatively inexpensive technique may be proximity projection, in which the images from individual panels are expanded by a small factor onto a common front plate. We believe that for most applications, the motivation to completely remove the seams will diminish as the number of pixels in each panel increases. This assertion needs to be tested.

In the meantime, we will work with suppliers to obtain panels with minimum extrusion of bevels and driver electronics and will mount the panels so that the seams are as unobtrusive as possible. We will try to determine the conditions under which it is better to maintain continuity of diagonal lines that cross seams, even though this necessitates the suppression of part of the image.

7. Resources.

We expect that the budget for the first year of this program will be around \$1M. This does not include the much larger cost of the many graphics engines that will be used in the generation of the computer images. In-kind contributions, in the form of equipment or collaborative effort, will be extremely welcome and will increase the value of the project to our industrial partners.