ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, IL 60439

The μMural: A Six-Projector Tiled Display

by

Ivan R. Judson

Mathematics and Computer Science Division

Technical Memorandum No. 251

June 2001

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-ENG-38.

Contents

Table of Figures

List of Tables

The μMural: A Six-Projector Tiled Display

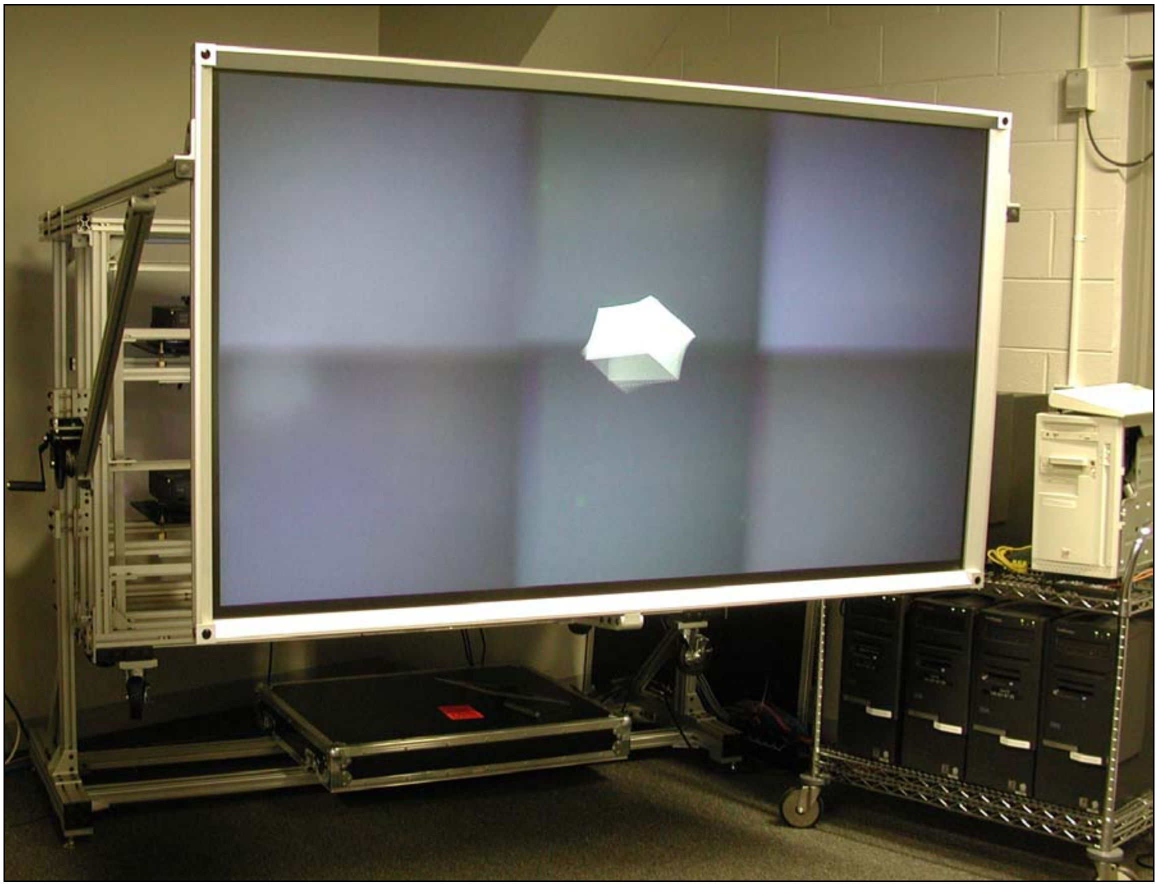

Ivan R. Judson
Futures Laboratory
Argonne National Laboratory
(fl-info@mcs.anl.gov)

Abstract

Tiled displays have become a recent technical solution to aggregating commodity displays in order to provide higher resolution displays. This document describes the background, design, and implementation of the micromural, a six projector tiled display developed at Argonne National Laboratory.

Overview

It is easy to generate data. In the past few years there has been an explosion in the amount of data produced, from scientific simulations generating terabytes of output, to new digitizers that can capture environmental data at very high speeds and resolutions, it is now easier than ever to create large amounts of data. In order to see this quantity of data, large displays are needed. Examples include the 2D Cellular Detonation simulation at the ASCI FLASH Center. This simulation is currently producing data that is 128x128x34000 data points in size [1]. This data is too large to be viewed on any existing display device, and it is only one of a large set of current computations that need to be explored.

Table 1: NVIDIA GeForce2 MX Display Resolutions.

Width (Pixels)	*Height (Pixels)*	*Depth (Bits)*	*Refresh Rate (Hz)*	*Total Bandwidth (MBytes/Second)*
640	480	32	120	148
800	600	32	120	230
1024	768	32	120	378
1280	1024	32	120	629
1600	1200	32	85	653
1920	1440	32	60	664

Consider the available desktop resolutions of a NVIDIA GeForce2 MX AGP graphics adapter, and the resulting uncompressed bandwidths shown in Table 1. The GeForce2 MX is already a generation or two old and can be purchased for less than $200. Current desktop display devices are show in Table 2. These two forces together show that it is currently feasible to interact with at most 1920x1440 pixels at a time; however, the current desktop trend is to move toward flat panels. Using these displays, it is currently feasible to use only 1024x768 resolution displays on the desktop. All of this data shows that, currently, the rate and scale at which data can be generated are far outpacing the rate and scale that we are developing higher-resolution desktop display systems.

The current landmark desktop display, IBM's Bertha, is a 3840x2400 pixel, 22-inch display, offering 200 pixels per inch, or 9.2 megapixels. This represents significant progress toward developing desktop displays that can show data in the quantities we used to generate. However, during the past six years, CPUs improved a factor of 13, memory improved a factor of 16, hard disk drives improved a factor of 20 and modems improved a factor of 13, whereas displays improved only a factor of 3. Thus it will take some time for displays to catch up in performance. For the foreseeable future, these display devices will not scale to the size required to view data that can be generated today. Unless the

Table 2: Current Desktop Display Technology Overview.

Width (Pixels)	*Height (Pixels)*	*Diagonal (Inches)*	*Dot Pitch (mm)*	*Type*	*Price*
1024	768	14 – 16	0.30	LCD	$700 - $1000
1280	1024	16-18	0.25	LCD	$2000 - $3500
1600	1200	17.3	0.23	LCD	$1500
1600	1024	22	-	LCD	$3000
Up to 1920	Up to 1440	14-24	0.20 - 0.28	CRT	$150 - $2000

rate of progress in display development increases to a rate faster than the ability to generate data, there may never be a desktop display capable of viewing the data sets that will be produced.

Since single displays do not have enough pixels to display current data sets, the only way to get enough pixels is to use multiple current displays and tile them to build a large display. This technique, similar to the Beowulf Linux cluster effort [2] to build large computers, has been used in the commercial display market for quite some time. However, the commercial display market seldom has to address the need for more pixels; instead, the goals include a large viewing distance. In order to build displays that are viewable in sports arenas, convention centers, and other large venues, multiple displays are usually tiled. Video cubes were invented to minimize the borders between tiles, while still providing relatively good manufacturability. Current video cube manufacturers include Synelec and Clarity Visual Systems [3]. The two problems with these solutions involve signal generation and cost. Commercial tiled display technology based on video cubes uses "signal processors" which take multiple video signals in and put out a signal for each video cube. These processors are usually architecturally limited in functionality, and also suffer bandwidth limitations. Moreover, these processors are expensive, ranging in price from $5000 to $100,000. Typical Synelec Video Cubes, at 1024x768, range in price from approximately $30,000 to $35,000 per cube, depending on size.

Table 3: Cost of Tiling Commodity Video Cubes.

Width (Tiles)	*Height (Tiles)*	*Total Resolution (MPixels)*	*Total Cost (No Processor)*
3	2	5	$180,000
5	3	12	$450,000
6	4	19	$720,000
8	8	50	$1,920,000

Recently, the video processors and software have been attacking the bandwidth and architectural issues. Synelec offers X11R6 and Windows solutions for video cube based walls. Synelec is also developing versions of the system that incorporate a more distributed model for content generation, although a single box is still needed to control the display. Even with these innovations, however, price is the major factor keeping the scientific community from utilizing this technology. A cost breakdown of various configurations is shown in Table 3, utilizing 1024x768, $30,000 video cubes.

To provide displays with large numbers of pixels at a cost that is scalable, we are contributing to a community of researchers who have been building tiled displays from commodity components. Participants in this community include Argonne National Laboratory [4, 5], Princeton University [6], Stanford University [7], the National Computational Science Alliance, and the University of North Carolina at Chapel Hill [8]. We have designed and fabricated various tiled display systems. These prove to be cost-effective, scalable solutions that provide users with more pixels than ever previously experienced. The μMural, described in this report, is one of these displays, with 3.5 megapixels provided by six projectors in a 3x2 configuration, with a comprehensive cost of approximately $46,600, as summarized in Table 4.

The µMural screen is 76 x 44.5 inches, with a viewable area of 72 x 40.5 inches. Additional framework makes the overall width 85 inches. The screen can be raised and lowered so that the distance from the bottom of the screen can change from 14 to 40 inches. The adjustable height makes the µMural usable both for standing with a group of 4-6 people or sitting with 2-3 people. The tiles are overlapped and blending with hardware that resides in the path of light between the projectors and the screen. This design does not use a large amount of space between the last blending layer and the screen, which is a distance of 44.5 inches. This µMural invites small groups of people to gather around it and discuss the content that is currently being displayed, rather than huddle around a monitor where the pixels are comfortable only when viewed from a relatively small distance by a single person.

Table 4: Cost Breakdown for the µMural.

Part	*Cost Per Unit*	*Quantity*	*Net*
Screen	$3,000	1	$3,000
Projectors	$4,250	6	$25,500
Positioners	$600	6	$3,600
Frame	$2,500	1	$2,500
Computers	$2000	6	$12,000
		Total	$46,600

This document describes the design, fabrication, and assembly of the µMural, including the pixel generation options available for driving tiled displays. We discuss what software has been modified to drive the µMural and what software has been created to support applications for tiled displays. We conclude by analyzing how well the µMural has accomplished the goal of providing a usable, cost-effective, scalable solution for creating large display systems.

Application and Design Goals

Over the past 15 years the supercomputing industry has gone through several contractions, it has shrunk to a small number of vendors who do not rely solely on supercomputers for revenue. This trend, which leaves Cray Research, IBM, and SGI as the only three large system vendors, has been further accelerated by the fact that Beowulf-class supercomputers are gaining popularity. With this popularity researchers are producing more software that makes these machines easier to use. This effect, known as the "Beowulf Effect" was predicted before it was even realized, and a similar effect applies to tiled display technology.

Difficult system architecture issues need to be addressed when designing a megapixel display system; A common challenge between building monolithic megapixel display systems and monolithic supercomputers is the difficulty of obtaining high-quality parts when the manufacturing process is pushing the limits of technology. In particular, liquid crystal display manufacturing has this problem to a greater extent than microchips, because of the size of the manufactured device. Bertha, which is facing these manufacturing challenges, is a good sign that manufacturers are trying to overcome them. To overcome the bandwidth limitation, Bertha is currently organized as tiles, each accepting a DVI video signal.

Another key challenge common to both display systems and supercomputers is bandwidth. For displays the issue is exacerbated by the fact that the video industry standards have not been looking at wall-sized displays. DVI, the latest among these standards, does not address displays larger than 1920x1080 (HDTV) at 60 Hz for a single-link DVI connection and 2046x1536 (QXGA) for a dual-link DVI connection [9]. However, the DVI specification is assuming that greater pixel resolutions are desired, thus, they incorporated the second parallel DVI link. This parallelism in the designs is a key to the future of megapixel displays.

Tiled Displays

The parallelism of the DVI specification and that in the Bertha hardware is a glimpse at how megapixel displays are going to be developed in the future. Bertha, at 9.4 megapixels requires 4 QXGA DVI signals. Imagine a graphics card that could produce 9.4 million pixels at 60 Hz: that is 564 million pixels per second with 24 bits per pixel that totals 1.7 gigabits per second. In order to build a 100 megapixel display, tiles will need to be used for the foreseeable future, since single display devices are not scaling fast enough to reach tens of megapixels any time soon; even Bertha will not be "commodity" for 3-5 years. Hence, two challenges must be tackled to keep the performance of the tiled display systems high enough to make them usable. First, rendering solutions need to be able to scale to the size of the display, and second, communications infrastructure needs to scale at a rate that makes it possible for rendering solutions to send data to the display devices.

Rendering, the process of taking a data and generating pixels that can be displayed, has made great strides in performance. It is now commonplace to view models with hundreds of thousands of polygons, and volumes of megavoxel sizes and to interact in real-time with these visualizations. However, these renderers usually live on a single machine, generating on average a megapixel stream, directly connected to the display via an analog cable. The renderers also generally contain all of the data in memory, so that access times are minimized. Currently, the huge amount of data cannot be rendered interactively: Either it won't fit into physical memory, or the visualization algorithms are computationally complex and thus require more time to operate. A common approach to solving this problem is to create parallel rendering solutions. Investigations into the scalability of these solutions are under way by various researchers in the visualization community [8, 10-18].

The other challenge to tiled display devices is scalable communications networks. A single current data stream of 1024x768 pixels, with 24 bits per pixel at 60 Hz, generates 1.1 gigabits per second. This uncompressed data stream is usually transformed to reduce the needed bandwidth of the network. VNC [19] has developed protocols to support the notion of a remote framebuffer and can bring this bandwidth down far enough to be usable via standard dialup connections. With new network technology such as Myrinet and 10 Gigabit Ethernet, increasing network bandwidth and decreasing latency continue to be goals driven by the market at large. This market-driven approach is one of the advantages of the cluster approach to both computing and display devices.

The μMural

The μMural is a display device that can be used by a single person or a small group, in an office setting or at a trade show. One reason the μMural was built was to satisfy the need for a portable tiled display. At Argonne National Laboratory the *Active*Mural is a larger tiled display, 18 x 8 feet, with 15 projectors tiled to form a 12 megapixel display. This display consumes the larger part of a room and weighs nearly a ton. Constructed from large steel beams that hold 4x8 sheets of screen material, it is far from portable. In contrast, the μMural can be packed into a 30.5” (D) x 53” (H) x 85” (W) shippable container and air freighted anywhere in the world. The μMural is designed to be simple to assemble and adjust, it does not try to be the most compact design, which might be more difficult to fabricate. It also provides more pixels per dollar than any existing display solution.

Hardware Overview

The major hardware components of a tiled display include the, projectors, positioners, screen, framework to hold these components, blending hardware, and the computers that generate the video signals. In deciding on what specific components would be incorporated into the μMural, we considered durability, transportability, ease of replacement or repair, and the ability to incorporate design modifications. Some design challenges were left for future work, such as invertible projector positioners.

Positioners

In order to make a tiled display seamless, each tile of the display needs to be pixel aligned with it is neighbors to avoid discontinuities that cause the human eye to perceive the tiles rather than the display as a whole. This can be done with software or with hardware. MIT and UNC have developed software approaches to distort the images projected by each projector, using well-calibrated cameras [20, 21]. Princeton has extended this work to work with uncalibrated cameras [22]. These techniques are generally applied to projectors that have been grossly aligned with physical positioners. UNC’s Office of the Future Group are exploring software solutions that do not require this gross alignment and thus handle alignment completely in software [17, 21].

The benefits of the software solution include speed in setup and reduced cost of hardware. However, the cost of software solutions is usually resolution. The software remaps the existing pixels in ways that are globally optimal for the display, which costs pixels on each local tile. In contrast, the physical positioning system involves more setup time and increased manufacturing costs, but it does not require that pixels be lost on each tile. The optimal solution involves both of these methods: physical positioning is used to bring the projectors into alignment, where the only misalignment is caused by the optics of the projectors, and then software warps each tile to invert the effects of the optics of the projector.

The positioners that are part of the μMural solve the physical positioning problem in a very cost effective manner. The stability, accuracy and repeatability obtained is excellent. These positioners can be adjusted at 1/10 of a pixel resolution in each of the six degrees of freedom, at a distance of 84 inches from the screen. (Pixels on the μMural are slightly smaller than 0.05 inches.) The positioners are based on a design developed by Intel and Princeton; however, it was heavily modified to be more accurate, stable, and less costly to manufacture. The design includes direct manipulation of knobs that control each of the six degrees of freedom; these knobs are located in places where the relationship between adjustment knob and effect is intuitive. The positioning is done via kinematic contact points that ensure repeatable and stable adjustments.

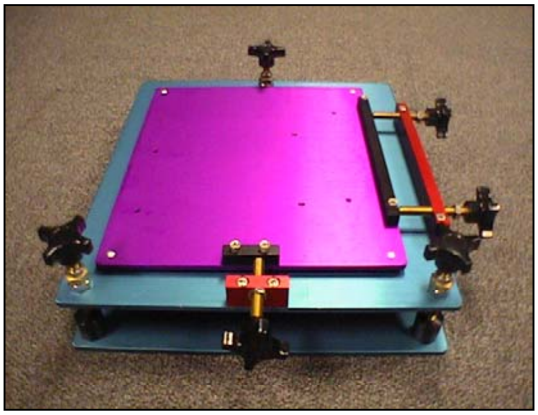
Figure 1: Projector Positioner Designed and Fabricated by Argonne National Laboratory.

Current commodity projectors seem to be asymptotically limited to minimum cost of about $4000. The positioners developed for the μMural cost approximately 10% of this, or $400. The μMural positioners are composed of a small number of inexpensive parts, where the current largest cost of manufacturing is labor. (Some of the subcomponents of the design should be available from various manufacturers.) This design is freely available for others to obtain, modify, and manufacture, as has been done by SDSC for the High Resolution Display Wall [23]. Blueprints for this positioner are included in Appendix C, Figure 13-Figure 18.

The stability of the projected image relative to the fixed the screen has been measured over a range of time scales from 1/15 of a second up to 5.5 hours. Experiments to extend these measurements to days, weeks, and months are under way. On these timescales, the combination of shelving, positioners, and screen frame proves to be very stable. The upper limit of the drift appears to be 0.005 inches/hour at the screen (or less than about 0.1 pixel/hour).

Projectors

In a tiled display each tile must have some device that converts the pixels to light. Bertha uses tiled LCD panels; the μMural uses commodity projectors. The advantage of projectors over LCD panels is they can be tiled with abutted edges or overlapped with no discontinuities on the screen. A wide variety of projectors are manufactured. At the low end, many of these are actually the same projector subsystem, manufactured by one of the few subsystem manufacturers, Plus, Epson, and Sharp. The international InfoComm conference is a good source for the latest in projector technology [24]. Projectors can be categorized into three classes: (1) low-cost LCD or DLP projectors designed for the high-volume portable business-presentation marketplace, (2) multiple-chip LCD or DLP projectors designed for permanent multimedia installations, and (3) high-end projectors designed for large-venue professional imaging marketplace using a variety of high-quality and high-performance image modulation technologies (CRT, LCD, DLP, Lightvalve, etc.). Another way to distinguish between these is by price: less than $8,000, $8,000-$20,000, and greater than $20,000. Properties important for projectors include resolution, brightness, color gamut, data interfaces, image quality, optical quality, noise, remote control interfaces, software controls, calibration and configuration control, stability (zoom, focus), image refresh rate and support for stereo, resampling algorithms (e.g., automatic resolution conversion), flatness of illumination field, keystone correction or off-axis projection capability, color convergence, image alignment or adjustment and power, cooling, weight, and size.

Current low-end projectors support native resolutions of 1024x768 pixels, with a move to 1280x1024 expected in the near future.

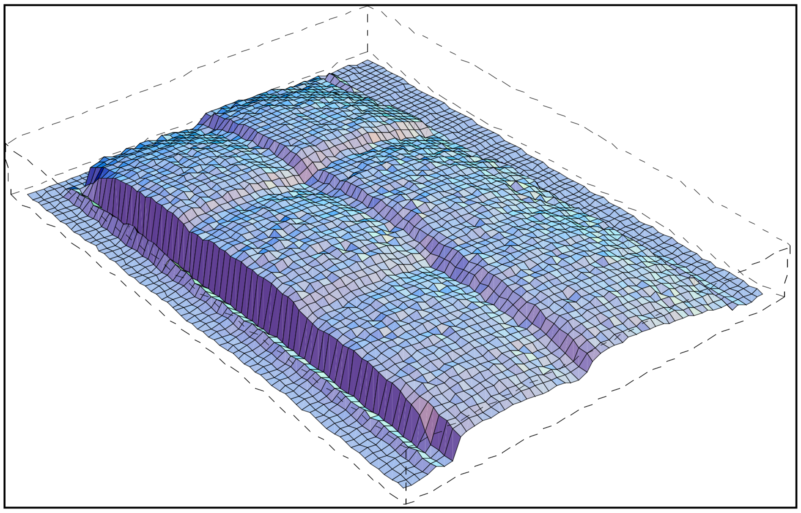

Figure 2: 3D Plot of Imaged Intensity of Maximum White Displayed on the μMural.

Brightness is typically in the 800-1400 lumen range, which is sufficiently bright that a 6+ projector rear-projection array will have no problem being viewed in a fully lighted room. A critical feature for tiling applications is the flatness of the illumination field (brightness falloff from the center of the image to the edges). Ideally the illumination field would have no falloff; in practice, falloff is one of the key image quality factors that must be addressed to create seamless displays. Falloff is clearly evident in Figure 2; however, the difference in intensity between the brightest parts of the screen and the dimmest parts of the screen is a factor of 10 percent, which is very small for commodity projectors.

Color gamut (the range of colors supported by the device) is also important, as is color uniformity or matching of colors between projectors. Depending on the application, color gamut testing might need to be done prior to projector selection. Color calibration and color gamut matching between projectors are important. Low-end projectors often have limited color correction capability and a high-degree of color gamut variability between projectors. As a result, a tiled display can be limited in its ability to achieve matched colors between tiles. Higher-end projectors tend to have more color calibration capability and are more likely to be color matched between projectors.

Data interfaces to the projector are also an important factor in image quality. Analog RGB inputs can suffer from cable noise and pose problems in video formatting when switching between different input sources, often requiring video format adjustments in the projector to preserve framing. Digital video interfaces are just now becoming available that offer the potential to address the noise and video format issues. It is important that the projectors provide a control interface and API that permits setting and changing projector parameters via a serial interface or other computer-based control that can be automated.

Low-end projectors often have relatively simple optics and should be tested for image quality (flatness of focus, sharpness, image distortion, etc.). A zoom lens can enable flexible throw distances but also creates the potential for instability if the zoom is power driven and cannot be manually locked. For multiple-chip projectors, color convergence (alignment of each color component of the pixel) is also important for image quality. Higher-end projectors often support convergence adjustments. Over time, projectors and bulbs age, thereby affecting both color and brightness.

Projectors also vary in image refresh rate and support for stereo. Current low-end projectors typically cannot refresh fast enough to support field sequential stereo (80-100 fps). Often the multiresolution resampling engine that enables variable input display resolution runs more slowly than the image modulator, and this limits stereo capabilities. An effort to identify and prioritize these critical issues for next generation projector designs was started with the first ASCI Advanced Projector Workshop [25]. In the process of constructing the μMural, about 6-8 different projectors were evaluated. Our previous work in this area had led to the conclusion that LCD projectors provided flatter illumination across the entire image but worse pixel component alignment. Also, DLP projectors tended to have darker blacks, but more light scattering from the optics. Argonne's first choice was to use Sharp Notevision 7's; however, as deadlines approached, it became evident Sharp would not be able to deliver in time. Therefore, upon the recommendation of the vendor, we were able to get Epson 7500c's in very short order. The projectors have been extremely good at providing the qualities we require for the μMural, and perhaps are a better choice than the Sharp Notevisions we originally planned to use.

Screen

Most tiled displays are designed to be rear projected. This is not a requirement, but the display geometries are greatly simplified if simple projector throw can be used. Since the μMural is of the simplest design, it has been built using this rear-projection technique. Therefore, our search for screen materials was limited to those that were rear-projection. Various rigid and nonrigid screen materials are available. The primary criteria for choosing screens include image performance (e.g., brightness, resolution, angle of view, contrast ratio), availability of large seamless sheets, type of mounting method, suitability for touch screen applications, rigidity (or degree of self-support in large-span applications), weight, fragility, portability, and cost. As a result of the popularization of rear-projected large-format home entertainment systems, there is an increasingly large market for advanced-projection screen materials [26]. However, these materials often are available in limited sizes and may be optimized for nonoverlapping or single-projector applications (e.g., requiring a Fresnel lens for increased viewing angle support). If the geometry is not planar, then one may need to build the screen in segments or use flexible screen that can be bent or stretched [27].

The screen material in the µMural is JenMar Visual Systems BlackScreen™. It has a resolution of greater than 200 lines per inch and a contrast ratio greater than 250 to 1. The ambient light rejection on the screen is greater than 96%. These factors together make the screen bright enough to be fully usable in a normal-sized room, with the room set for normal ambient conditions. One drawback of this screen material is that it appears not to preserve the polarity of the light it passes. This reduces its utility for doing passive stereo displays. However, since the µMural is not a stereo display device, the lack of stereo functionality was considered irrelevant, and the screen has performed very well since installation. Another drawback of the JenMar screen is the illumination falloff that occurs because of the viewing angle. In Figure 2 and Figure 3, the projector falloff due to the angles between the camera and the screen and the screen and the project is easy to see. The brightest part of each projector is not in the same place with respect to each tile: in the outer tiles, the brightest region is moved slightly toward the center of the screen horizontally.

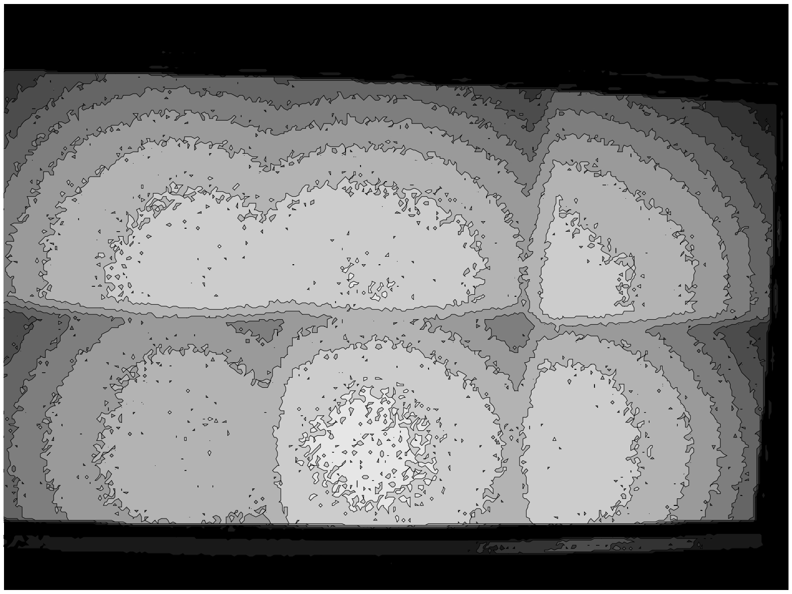

Figure 3: Contour Plot of the µMural Displaying Solid White.

Structural Frame

Each of the components of the tiled display needs to be held firmly in place. The components should not drift, and the structure should be engineered to support the collective weight. The material used should be simple to use to build complex systems. Cutting is reasonable, but materials that need complex connections, such as welding, are too difficult to work with. The projectors weigh in the range of 8 pounds; with the positioners this weight can increase to 20 pounds. There are six of these assemblies for a total of 120 pounds. The screen weighs approximately 50 pounds. All of these components need to be held at a range of heights from 14 to 40 inches from the floor.

Most tiled displays uses Metro™ shelving so are not portable systems. Other portable systems generally do not include multiple projectors. Choices of material include commercial shelving, custom wood solutions, or some system of parts that can be assembled into custom shapes.

The structure of the µMural is made entirely of aluminum extrusion manufactured under the name FrameWorld, by Barrington Automation. The µMural uses the 1.5 inch cross-section variety of FrameWorld components, which provides significant strength under all foreseeable load conditions. This material comes in standard 96 inch lengths but can be ordered in any length shorter than the standard. A variety of connectors and accessories are available for this material. It is a virtual erector set for engineers.

Blending Hardware

One choice in the design of tiled displays is what to do where tiles meet. Video Cube technology offers no options and generally suffers from a very small black edge around each cube, that, when doubled for each cube on an edge, causes the eye to perceive a discontinuity in the large display. Another option afforded by projector-based tiled displays is to overlap the tiles by some amount. This technique costs pixels in the end, but it has been done successfully by using a variety of methods to manage the intensity of the overlapped regions. Overlapping image tiles and tapering the brightness of the image from each projector can result in a smooth intensity transition from tile to tile. This effect can be achieved in signal electronics [28, 29] or in software [20, 21]. As has been pointed out in the literature [22], unvignetted light from the projector results in a brighter-than-black level in and around the projected image. Overlapping or even abutting such images results in a bright region that can be eliminated only by adding baffling to the light path. One approach, referred to as aperture modulation, interposes a baffling window between the projector and screen so as to remove stray light while simultaneously grading the light from one projector so that its overlap with its neighbor results in a continuous and smooth transition.

Residual misalignment, image zoom error, and distortion also drive our design to include blending (either hardware or software). An added benefit of overlapping and blending adjacent projectors is that these problems are somewhat ameliorated by the soft averaging of errors. The μMural incorporates optical blending into the structural assembly. (Princeton incorporates the blending solution into the projector positioners.) For most purposes the distance from the projector to the blending mask hardware is not a critical parameter. Accordingly, the blending hardware in the μMural was designed as a very inexpensive system of adjustable bars based on off-the-shelf lightweight extruded aluminum components. The systems are easy to assemble and easy to adjust. This blending solution costs approximately $100.

Experimentation has been done with alpha masks computed to compensate for falloff from a fixed position. These masks are a software method to achieve the same effect as the optical blending. The one drawback of the software method is that the darkest black is the sum of the brightest overlapped regions in the display. This can lead to a reduced dynamic range if the projectors are particularly bad at producing very dark blacks. Further experiments were done to explore the possibility of combining hardware and software blending. The ability of the blending solution to match illumination intensities in regions where projectors overlap can be seen in Figure 3, where the projectors overlap approximately 8% between the three vertical rows, and approximately 24 percent between the two horizontal rows. The contour plot indicates that the blending hardware is providing a smooth transition between tiles, since there are no hard lines where the contours change. Hard lines would be produced by the sharp decline in illumination where the active area of a projector terminates.

Display Interfaces

The µMural projectors can be driven from a number of sources: three to six Infinite Reality™ graphics pipes, a single eight-channel pipe, a six-headed PC, or a cluster of PCs. It is normally driven by six of the visualization nodes of Chiba City, the large-scale Linux cluster at Argonne National Laboratory. The reason we use multiple rendering engines is to provide different types of functionality through the µMural. One type of functionality, which users are familiar with, is the desktop metaphor. Using a multi-headed single Linux PC or the multichannel option on a single SGI IR, this functionality can be realized on the µMural. Higher-performance methods of using the display are available when each projector is connected to a dedicated rendering engine. This configuration, however, currently cannot support the single logical desktop users are familiar with. Here again lessons can be learned from the history of supercomputing architectural research.

In the parallel computing area much work has been done to explore, enumerate and understand the programming models for the various architectures of supercomputers that have existed [30]. Since hardware architecture has not changed substantially during this era, the set of programming models has not exploded. Parallels can be drawn between the programming models for supercomputers and the programming models for tiled displays, since the computers generating the pixels for the µMural are generally variations on the hardware architectures of the current generation supercomputers.

The goal of the µMural rendering system is to generate pixels at a sufficiently fast rate to give the user a realistic and compelling experience. Current commodity projectors operate natively at 1024x768 at 60Hz. This is a pixel bandwidth maximum of 4.7 megapixels/second per tile of the display. However, current graphics adapters, both commodity and custom, can generate pixel outputs for nongraphics-intensive applications at rates of up to 1920x1280 at 60 Hz. These numbers indicate that the graphics adapters can outperform commodity projectors. The reason this performance is not realized is that applications are more complex that state-of-the-art graphics pipelines. Below I discuss the types of systems I have tested in conjunction with the µMural. The organization by shared-memory and distributed-memory follows the parallel programming literature.

Shared-Memory Rendering Engines

A shared-memory computer is a computer in which multiple processors share the same physical memory. Similarly, a shared-memory rendering engine is a system in which all communication between the processes responsible for generating pixels is done via shared memory. These systems typically have multiple graphics outputs; in most cases these systems try to have one processor per graphics output at the minimum. Examples of these systems include the SGI Origin2000, with multiple Infinite Reality Graphics Systems, and personal computers with multiple CPUs and multiple graphics adapters. The benefits and deficits of each of these systems are discussed below.

PC/Windows

Personal computers running the Microsoft Windows™ operating system are one of the three types of shared-memory rendering engines I will address. My experience with these computers leads me to believe they are the least beneficial of the shared memory systems. The benefits of the Microsoft Windows-based multiple-CPU, multiple graphics-adapter system include excellent device integration, operating system-supported hardware graphics acceleration, general ease-of-use since it is a familiar system, and rapid availability of much hardware and software, since Microsoft Windows and Microsoft Windows Certified Hardware dominate the marketplace. The drawbacks of the Microsoft Windows-based multiple-CPU, multiple-graphics adapter system include no overlap support for multiple graphics adapters, steep learning curve to get optimal performance when graphics adapters have varying performance, and cost of the programming resources needed to harness the highest (since they are geared at the game developers).

PC/Linux

Personal computers running the Linux Operating Systems and Xfree86 version 4.0 or higher [31] are the second of the three types of shared-memory rendering that can be used to generate pixels for tiled displays. I have extensive experience with Linux and am possibly predisposed to favoring Unix-based operating systems over Microsoft Windows. I believe that the Linux solution is by far the least expensive and most flexible of the rendering solutions; however, I think that the shared-memory rendering model is limited in its scalability. The benefits of the Linux-based multiple CPU multiple graphics adapter system include widely available development tools, large communities of developers, and a single logical desktop. The deficits of such a system are typical of open source software solutions: slower driver development; complicated configuration of driver to get the optimal performance; and fragmented, often splintered development efforts.

SGI/Irix

Silicon Graphics Origin series computers with Infinite Reality graphics engines are the final type of shared-memory, multiple graphics adapter system; I have used to generate pixels for the µMural. This system has some definite advantages, including excellent integration of graphics and processing, vendor-supported development tools and graphics libraries, and exotic extensions in the graphics subsystem for doing compute-intensive visualization fast. The disadvantages of the SGI solution include a closed-system model (including the graphics subsystem), extremely expensive hardware and software, and a poor record for tracking the same price/performance curve of commodity graphics systems. One of the few key applications the Infinite Reality graphics systems are still being used for is volume visualization. In order to do this fast, large memory is usually required as an integrated component in the graphics subsystem. Currently, the Infinite Reality systems have the largest amounts of memory of any graphics adapters, with 256 MB and even 512 MB. I believe this trend will continue, with Infinite Reality systems to contain 1 GB and larger memory capacities.

In summary, I believe the lessons learned from the parallel computing community with regard to the shared-memory system model apply directly to the shared-memory model for doing graphics-focused computing. First, scalability is an issue. Current shared-memory systems without non-uniform memory access (NUMA), such as current SGI systems use, stop being cost-effective solutions at 16- or even 8-way parallelism. When NUMA architectures are used, as in the Origin series of computers, systems can be built as large as 512 or 1024 processors; however, the programming complexity required for exacting the highest performance from applications begins to rise sharply [32]. Given these facts, and the fact that processor speeds have been steadily increasing while bus bandwidths have been stalled, I believe that these systems are not currently the best choice for generating pixels for tiled displays.

I also do not believe that a single PC can generate the requisite 142.5 megapixels/second/tile that the μMural can consume operating at the modest speed of 30 frames per second, with 1024x768 pixels per frame.

Distributed-Memory Rendering Engines

Distributed-memory rendering engines are the other type of system I have investigated in conjunction with the μMural for generating pixels. In these systems I have tried to follow the same logic introduced by the Beowulf community in developing Linux clusters to address the price-performance problem of large-scale computing. In this architectural arena there are only two choices, which exist on the same PC hardware. The two choices are Microsoft Windows-based operating systems or Linux. Since Princeton is investigating Microsoft-based clustering for tiled displays, I decided to focus on Linux-based clusters; therefore I describe only this particular type of cluster below.

The μMural is usually driven by six of the thirty-two nodes of Argonne National Laboratory’s Visualization Cluster, which is one “town” of Chiba City, a large Linux system maintained by the Mathematics and Computer Science Division. This cluster is focused on providing the visualization capabilities for the computational nodes, and research is underway to understand how to have applications that involve both computation and visualization at the same time.

The distributed-memory rendering model is highly scalable. However, two specific issues might slow the progress of distributed-memory rendering: the distributed scene graph, more generally known as load-balancing the large scene database, and the slowing of progress on network performance. These issues are not discussed in this paper; however, they pose a serious problem for visualization researchers in the near future.

PC/Linux

The μMural is usually driven by six IBM Intellistation M-PRO (Model 6889-AG5) PC’s, each with a single 500MHz Pentium III CPU, 384 MB of memory, 13 GB 7200RPM IDE primary hard drives, and 18 GB 15000RPM SCSI secondary hard drives. The motherboard chipset is customized by IBM but based on the

Intel 440BX chipset, which has an onboard 10/100Mbit Ethernet, using the Intel 82557 chip. While the ethernet is the primary interface for these nodes, they are each also equipped with a Myricom M2L-PCI64B Myrinet 1280 network interface.

Each node of this cluster has a commodity graphics adapter; currently it contains Elsa Gladiac GeForce2 GTS graphics adapters each with 32 MB of memory. These graphics adapters outperform SGI Infinite Reality graphics systems on a large subset of the SPECglperf benchmarks [33].

Software Overview

The software required to drive the µMural can be broken into three types: system maintenance tools, the Single Logical Desktop, and the Argonne MPI/Glut Framework. The system maintenance tools include programs that can be used to measure and adjust the color balance of the projectors and programs that can measure the alignment of the projectors and create configuration files that represent the configuration. Others who have built tiled displays have attacked these software issues in other ways. Stanford has developed WireGL [7], which provide remote OpenGL rendering; the rendering nodes can be clusters that drive the tiled display. Princeton researchers, who are now adopting WireGL as part of their software infrastructure, previously developed virtual device drivers that presented a single logical Framebuffer to each host of the cluster that drove the tiled display.

Tiled Display Management Tools

The installation and maintenance of a tiled display require software tools that can be routinely used to measure and adjust color balance and position of each tile. Also, once these procedures are complete, measurements can be taken to compute an image warp for each projector to correct for the distortions introduced by the projector optics. With the µMural, experiments have been conducted that show systems involving automated processing and commodity digital cameras can be used to measure and correct the color balance of tiles in a tiled display. Other work in this area is being done at the University of North Carolina at Chapel Hill [34]. Positional measurement of the tiles is a bit harder, and efforts are under way by various groups to solve these problems. Argonne is looking to adapt methods developed at UNC for measurement of projector position to provide input to electronically controlled projector positioners. The work involving "undistorting" images to correct for projector optics has not yet been started, since it is an integration of previous research in the area and the benefit is really the last 2% of the work.

Single Logical Desktop

One of the most obvious uses for tiled displays is to expand the capacity of users who are bound by desktops that lack enough pixels to get their work done. These users need a seamless desktop with more pixels, in order to run numerous applications at the same time. Often these applications include audio and video conferencing in order to collaborate with colleagues on shared work. These users need the system that operates the tiled display to behave identically to the computer that they use daily, with a single keyboard and a single mouse. (It might become beneficial, in the future, to enable multiple keyboards and multiple mice, as these systems grow. It is easy to imagine a desktop with enough pixels to support multiple users operating applications simultaneously, sharing when necessary.)

Currently, the μMural can operate this way using the shared-memory rendering model under Linux. We use the Xinerama X11R4 extension to allow multiple graphics adapters to operate a single logical desktop. The only other systems that we know about that operate this way are the SGI Irix multi-channel option for the Infinite Reality graphics subsystem, and the Hewlett-Packard Single Logical Screen (sls) and Single Logical Screen/Distributed (sls/d) [35] software. I have not been able to use the HP system because it operates only on HP/Unix, which works only on HP workstations, which are not readily available for evaluation.

The SGI multi-channel option for the Infinite Reality graphics system and the Xinerama mode for Xfree86 v4.0 both suffer from problems that limit the usability for fast graphics usage. The SGI solution divides the performance of the IR among the channels being used. On a tiled display, this means only one graphics pipeline can be used for the entire display, providing less performance than can be obtained from multiple graphics pipelines. When multiple graphics pipelines are used on the SGI system, each pipeline operates as if functionally distributed with respect to the graphics. The user cannot drag windows across pipeline boundaries, which is a common operation on a large desktop. This is a very clumsy behavior for users who want to interact with a single desktop.

On the other hand, the Linux solution currently does not have OpenGL hardware acceleration; hence, visualizations do not perform at speeds that could be obtained otherwise. This is a smaller limitation and one hindering only those who are running OpenGL-based applications. For users of the Linux-based single desktop, the first obstacle they usually encounter is the overall system performance. With such a large desktop, it is easy to operate enough applications simultaneously to swamp even the fastest of current processors, or to run out of memory, causing a slowdown while applications are getting swapped out to disk. These limitations point directly at the distributed rendering solution as the only solution to provide enough system performance to keep the display and the human operating at as fast a speed as possible.

OpenGL Applications

Another of the legacy-like interfaces tiled displays have to support is applications that render via the OpenGL API. This contributes directly to the success of these display systems, since currently most of the high-end graphics research is being conducted based on the OpenGL graphics standard. Not supporting applications that use OpenGL would eliminate the utility of these displays for most, if not all, serious researchers in the scientific community. Two areas we have investigated are using the CAVE Library and a new applications framework I developed specifically for tiled displays.

CAVE Library

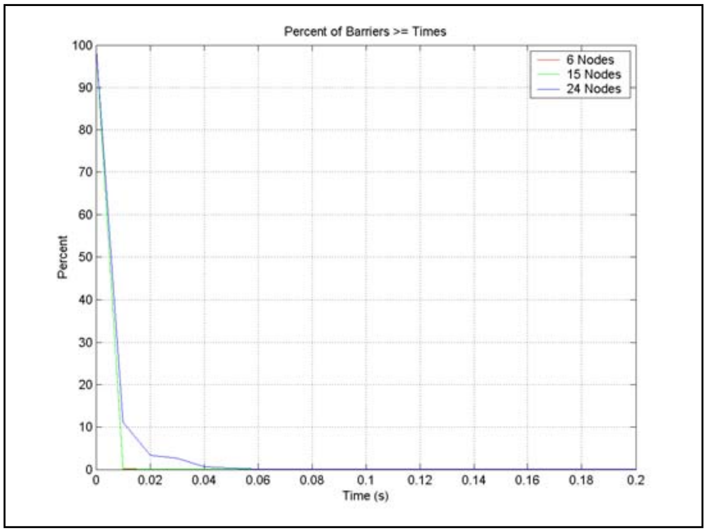

Figure 5: MPI Barrier Tests for various-sized clusters.

The CAVE Library has a well-documented history of use for immersive visualization in CAVEs, ImmersaDesks, and PortaDesks and more recently on a wide variety of other Virtual Reality devices. Since the µMural is a virtual reality device in terms of its field of view, and not in its ability to do stereo video, most of the functionality of the CAVE library is not used. The benefits of supporting the CAVE library on the µMural are the support for legacy applications, the ability of CAVE programmers to write applications for a new device, and the immediate ability to understand screen geometry issues. These advantages should not be considered small. For instance, we can run applications that involve multiple users, where one user interacts on the µMural and the other users are using CAVEs or more traditional display devices. This shows how flexible the CAVE library truly is. The other immediate realization when utilizing the CAVE library is that it contains many programmatic constructs that show an affinity for shared-memory rendering solutions. This is probably due to the relatively short life span of the software and the SGI-centric development cycles that it has undergone.

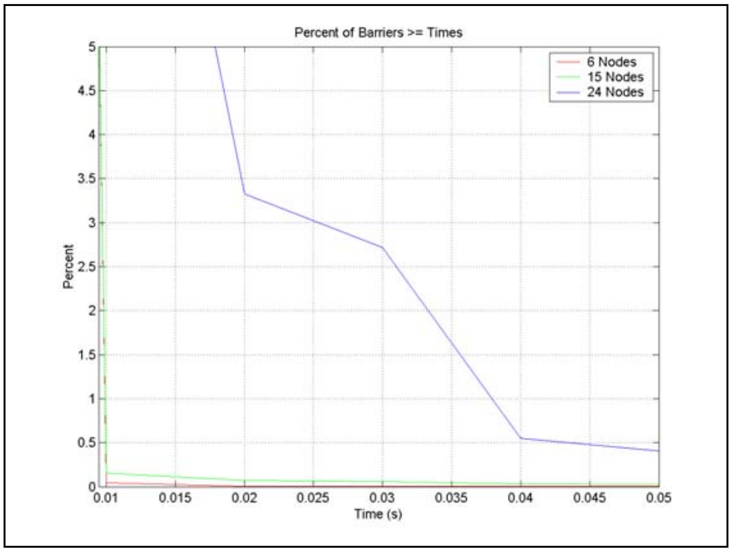

Figure 4: Details of MPI Barrier Tests.

MPI/Glut Application Framework

We have developed a framework for building distributed rendering applications that uses the Message Passing Interface (MPI) [36] and the GL utility (GLUT) library [37]. This framework allows the user to write applications that duplicate the scene database on each node of a cluster of computers and render in a frame-synchronized fashion. The framework has been used to develop test applications, demos, and movie players.

During the development of this framework it was important to understand the characteristics of the MPI Barrier and whether it would be sufficient for frame synchronization on Beowulf type clusters. Some results of tests done to verify the low cost of MPI Barriers to synchronize frames are shown in Figure 5 and Figure 4. These plots show the percentage of the barriers that took as long or longer than the time indicated along the x-axis. The tests executed 1 million barriers on 6, 15, and 24 nodes of a Linux cluster. The test shows that the barriers for clusters of size 6 and 15 are almost identical, while the 24-node cluster performed slightly worse. However, even for the 24-node cluster, only 3.3 percent took longer than 0.02 seconds, and only 2.8 percent took longer than 0.03 seconds. For the 6 and 15 node clusters, less than 0.15 percent of the barriers took longer than 0.01 seconds. From this analysis it is clear that the performance of the MPI Barrier is sufficient to support frame synchronized graphics applications.

Conclusions

I believe the design, fabrication, and development of software for the μMural Tiled display has helped crystallize some of the software architecture issues that currently inhibit larger resolution display devices from being developed and used for large-scale visualization. More work is required to develop an infrastructure that supports the single logical desktop model across a cluster of computers. However, the issues have been illuminated by the creation of a device that can be used in this fashion. A combination of WireGL and some Xfree86 4.0 extensions might solve 80-90 percent of the problem.

While the trend for more pixels continues, obstacles still must be overcome before the performance scales as the resolution does. Scalable rendering solutions are still lacking, and scene graph distribution and load balancing are still very difficult. The research community has not embraced the truly difficult research problems yet, and until that happens, we will continue to be able to build displays with more pixels that can be efficiently rendered.

A large need remains for distributed parallel rendering software. Since the earliest attempts at these systems back in the mid-eighties there has not been much improvement in software. I believe this points to the fact that we know how to do the simple static decompositions of the graphics problems, but we are not yet capable of understanding the more dynamic decomposition problems that are required to get the next order of magnitude of performance from graphics and visualization applications. This dynamic decomposition is the dynamic load balancing, the distributed scene graph, and the space-based decomposition of the rendering problem, which ideally re-decomposes the rendering at every frame.

The μMural display system designed, built and deployed at Argonne National Laboratory has been a success. Astrophysicists use the μMural to view datasets; computer scientists use it to view performance data of large computations; and use it to develop the software needed to efficiently use the display. Since its completion, optimizations have been incorporated into a new version, the μMural2. This version of the tiled display also has six projectors. However, its overall depth has been reduced from 76 inches to 30 inches, making it an office-sized display. The fact that the µMural2 has followed so quickly is a sure sign that tiled displays are one of the possible methods for getting more usable pixels for the user in a cost-effective way.

Acknowledgments

I thank Mark Hereld, Michael E. Papka, Joe Paris, and the rest of the Futures Laboratory, led by Rick Stevens at Argonne National Laboratory. They have supported this work with large amounts of time for discussion, debugging, and physical assembly of the µMural. This project was truly a team effort, and I would be remiss in not thanking the rest of the team for the labor they contributed. I also thank Gail Pieper, in the Mathematics and Computer Science Division, and John Koss, in Central Shops, both at Argonne, for their work involved in making this project a success. This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38.

References

[1] F. X. Timmes, M. Zingale, K. Olson, B. Fryxell, P. Ricker, A. C. Calder, L. J. Dursi, H. Tufo, P. MacNeice, J. W. Truran, and R. Rosner, "On the Cellular Structure of Carbon Detonations," *The Astrophysical Journal*, vol. 10, pp. 938-954, 2000.

[2] T. Sterling, D. Becker, and D. Savarese, "BEOWULF: A Parallel Workstation for Scientific Computation," presented at International Conference on Parallel Processing, 1995.

[3] "Clarity Visual Systems - Video Cubes," Clarity Visual Systems, 2000.

[4] M. Hereld, I. R. Judson, J. Paris, and R. L. Stevens, "Developing Tiled Projection Display Systems," presented at Fourth International Projection Technology Workshop, IPT2000, Ames, Iowa, 2000.

[5] M. Hereld, I. R. Judson, and R. L. Stevens, "Introduction to Building Projection-based Tiled Display Systems," *IEEE Computer Graphics and Applications*, 2000.

[6] K. Li, H. Chen, Y. Chen, D. W. Clark, P. Cook, S. Damianakis, G. Essl, A. Finkelstein, T. Funkhouser, T. Housel, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J. P. Singh, G. Tzanetakis, and J. Zheng, "Early Experiences and Challenges in Building and Using a Scalable Display Wall System," *IEEE Computer Graphics and Applications*, 2000.

[7] G. Humphreys and P. Hanrahan, "A Distributed Graphics System for Large Tiled Displays," presented at IEEE Visualization '99, 1999.

[8] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs, "The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays.," presented at Siggraph '98, 1998.

[9] "Digital Visual Interface," Dell Computers, 2001.

[10] S. Molnar, M. Cox, D. Ellsworth, and H. Fuchs, "A Sorting Classification of Parallel Rendering," *IEEE Computer Graphics and Applications*, pp. 23-31, 1994.

[11] G. Humphreys, I. Buck, M. Eldridge, and P. Hanrahan, "Distributed Rendering for Scalable Displays," presented at Supercomputing, 2000.

[12] H. Igehy, G. Stoll, and P. Hanrahan, "The Design of a Parallel Graphics Interface," presented at SIGGRAPH 98, 1998.

[13] J. A. Friesen and T. D. Tarman, "Remote High-Performance Visualization and Collaboration," *IEEE Computer Graphics and Applications*, 2000.

[14] T. W. Crockett, "Parallel Rendering," in *Encyclopedia of Computer Science and Technology*, vol. 34, A. Kent and J. G. Williams, Eds.: Marcel Dekker, 1996, pp. 335-371.

[15] D. R. Schikore, R. A. Fischer, R. Frank, R. Gaunt, J. Hobson, and B. Whitlock, "High-Resolution Multi-Projector Display Walls and Applications.," *IEEE Computer Graphics and Applications*, 2000.

[16] R. Samanta, J. Zheng, T. Funkhouser, K. Li, and J. P. Singh, "Load Balancing for Multi-Projector Rendering Systems," presented at SIGGRAPH/Eurographics Workshop on Graphics Hardware, 1999.

[17] R. Raskar, M. Cutts, G. Welch, and W. Sturzlinger, "Efficient Image Generation for Multiprojector and Multisurface Displays," presented at 9th Eurographics Workshop on Rendering, 1998.

[18] B. Wei, C. Silva, E. Koutsofios, S. Krishnan, and S. C. North, "Information Visualization with Large Displays," *IEEE Computer Graphics and Applications*, 2000.

[19] "Virtual Network Computer," AT&T UK Research.

[20] R. Surati, "Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays." Cambridge: Massachusetts Institute of Technology, 1999.

[21] R. Raskar, M. S. Brown, R. Yang, W.-C. Chen, G. Welch, and H. Towles, "Multi-Projector Displays Using Camera-Based Registration," presented at Visualization, 1999.

[22] Y. Chen, D. W. Clark, A. Finkelstein, T. Housel, and K. Li, "Methods For Avoiding Seams On High-Resolution Multi-Projector Displays Using An Un-calibrated Camera," Princeton University, Princeton TR-618-00, April 2000.

[23] "High Density Display," San Diego Supercomputer Center, 2000.

[24] "InfoComm International," International Communications Industry Association, Inc., Anaheim, CA June 15-17, 2000 2000.

[25] "ASCI Advanced Projector Workshop," San Diego, CA 2000.
[26] "BlackScreen," JenMar Visual Systems http://www.jenmarvs.com/, 2000.
[27] "Stewart Filmscreen," Stewart Filmscreen Corporation http://www.stewartfilm.com/, 2000.
[28] "Panoram Technologies," http://www.panoramtech.com/.
[29] "Trimension," http://www.trimension-inc.com/.
[30] R. Butler, J. Boyle, T. Disz, B. Glickford, E. Lusk, R. Overbeek, J. Patterson, and R. Stevens, *Portable Programs for Parallel Processors*. New York, NY: Holt, Rinehart and Winston, Inc., 1987.
[31] "XFree86," http://www.xfree86.org/ 2000.
[32] U. Ramachandran, G. Shah, A. Sivasubramaniam, A. Singla, and I. Yanasak, "Architectural Mechanisms for Explicit Communication in Shared Memory Multiprocessors," presented at Supercomputing, San Diego, CA, 1995.
[33] "SPECglperf 3.1.2," http://www.spec.org/gpc/opc.static/glperf.htm, 2000.
[34] A. Majumder, Z. He, H. Towles, and G. Welch, "Color Matching of Projectors for Multi-Projector Displays," presented at Eurographics 2000, 2000.
[35] J. Walls, "Understanding X Features: Distributed Single Logical Screen," Hewlett-Packard September 23 1998.
[36] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*: MIT Press, 1994.
[37] M. Kilgard, *OpenGL Programming for the X Window System*, 3rd ed. Reading, MA: Addison-Wesley Publishing Company, 1997.

Appendixes

Appendix A: Parts Lists

Table 5: Parts from AJ Hudson (some customized by Argonne Shops).

Qty	Part	Description for Ordering	Drawing Number	Functional Description	Screen	Screen Arms	Frame 1	Frame 2	Projector Cage	Base	Lift	Skin
2	EB-7-18	1.5x1.5 Extrusion Corner (18")		Lift Post Brace							2	
4	EBP-5	3" & 4" Caster Plate		Caster Plates					4			
4	ECC-A	3-Way Castle Corner		Screen Corners	4							
3	EB-46	Adjustable Corner Post		Screen /Screen Arm Attachment		2						
6	EB-48	Adjustable Corner Post		Screen /Screen Arm Attachment	2	4						
4	EAL-3	Adjustable Leg		Levelling foot						4		
12	EB-1	Angle Bracket		Screen Arm Connectors/Shipping Connectors		12						
2	EB-18	Angle Bracket		Top Horizontal Shadow Mask Connector			2					
2	EB-2	Angle Bracket		Frame to Cage Connector				2				
20	EB-3	Angle Bracket		Shadow Mask to Frame Connectors		2	6	12				
6	EX-18-41	Angle Bracket Profile		Shadow Mask				6				
4	EX-18-72.5	Angle Bracket Profile		Shadow Mask			4					
36	EB-11	Corner Gusset		Corner Connectors			4		28	4		
4	EB-30	Connector Angle		Projector Lift Connectors		4						
4	EB-55	Connector Angle		Projector Cage/Base Connector					4			
1	EB-57	Connector Angle (Customized)	18	Screen to Screen Arm Connector								
1	EB-58	Connector Angle (Customized)	19	Screen to Screen Arm Connector								
16	EB-22	Connector Plate		Projector Mount Plate					16			
6	EB-23	Connector Plate		Projector Mount Plate					6			
4	EB-24	Connector Plate		Base Spanner Connector						4		
4	EB-28	Connector Plate		Frame1 to Cage Connector			2					
4	EB-29	Connector Plate		Screen Arm to Cage Connector					4			
2	EB-33	Connector Plate		Frame1 to Cage Connector				2				
50	EF-2D	Double Econo T-nut		Nuts, Oh my!								
500	EF-2	Economy T-nut		Nuts, Oh my!								
4	EBP-1	Leveling Plate		Foot Plate						4		
4	EBP-6	Leveling Plate		Caster Plate					4			
70	EF-19-A	Plastic End Plates		Finish Plates								
20	EF-19-B	Plastic End Plates		Finish Plates								
10	EF-19-C	Plastic End Plates		Finish Plates								
2	EF-15B	Plastic Handle		Lift Handles							2	
100	EF-14	Push Lock Fasteners		Finish Plate Fasteners								
4	EB-93	Single Mount Linear Bearings		Linear Lift Slide							4	
2	EX-2-30	Structural Aluminum Extrusion		Castor Support Bar					2			
2	EX-2-43	Structural Aluminum Extrusion		Lift Post							2	
2	EX-2-77.5	Structural Aluminum Extrusion		Base Spanner						2		
2	EX-3-30	Structural Aluminum Extrusion		Base Leg						2		
4	EX-6-17	Structural Aluminum Extrusion		Projector Cage (Front to Back, Internal)					4			
4	EX-6-30	Structural Aluminum Extrusion		Projector Cage (Front to Back, External)					4			
2	EX-6-30	Structural Aluminum Extrusion	13	Projector Cage (Bottom, Filler bars)					2			
10	EX-6-41.5	Structural Aluminum Extrusion		Projector Cage (Top to Bottom)					10			
4	EX-6-72	Structural Aluminum Extrusion		Projector Lift (Side to Side)					4			
10	EX-6-76	Structural Aluminum Extrusion		Projector Cage (Side to Side)		2			8			
2	EX-7-41.5	Structural Aluminum Extrusion		Vertical Member – Beveled Screen Frame	2							
2	EX-7-73	Structural Aluminum Extrusion		Horizontal Member – Beveled Screen Frame	2							

Table 6: Parts from McMaster & Carr.

Qty	Part (from catalog 99)	Description	Screen	Screen Arms	Frame 1	Frame 2	Projector Cage	Base	Lift	Skin
2	3644T35	Standard Duty Single Speed Hand Winch							2	
2	3308T11	Cable Assembly for Hand Winch							2	
12	1909A52	Adhesive Backed white inch/mm ruler					4			
5	8859K84	0.032" x 0.75" x 12" Brass Strips (6 in package)			3	8				
4	?	Casters						4		

Table 7: Parts Made by Argonne Central Shops.

Qty	Description for Ordering	Drawing Number	Functional Description	Screen	Screen Arms	Frame 1	Frame 2	Projector Cage	Base	Lift	Skin
1	30"x72"x1/16" Aluminum Sheet		Top Cover								1
1	30"x64"x1/8" Aluminum Sheet		Bottom Cover								1
4	4"x3"x1/2" Aluminum Block	17	Linear Lift Bearing Clamp							4	
4	Custom Aluminum Blocks	14	Lift Block							2	
2	1"x1/4"x41.5" Aluminum Bar	15	Vertical Screen Edge Clamp	2							
2	1"x1/4"x73" Aluminum Bar	16	Horizontal Screen Edge Clamp	2							

Appendix B: Drawings of Custom Parts

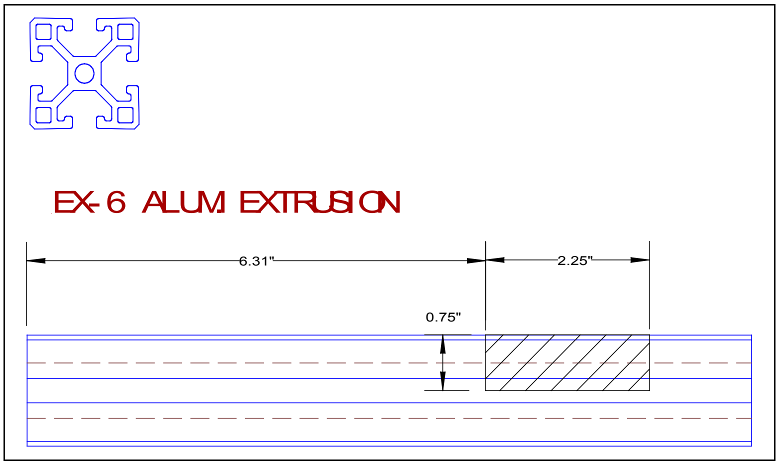

Figure 6: Projector Cage Base Gap Plugger – Modified extrusion.

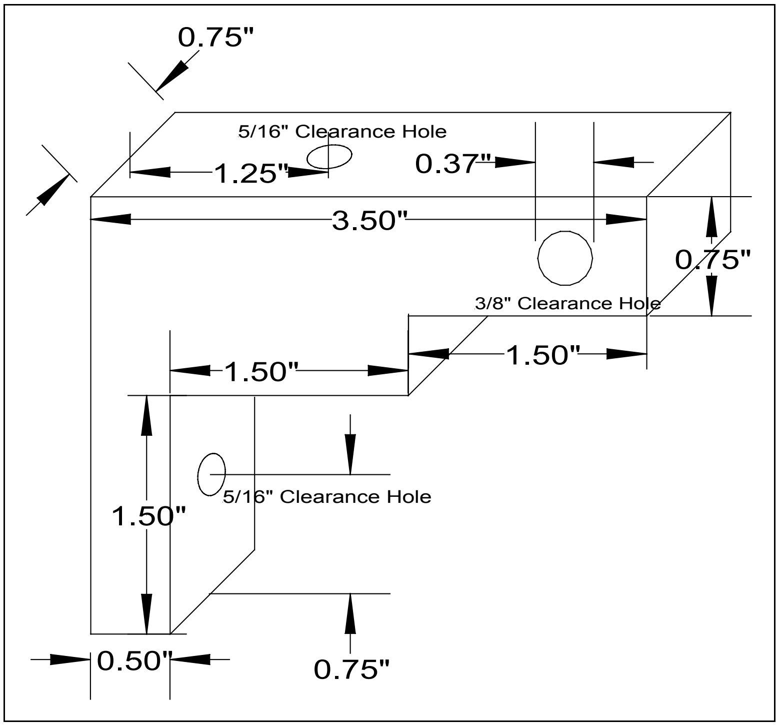

Figure 7: Custom Lift Block – 2 pieces (left and right).

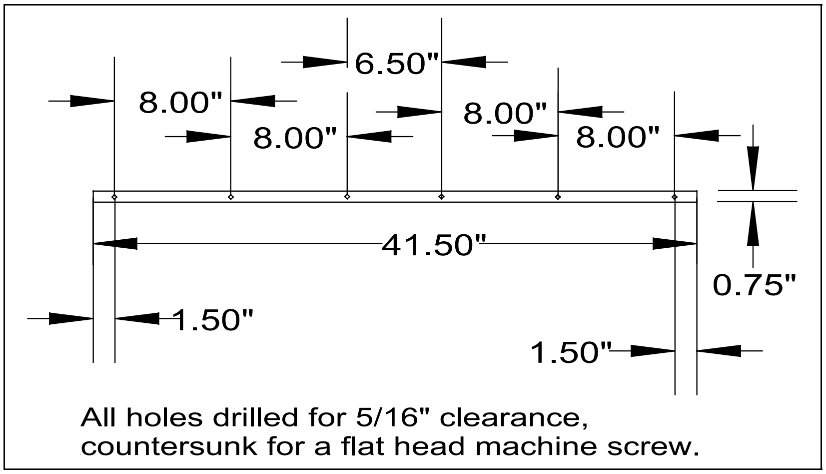

All holes drilled for 5/16" clearance,
countersunk for a flat head machine screw.

Figure 8: Vertical Screen Edge Clamp – 2 pieces (left and right).

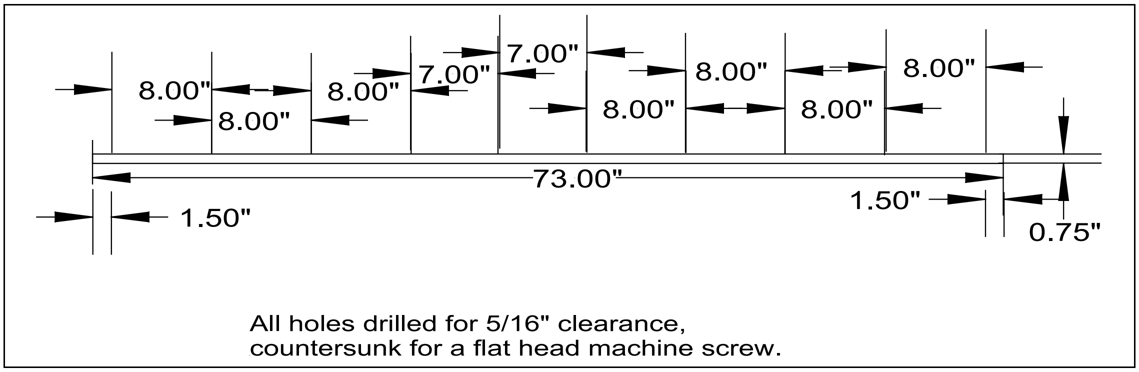

All holes drilled for 5/16" clearance,
countersunk for a flat head machine screw.

Figure 9: Horizontal Screen Edge Clamp – 2 pieces (top and bottom).

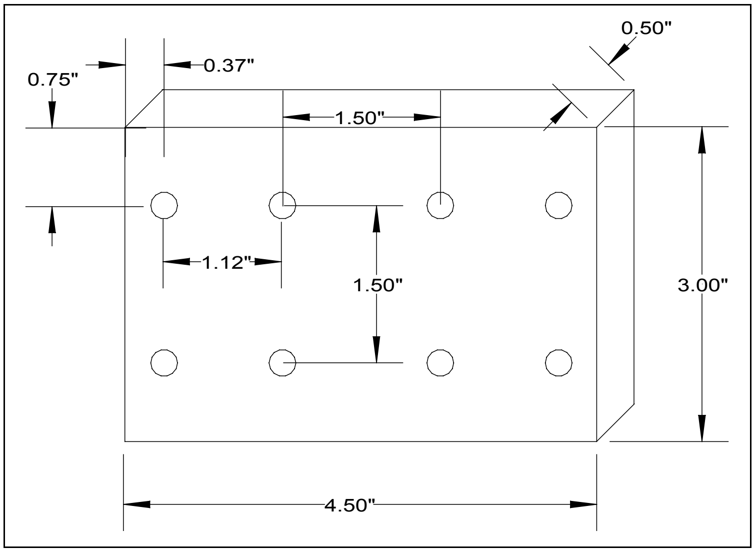

Figure 10: Linear Lift Bearing Clamp – 4 pieces (2 on each lift post).

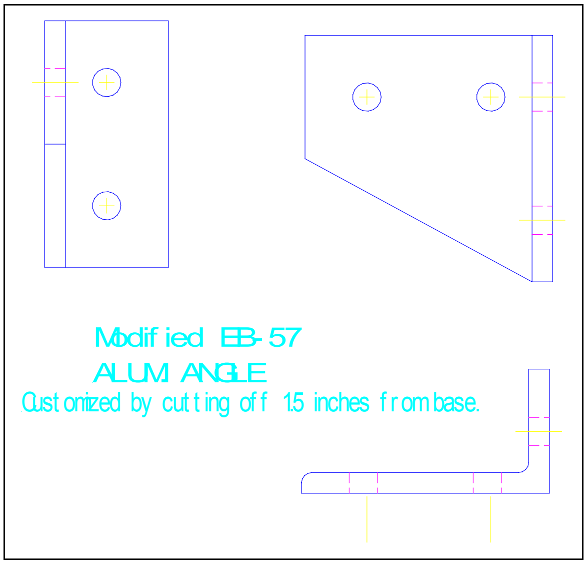

Figure 11: Modified EB-57 holds the screen on the Right Screen Arm.

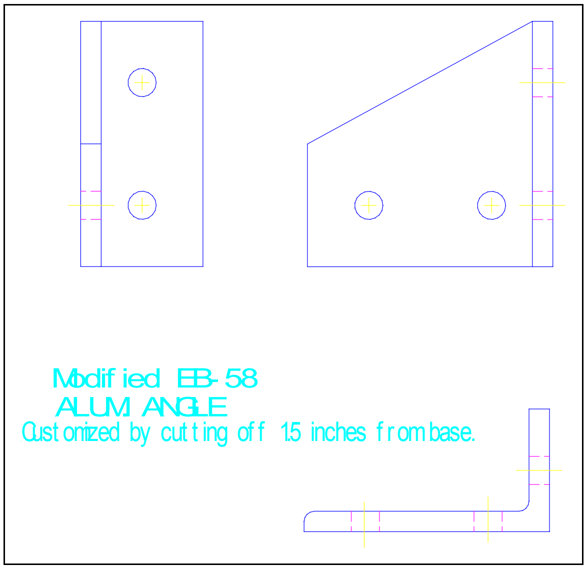

Figure 12: Modified EB-57 holds the screen on the left screen arm.

Appendix C: Positioner Blueprints

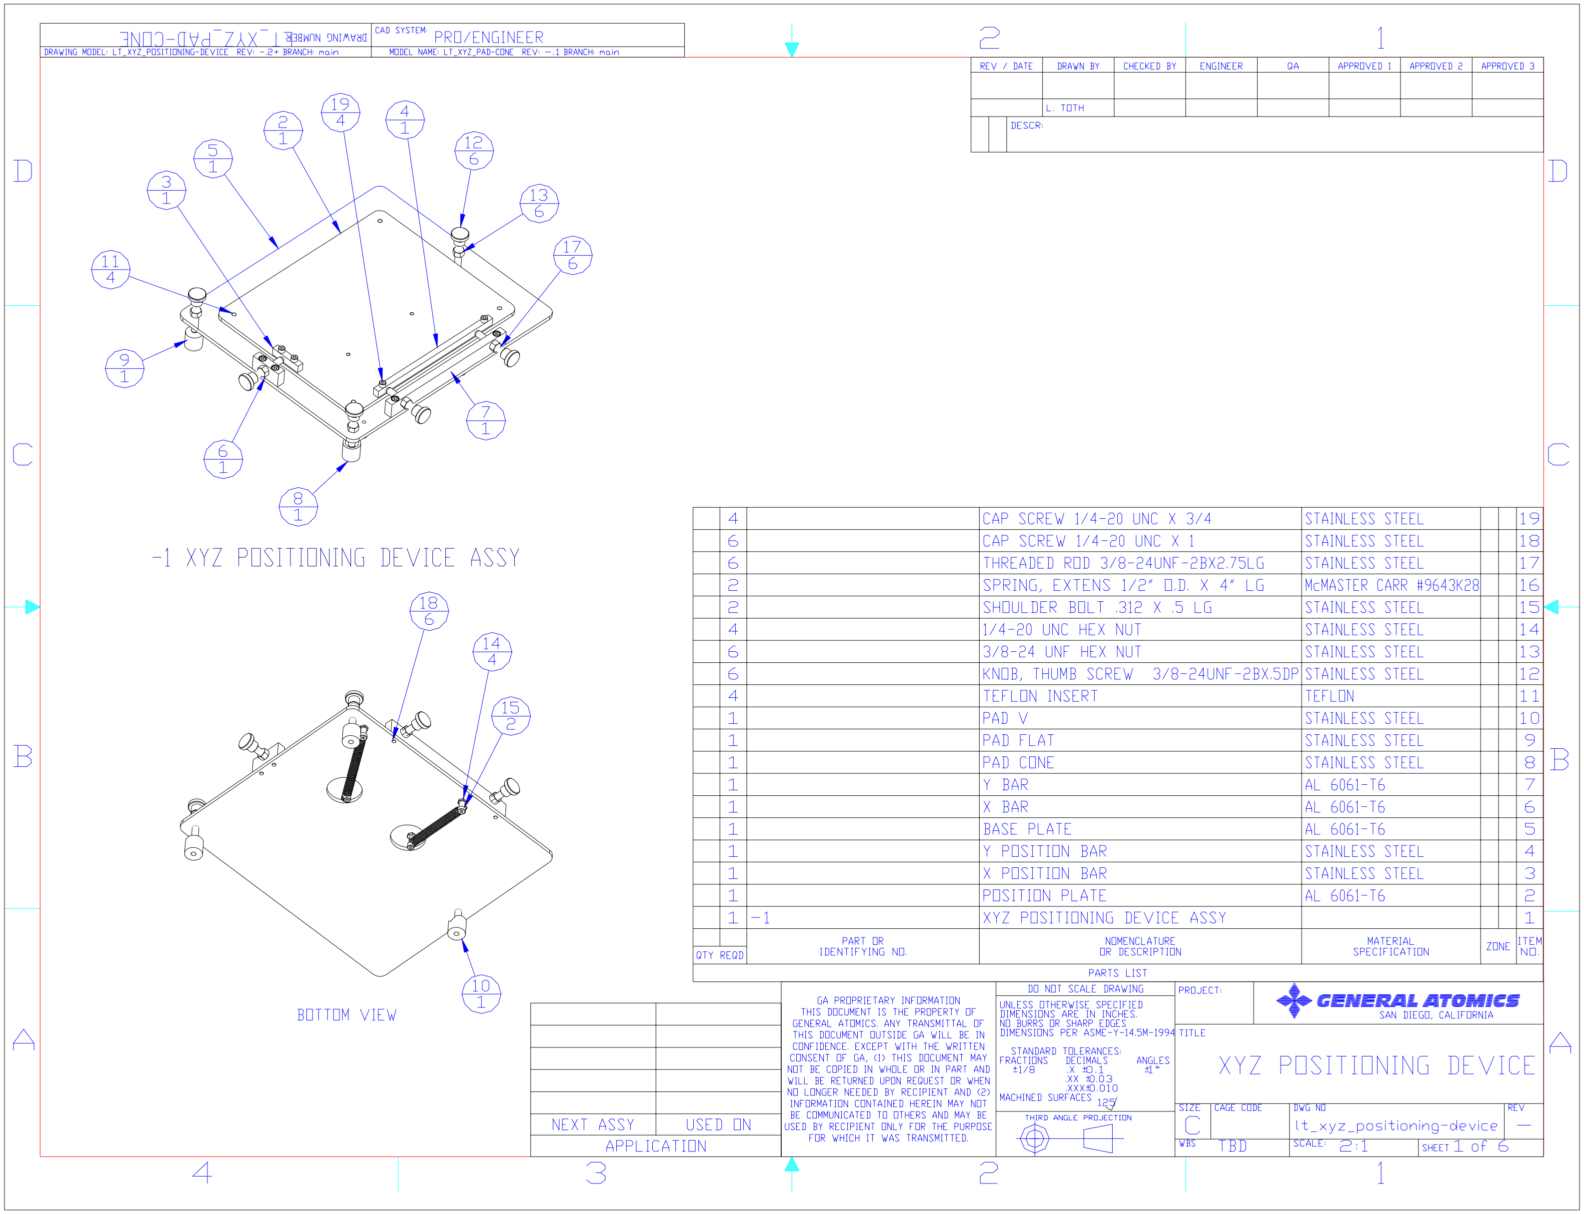

QTY REQD	PART OR IDENTIFYING NO.	NOMENCLATURE OR DESCRIPTION	MATERIAL SPECIFICATION	ZONE	ITEM NO.
4		CAP SCREW 1/4-20 UNC X 3/4	STAINLESS STEEL		19
6		CAP SCREW 1/4-20 UNC X 1	STAINLESS STEEL		18
6		THREADED ROD 3/8-24UNF-2BX2.75LG	STAINLESS STEEL		17
2		SPRING, EXTENS 1/2" O.D. X 4" LG	McMASTER CARR #9643K28		16
2		SHOULDER BOLT .312 X .5 LG	STAINLESS STEEL		15
4		1/4-20 UNC HEX NUT	STAINLESS STEEL		14
6		3/8-24 UNF HEX NUT	STAINLESS STEEL		13
6		KNOB, THUMB SCREW 3/8-24UNF-2BX.5DP	STAINLESS STEEL		12
4		TEFLON INSERT	TEFLON		11
1		PAD V	STAINLESS STEEL		10
1		PAD FLAT	STAINLESS STEEL		9
1		PAD CONE	STAINLESS STEEL		8
1		Y BAR	AL 6061-T6		7
1		X BAR	AL 6061-T6		6
1		BASE PLATE	AL 6061-T6		5
1		Y POSITION BAR	STAINLESS STEEL		4
1		X POSITION BAR	STAINLESS STEEL		3
1		POSITION PLATE	AL 6061-T6		2
1	-1	XYZ POSITIONING DEVICE ASSY			1

PARTS LIST

GA PROPRIETARY INFORMATION
THIS DOCUMENT IS THE PROPERTY OF GENERAL ATOMICS. ANY TRANSMITTAL OF THIS DOCUMENT OUTSIDE GA WILL BE IN CONFIDENCE. EXCEPT WITH THE WRITTEN CONSENT OF GA, (1) THIS DOCUMENT MAY NOT BE COPIED IN WHOLE OR IN PART AND WILL BE RETURNED UPON REQUEST OR WHEN NO LONGER NEEDED BY RECIPIENT AND (2) INFORMATION CONTAINED HEREIN MAY NOT BE COMMUNICATED TO OTHERS AND MAY BE USED BY RECIPIENT ONLY FOR THE PURPOSE FOR WHICH IT WAS TRANSMITTED.

UNLESS OTHERWISE SPECIFIED
DIMENSIONS ARE IN INCHES.
NO BURRS OR SHARP EDGES
DIMENSIONS PER ASME-Y-14.5M-1994

STANDARD TOLERANCES:
FRACTIONS ±1/8
DECIMALS .X ±0.1 .XX ±0.03 .XXX±0.010
ANGLES ±1°
MACHINED SURFACES 125

Figure 13: Assembly Drawing of Argonne Projector Positioner.

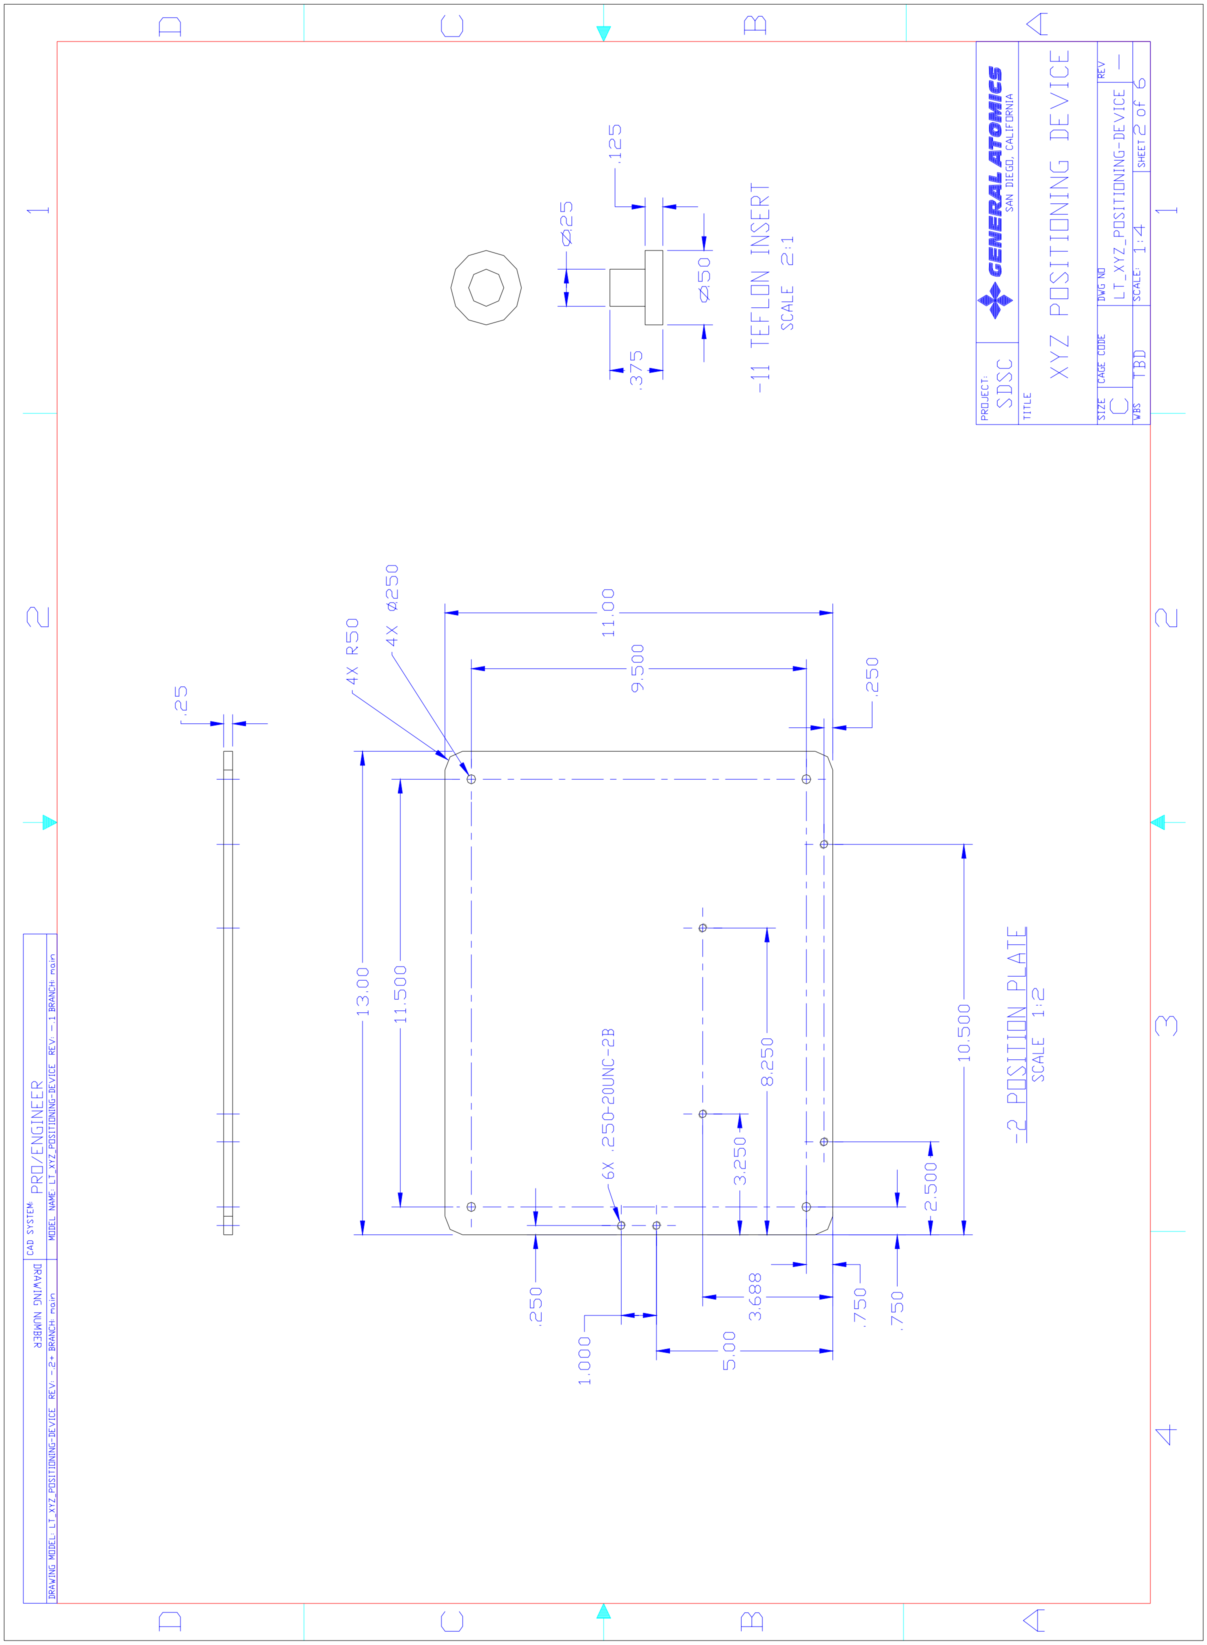

Figure 14: Machine Drawing of Positioner Plate.

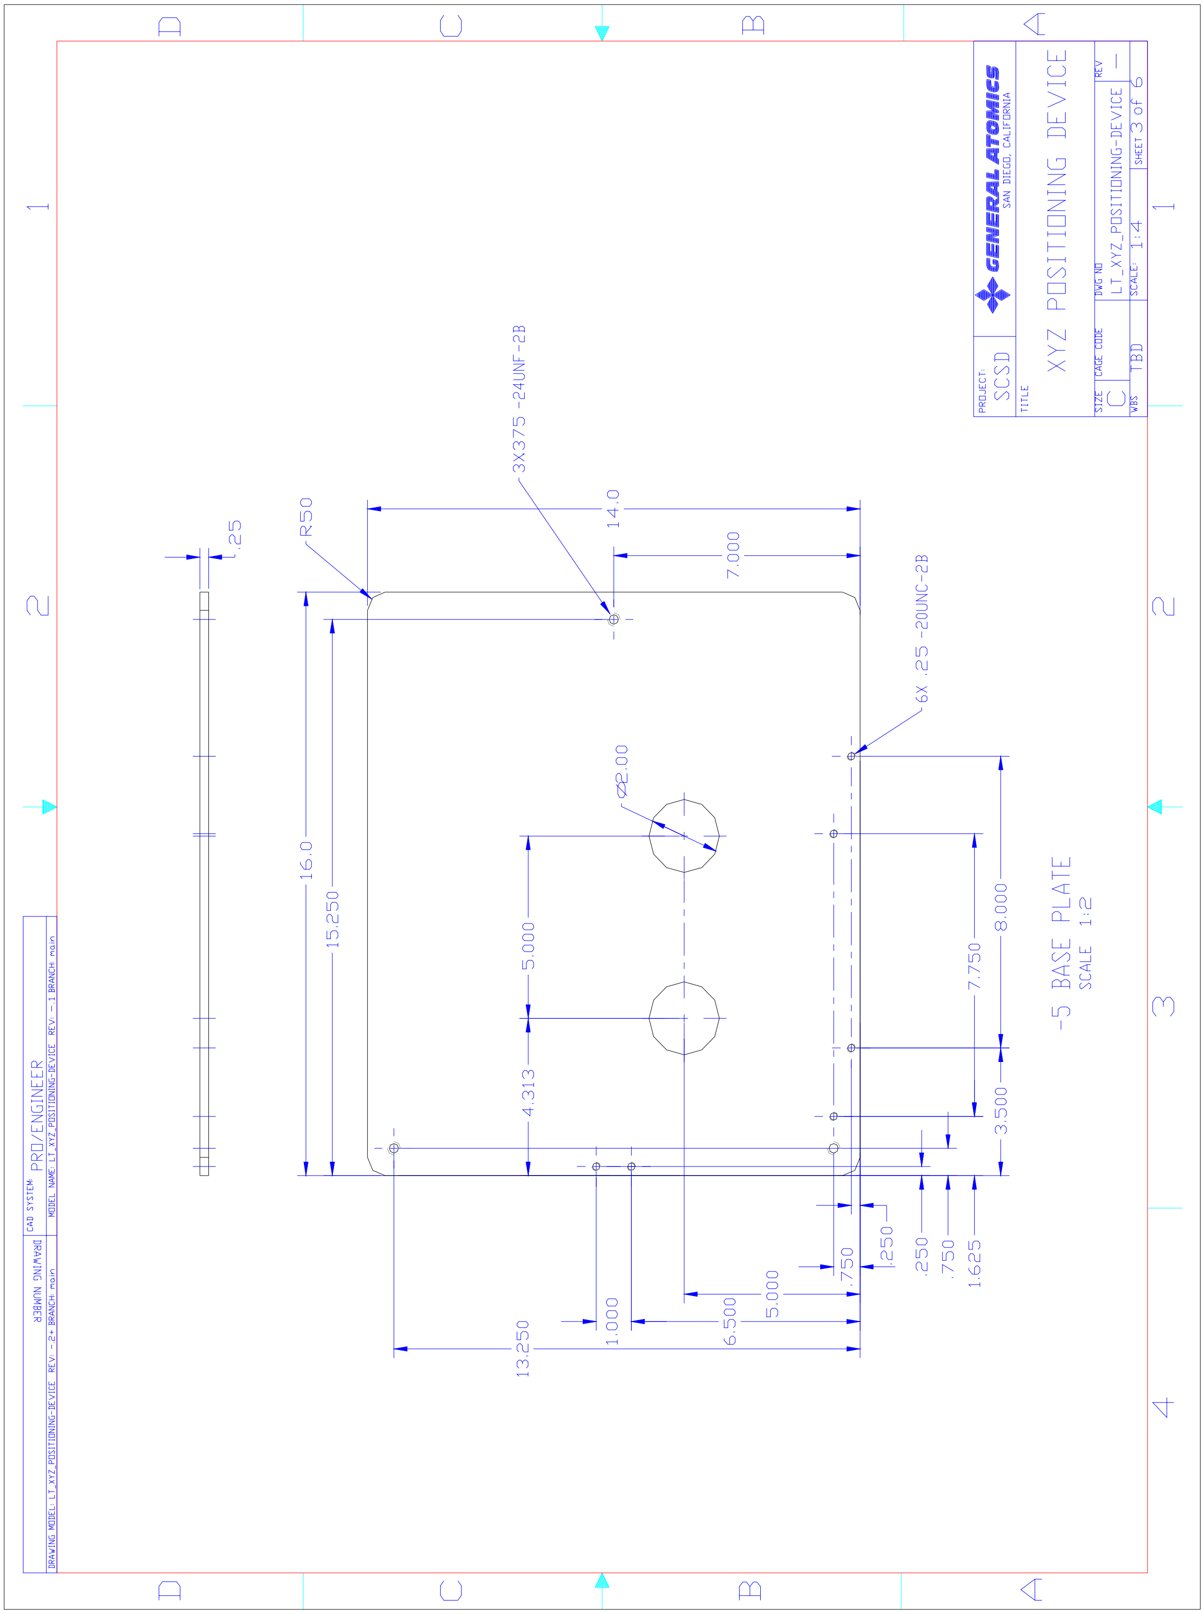

Figure 15: Machine Drawing of Base Plate.

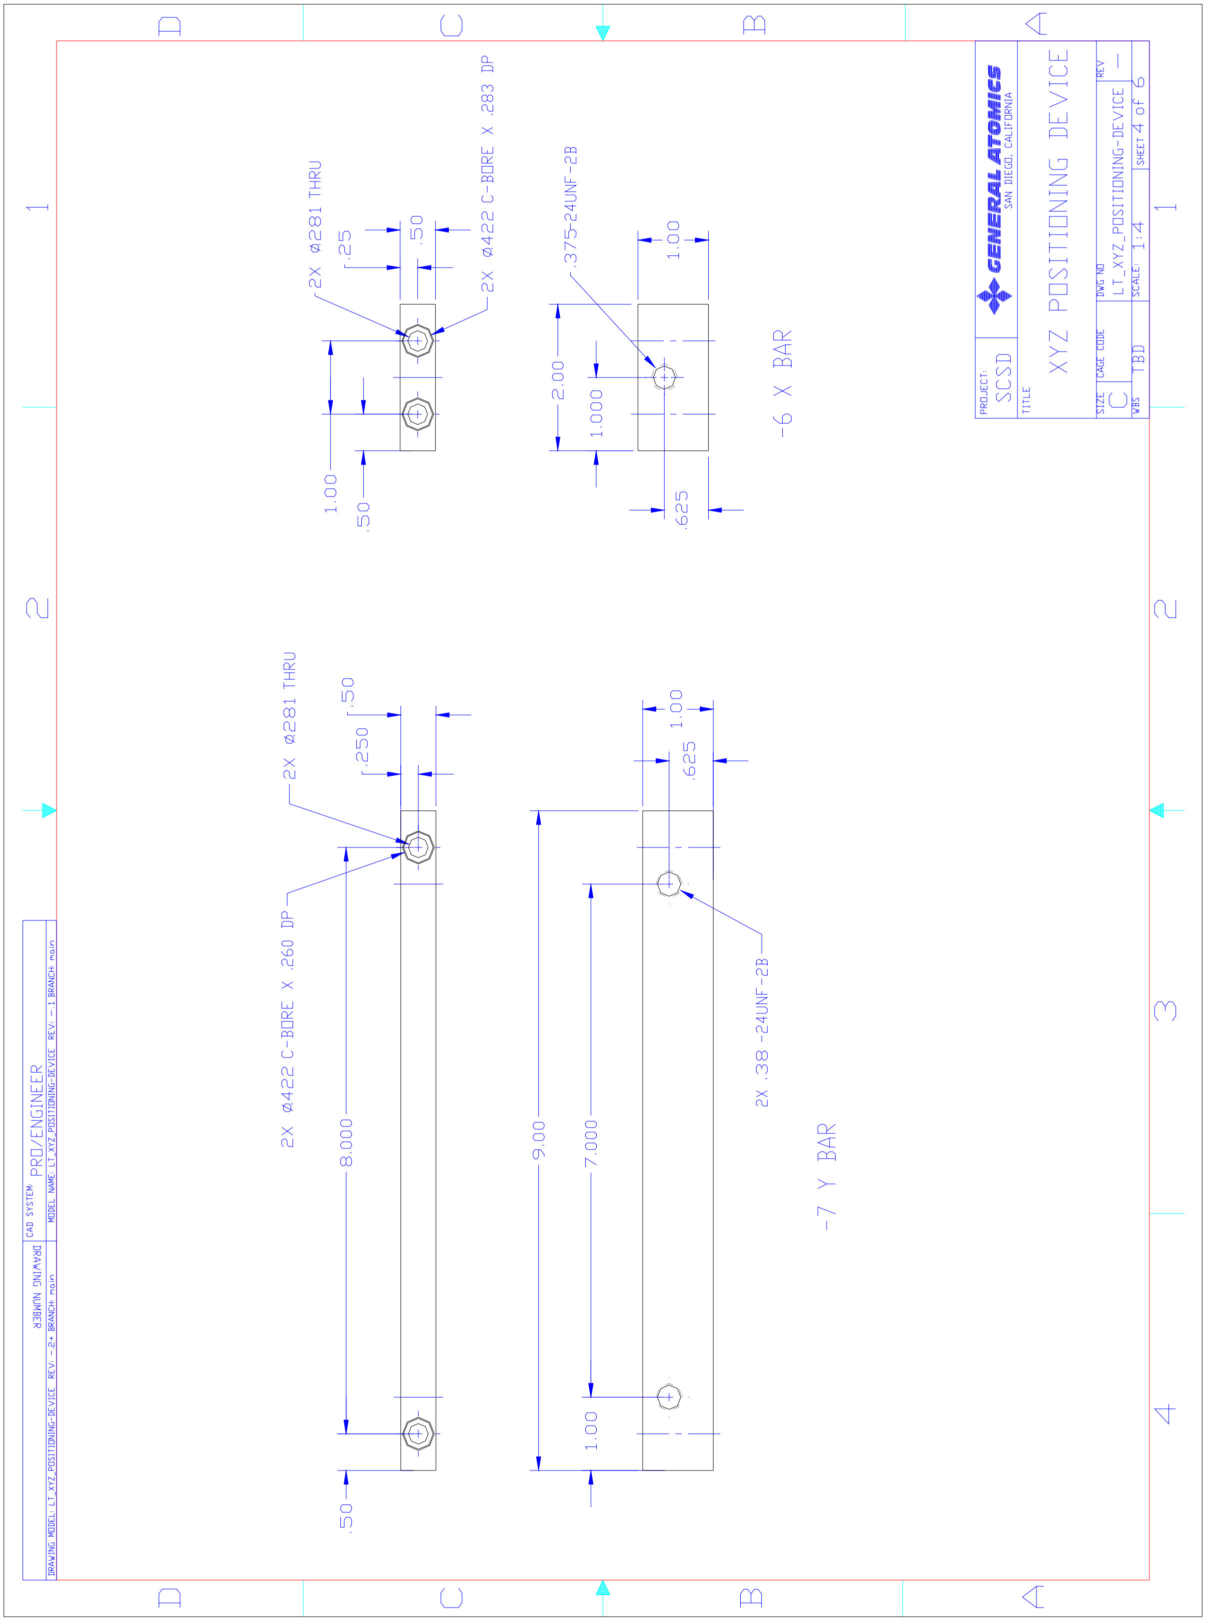

Figure 16: Machine Drawings of X and Y Bars.

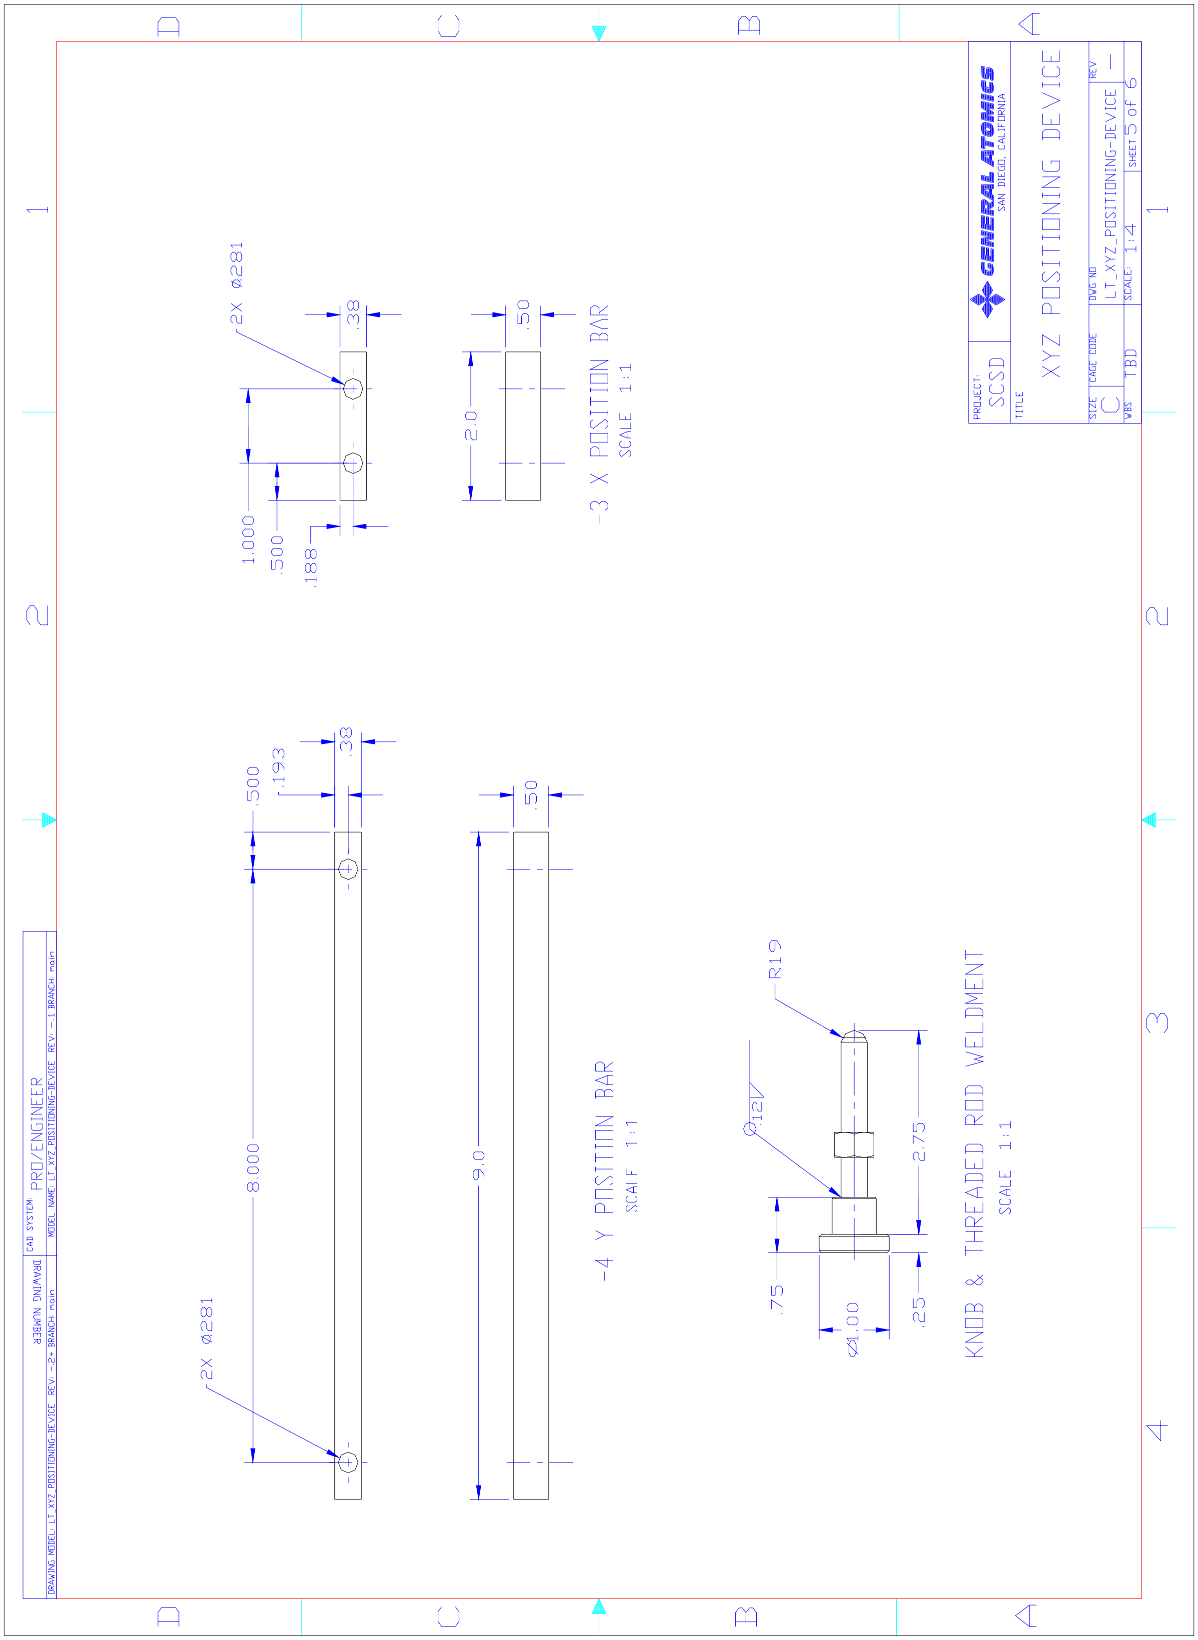

Figure 17: Machine Drawings of X and Y Position Bars and Threaded Ball-End Screw Assembly.

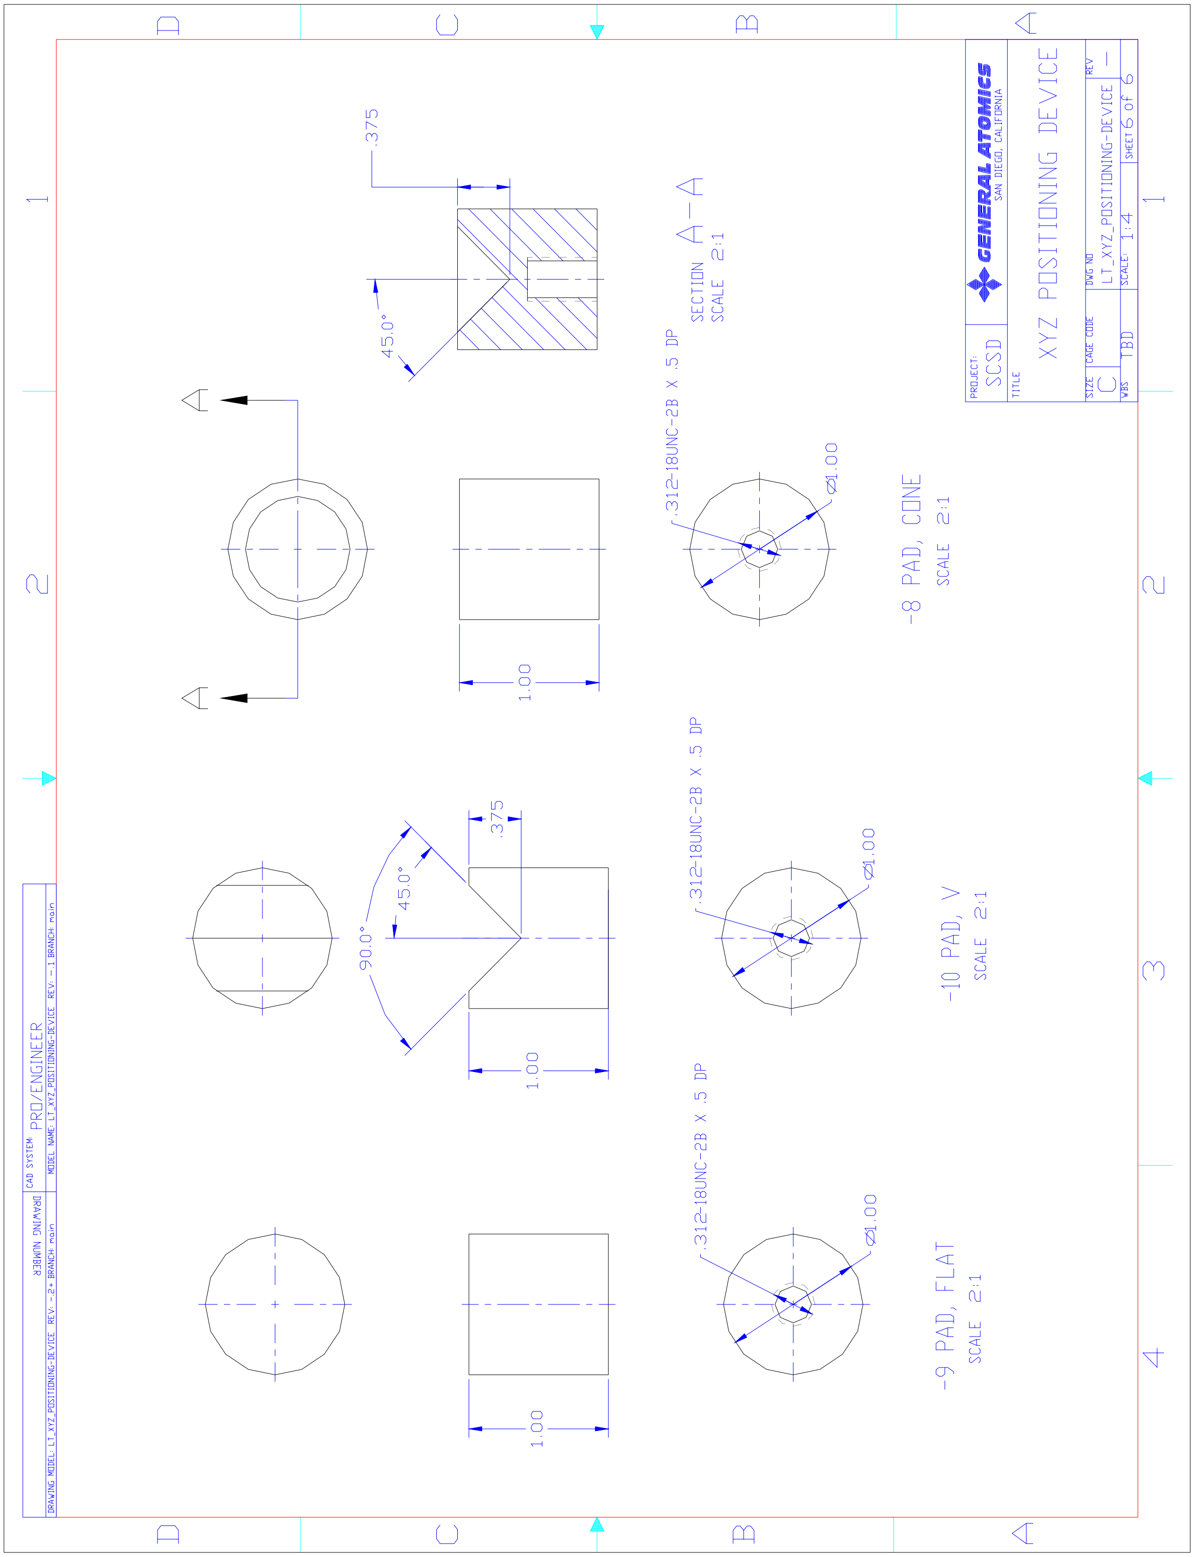

Figure 18: Machine Drawings of Kinematic Landing Pads.

Appendix D: Vendor Contact Information

Frame World Aluminum Structural Framing Components

Barrington Automation
780 Tek Drive
Crystal Lake, IL 60014
815 477 1400
815 477 9818 FAX
http://www.frame-world.com/
info@frame-world.com

Distributed by:

A J Hudson Co.
907 W. Liberty Drive
Wheaton, IL 60187-4846
630 665 6920
630 665 4840 FAX
http://www.ajhudson.com/

Screen Material

JenMar Visual Systems
5349 Randall Place
Fremont, CA 94538
510 249 2060
510 249 2193
http://www.jenmarvs.com/

Miscellaneous Parts

McMaster-Carr Supply Company
P.O. Box 4355
Chicago, IL 60680-4355
708 833 0300
708 834 9427 FAX
http://www.mcmaster.com/

Projectors

CompView
10035 SW Arctic Drive
Beaverton, OR 97005
800 448 8439
503 626 8439 FAX
http://www.compview.com/