Evolution of the Varrier™ Autostereoscopic VR Display: 2001-2007

Tom Peterka*, Robert L. Kooima, Javier I. Girado, Jinghua Ge, Daniel J. Sandin, Thomas A. DeFanti

Electronic Visualization Laboratory
University of Illinois at Chicago, Chicago IL 60607

ABSTRACT

Autostereoscopy (AS) is an increasingly valuable virtual reality (VR) display technology; indeed, the IS&T / SPIE Electronic Imaging Conference has seen rapid growth in the number and scope of AS papers in recent years. The first Varrier paper appeared at SPIE in 2001, and much has changed since then. What began as a single-panel prototype has grown to a full scale VR autostereo display system, with a variety of form factors, features, and options. Varrier is a barrier strip AS display system that qualifies as a true VR display, offering a head-tracked ortho-stereo first person interactive VR experience without the need for glasses or other gear to be worn by the user.

Since Varrier's inception, new algorithmic and systemic developments have produced performance and quality improvements. Visual acuity has increased by a factor of 1.4X with new fine-resolution barrier strip linescreens and computational algorithms that support variable sub-pixel resolutions. Performance has improved by a factor of 3X using a new GPU shader-based sub-pixel algorithm that accomplishes in one pass what previously required three passes. The Varrier modulation algorithm that began as a computationally expensive task is now no more costly than conventional stereoscopic rendering. Interactive rendering rates of 60 Hz are now possible in Varrier for complex scene geometry on the order of 100K vertices, and performance is GPU bound, hence it is expected to continue improving with graphics card enhancements.

Head tracking is accomplished with a neural network camera-based tracking system developed at EVL for Varrier. Multiple cameras capture subjects at 120 Hz and the neural network recognizes known faces from a database and tracks them in 3D space. New faces are trained and added to the database in a matter of minutes, and accuracy is comparable to commercially available tracking systems.

Varrier supports a variety of VR applications, including visualization of polygonal, ray traced, and volume rendered data. Both AS movie playback of pre-rendered stereo frames and interactive manipulation of 3D models are supported. Local as well as distributed computation is employed in various applications. Long-distance collaboration has been demonstrated with AS teleconferencing in Varrier. A variety of application domains such as art, medicine, and science have been exhibited, and Varrier exists in a variety of form factors from large tiled installations to smaller desktop forms to fit a variety of space and budget constraints.

Newest developments include the use of a dynamic parallax barrier that affords features that were inconceivable with a static barrier.

Keywords: Varrier, autostereoscopic 3D display, virtual reality, camera-based tracking, parallax barrier

* address correspondence to: tpeterka@evl.uic.edu

1. INTRODUCTION

The annual IS&T / SPIE Electronic Imaging Conference experienced rapid growth in the number and scope of autostereo (AS) papers in the last six years. The first Varrier paper to appear at SPIE was in 2001 [2], and what began as a small single-panel prototype system has matured to a full-scale VR AS display system, available in a variety of form factors with multiple features and options. The growth process resulted in notable paper publications and high-profile conference demonstrations, but more importantly, in solutions to key science and engineering problems. These range from breakthroughs in performance and quality to more routine engineering tasks of construction, configuration, and calibration.

This paper summarizes the last six years of progress, following a problem-solution timeline of Varrier. It is a progression that continues today; new problems arise and leading edge technologies are being tested. For readers unfamiliar with autostereoscopy, Section 2 presents background material on the optics of barrier strip AS, the computational algorithms used to render spatially multiplexed images, and a brief coverage of tracker technology. The background is necessarily brief, although readers are referred to specific sources for more detailed coverage. Section 3 presents a timeline of Varrier progression, from the first prototype in 2001 to current research being conducted at EVL. Each of the Varrier form factors produced thus far is described in chronological order. Section 4 summarizes some of main problems solved along the way, with a table that quantifies the various form factors along several orthogonal dimensions. Finally, future research is outlined in Section 5.

2. BACKGROUND

2.1 Barrier strip AS

A static parallax barrier display is constructed by printing a grid of opaque and transparent strips onto photographic film, laminating this film to a pane of glass, and attaching the assembly to the front of an LCD display. A parallax barrier functions equivalently to a lenticular screen, as shown in Figure 1. By rendering appropriate strips of left and right eye perspective images onto the LCD display, two disparate images are viewed simultaneously, one by each eye without the need for glasses. Tilting the barrier at an angle from vertical mitigates the Moiré interference pattern that results between the barrier and pixel grid [9], [10]. Other AS methods besides the parallax barrier and lenticular screen are surveyed in [11], [12].

One of two design paradigms is commonly employed in parallax barrier AS: tracked two view and untracked multi-view. Tracked systems [13], [14], [15] direct a stereo pair of views to the eye positions, and follow those positions from frame to frame. Until now, they have been limited to a single user. The other popular design pattern, the untracked multi-view panoramagram, is a set of perspective views rendered from slightly varying centers of projection. With this approach, limited look-around capability is simulated for any number of untracked viewers [16]. A comparison of tracked and untracked AS appears in [3].

2.2 Computational Algorithms

AS algorithms fall into three categories: sorting, occlusion, and combining. Original barrier strip AS was performed by sorting, or electronically cutting the perspective views into columns and forming a composite image from the interleaved columns of the original images. The phscologram [17] and the interzig™ [16] algorithm of the Synthagram™ are early examples of sorting algorithms. Sorting algorithms by definition operate in the integer domain on image pixels.

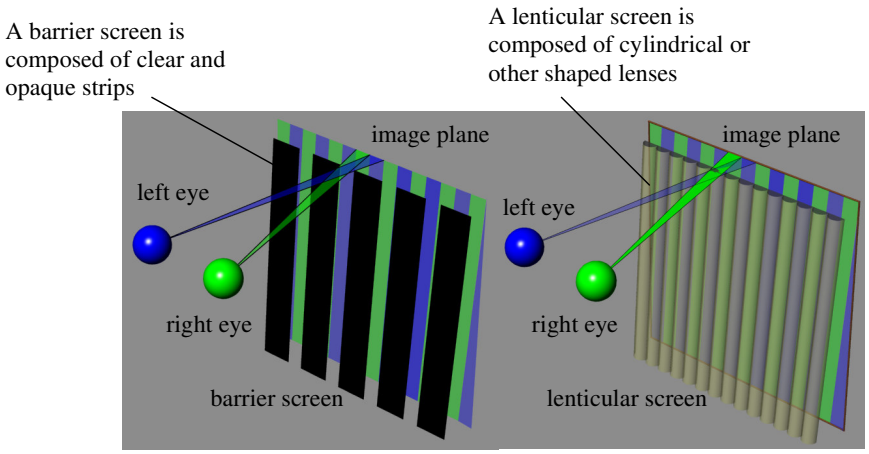

Figure 1: Varrier incorporates a static parallax barrier (left), where the left eye is represented by blue stripes and the right eye with green. The image stripes on the rear screen vary with viewer's motion while the barrier remains fixed. Functionally, the parallax barrier is equivalent to a lenticular screen (right).

In contrast, later approaches use the floating point domain, not only avoiding rounding errors but permitting sub-pixel accuracies to be achieved. A straightforward method, and the one used in Varrier for several years models the parallax barrier as a collection of polygons within the virtual world, so that the *virtual linescreen* [2] occludes parts of the scene behind it. Sandin et al. examine three Varrier occlusion algorithms according to the number of passes performed, and appraise the comparative benefits in terms of quality and performance in [3].

The newest algorithm leverages the graphics programmable pipeline by spatially multiplexing eye channels in GPU hardware. Varrier Combiner [1], [8] is named for its deferred order of operation. Individual eye images are rendered to off-screen buffers, and then these buffers are *combined* in a final step. Like the occlusion methods, computation is carried out with floating point accuracy to sub-pixel precision, and performance advantages result from execution in GPU hardware and from the down-sampled size of the off-screen buffers.

2.3 Tracking

From its inception, Varrier was intended to be a VR system, specifically to become the AS successor to the CAVE [18]. One of the principal components of any VR system is tracking; the system must first know the eye positions in space before it can steer imagery to the eyes. The original Varrier implementations relied on tethered acoustic-inertial trackers, while later systems incorporate camera-based trackers with and without markers. The last category is the nearest approximation to completely unencumbered VR.

The Intersense 900 [19] was the first Varrier tracking system, and several Varrier implementations continue to use it. In fact, the original Varrier prototype was developed within the physical space of the CAVE, so that the CAVE's tracker could be borrowed for Varrier without being moved and re-calibrated. Acoustic-inertial trackers such as Intersense, along with electromagnetic trackers such as Ascension's Flock of Birds [20] are aggregately termed *tethered* systems, regardless of whether the sensor communicates with the controller by a cable or a battery-powered RF transmitter. In contrast, camera-based systems require no transmitting devices to be worn by the user.

Most camera-based trackers require some form of markers or targets to be worn, for which the camera system has been calibrated. Examples are neon colored or retro-reflective surfaces, or even LED light sources. Computer vision techniques analyze the resulting video imagery for the presence of the targets, and positions of the same target in multiple camera images are triangulated to compute depth. Examples are the ART [21] and Vicon [22] trackers. Markers are generally attached to a headband, baseball cap, or other lightweight headgear.

The ultimate in freedom is a tracker that requires no additional accessories to be worn, as in the camera based face tracker of [23], [29]. Developed at EVL for Varrier, it is the current state of the art in modern Varrier installations. It has

been deployed successfully in a variety of form factors, from the single-panel Personal Varrier up to the 65-panel cylindrical Varrier, and has proved to be comparable in accuracy, jitter, and latency to high-end commercial systems such as Intersense. Note that this is presently a 3 degree of freedom (DOF) tracker providing positional data only. The following descriptions show that this is sufficient for most smaller systems, although the lack of rotational data will re-emerge in the discussion of the CV-65.

3. SYSTEMS

3.1 V-1 (2001)

The original Varrier [2] was a 17 inch single-panel LCD display fitted with a parallax barrier mounted on an acrylic substrate. After extensive tests with various barrier screens, acceptable ghost levels of 5% were attained, although color distortions were significant. The monitor was positioned within the physical space of the CAVE and borrowed its tracking system. The frame from LCD shutter glasses minus the LC elements supported the tracking sensor, and later the sensor was attached to a headband, as shown in Figure 2. The acrylic substrate suffered from thermal instability due to heat generated by the LCD panel, distorting the resulting optics. However, as a proof of concept the system functioned surprisingly well. A major contribution of this work was the theoretical basis for the Varrier algorithm, which occludes regions of left and right images via the graphics depth buffer and a virtual model of the physical barrier. Hence, two channels are interleaved in floating point space as opposed to a sorting operation in pixel space. A single PC running Windows 2000 controlled the display.

3.2 CV-4 (2004)

The next iteration was a 4-panel tiled display on a 3-node Linux cluster with a socket-based ad-hoc message passing scheme. Acrylic substrate was replaced by glass for improved thermal stability. The color distortions were understood as a bi-product of the Varrier 1-pass algorithm which operates at pixel resolution only. The solution was to execute the algorithm three times per channel, with a 1/3 pixel shift of the virtual barrier between each pass. However, the performance was penalized by factor of 3, a tradeoff that would remain for several years to come. Intersense head tracking, still shown in Figure 3, was soon replaced by camera-based face recognition and tracking [23]. By the time of the IEEE VR'04 Open House, the system functioned as a 4-panel VR AS system with camera-based tracking requiring no additional gear to be worn by the user. An interactive VR chemical engineering application [24] was demonstrated at that venue.

Figure 2: The original Varrier was a single-panel system that borrowed head tracking from the CAVE in which the system was located.

Figure 3: The Varrier 4-panel featured a clustered tiled display system capable of viewing an interactive VR application.

3.3 CV-35 (2004)

A second, larger system was simultaneously unveiled at IEEE VR'04, the 35 panel Cylindrical Varrier. Because of its large size, the CV-35 continued to use the Intersense 900 tracker. CV-35 was featured in Sandin et. al [3] the following year, although by that time the camera-based face tracker had been tested on the CV-35 as well.

Shown in Figure 4, it is driven by an 18 node Linux cluster, utilizes the MPICH-2 [26] message passing interface for inter-process communication and synchronization, and in 2004, applications were written using the CAVELib [25] library. In recent years, EVL has migrated to a newer VR applications platform called Electro [27]. Graphics hardware consists of a dual-output NVIDIA Quadro FX3000 card per node. Image channels were rendered initially with the floating point Varrier 1/1 algorithm [3], then with the Varrier 3/3 and Varrier 4/1 subpixel algorithms, and currently by Varrier Combiner [1], [8] algorithm. The CV-35 still supports research and is a popular demonstration stop for EVL visitors.

As researchers used the CV-35 on a daily basis, a few disadvantages became evident, some common to all barrier strip AS systems and others specific to tracked two-channel systems. The color distortions were already eliminated, but at substantial performance cost. Perhaps the most obvious drawback was the sensitivity to system latency during head movement; even relatively slow movements cause the user to see black guard bands before the system can update. The fact that tracked barrier strip AS was inherently a single-user system became the other conspicuous limitation. The CV-35 exaggerates this disadvantage because of its large physical size, 5 feet in radius. Its physical space accommodates multiple persons, yet correct stereo imagery is directed to only one.

3.4 PV-1 (2005)

The large, permanent structure of the CV-35 together with its cost and single-user limitation suggested the need for a smaller, more portable form factor that could travel to conferences and other remote venues. Meanwhile, desktop monitors grew to a 30" size, so a single panel display was relatively immersive. Thus, the single-panel Personal Varrier, or PV-1, debuted at the iGrid'05 conference in San Diego, and was also demonstrated at Supercomputing'05 in Seattle.

A single computer acts as the rendering node, equipped with dual AMD Opteron™ processors and an NVIDIA Quadro™ FX4400 graphics card. The GPU has recently been replaced by the GeForce™ 7900 card. The camera-based face tracker is a standard feature on this display, and the smaller working area of the desktop system is compatible with the tracking range. Additionally, a medium-range Flock of Birds™ and a hand-held device called a Wanda™ provide tracked navigation.

Figure 4: The Varrier 35-panel system is a large size, cylindrical, highly immersive tiled display.

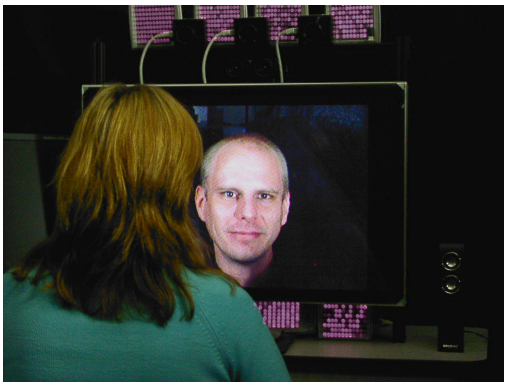

Figure 5: The single panel Personal Varrier display features a compact desktop form factor with AS teleconferencing capability.

The iGrid conference stresses distributed computing, and two distributed applications were demonstrated on the PV-1. The Point Packing and Visualization Engine (PPVE) [28] is an architecture that de-couples remote computation from local rendering, specifically targeted to the visualization of large datasets in Varrier. PPVE transfers color and depth maps from server to client, whereby the client re-constructs a subset of data for its local real-time interactive viewing [6].

The second distributed application was 3D AS teleconferencing, shown in Figure 5. With the addition of stereo video cameras and custom network software, stereo pairs of images plus synchronized audio were sent over a dedicated optical network called CAVEWave [6] in real time. Participants chatted with a live autostereo representation of their remote partner, and the interaction with a talking, floating head was quite compelling to conference attendees.

3.5 FV-6 (2006)

In 2006, EVL's Varrier group produced an intermediate size display for the Technology, Research, Education and Commercialization Center (TRECC) in West Chicago. Termed the FV-6, or flat Varrier 6-panel display, it is a floor mounted model like the CV-35. Driven by a small cluster of 4 rendering nodes plus two tracking nodes, the system incorporates the latest features of the PV-1 such as tracked wand interaction and stereo teleconferencing. Scientific and artistic visualizations are demonstrated at TRECC on the FV-6, along with live 3D teleconferences.

Figure 6 illustrates some of the accessories that surround the display panels. Infrared LEDs illuminate the subject; fourteen of these light gray illuminators are shown. Two speakers, one each on the left and right sides, provide audio for teleconferencing and other VR applications. Two teleconferencing video cameras are mounted directly above the displays, and three tracking cameras appear below. Finally, a beige-colored transmitter for the wand tracking appears near the bottom center of Figure 6.

3.6 CV-65 (2006)

The University of California at San Diego (UCSD) houses the largest Varrier to date, and perhaps the largest parallax barrier strip AS system in the world. Driven by 34 Linux nodes, it is an extension of the CV-35 to encompass a full 180 degree field of view. (The actual panels span just over 200 degrees at the outer edges). The nodes contain dual AMD Opteron™ processors and NVIDIA GeForce™ 7900 graphics cards. The Varrier Combiner algorithm leverages the features of this modern generation of graphics hardware to render autostereo in real time. Not only is the CV-65 the largest Varrier, but a fine-pitch barrier screen enhances image quality. The CV-65 is built around a 560 lines per foot (lpf) barrier, instead of the previous 280 lpf.

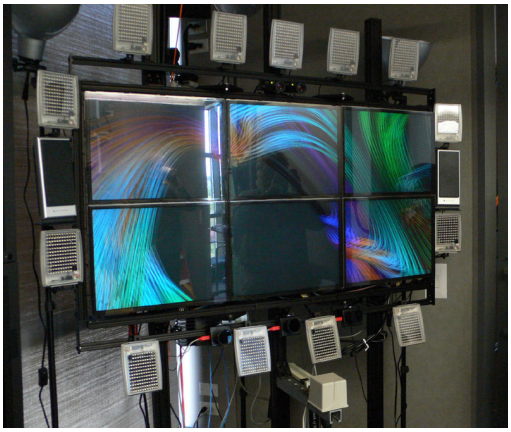

Figure 6: The Varrier 6-panel display is a smaller version of the 35 panel system. It also offers camera-based head tracking and AS teleconferencing.

Figure 7: The Varrier 65-panel display is currently the largest Varrier to be deployed. It improves over previous Varrier installations not only in size but in quality as well. Application depicting fly-over of Mars topography courtesy of Robert Kooima.

The 2X increase in barrier frequency affords a visual acuity increase of 1.4X [8], and the lines are barely noticeable from the average viewing distance at the center of the 5 foot radius of the screens. The CV-65 has a wide field of view that requires substantial head rotations in order to view the side panels. Therefore, the position-only data limitation of the 3 DOF camera-based face tracker now becomes significant. A software solution ameliorates the lack of rotation, but research is ongoing to add the additional three degrees of freedom to the face tracker.

3.7 PV-3 (2007)

Completed in January of 2007, the newest Varrier is an expanded 3-panel desktop system. Two of the panels are autostereoscopic; the third is a 2D console for viewing text and other 2D content. Apple™ 30-inch cinema LCD monitors comprise all three displays. A 3D teleconferencing feature is included, but unlike the PV-1, participants can discuss data shown on the second Varrier screen. When not teleconferencing, visualizations may span both Varrier displays, providing additional resolution. The PV-3 is shown in Figure 8. A single machine with dual Opteron™ processors and dual GeForce™ 7900 graphics cards serves as the rendering engine.

4. RESULTS AND CONCLUSIONS

The preceding timeline encapsulates a progression of problems and solutions, and it is instructive to highlight some of the key results.

4.1 Performance and image quality

Varrier Combiner results in at least a 3X performance speedup by replacing three rendering passes with one, by leveraging the programmable graphics pipeline of modern GPUs, and by down-sampling the off-screen buffers. In practice, even higher increases have been measured with scene complexities of over 100K vertices rendered in real time [8]. Autostereo rendering is now no more expensive than quad-buffered stereo rendering of the same resolution, and performance is GPU bound, hence it is expected to continue to improve with successive generations of graphics cards. Likewise, image quality has improved with performance. Earlier, color distortions were eliminated at significant cost, but today that cost overhead has disappeared while the quality improvements remain.

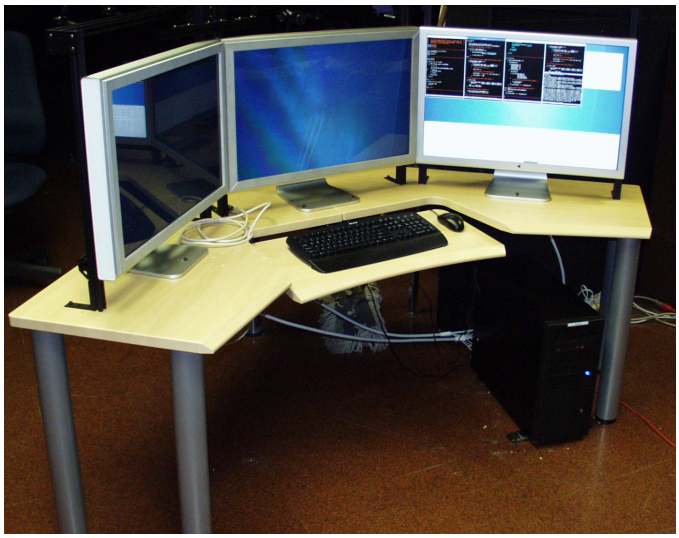

Figure 8: The 3-panel Personal Varrier is a larger desktop system with two AS panels and one conventional console.

4.2 Net effective resolution

The sub-pixel nature of the rendering algorithm permits the use of higher resolution line screens; currently the frequency of the barrier is approximately 560 lines per foot, or twice that of earlier models. While the same fraction of the native screen pixels is visible (approximately ¼ in the horizontal direction), that fraction is divided into a higher number of smaller elements. Note that color information is incomplete in the sub-pixel size elements, but this is not a serious limitation since color data tends to be of a lower spatial frequency. This higher number of smaller elements translates into an increased effective resolution and ultimately into better visual acuity, as shown in [8] with an improvement in acuity from 20/40 to 20/30.

The actual computation of net effective system resolution can be confusing, as several approaches are possible. One may compute the fraction of screen pixels (eg. ¼) and then apply some corrective factor to account for the higher number of smaller sub-pixels. Alternatively, beginning with the visual acuity, one may compute the minutes of arc resolved and convert this to a minimum number of pixels that are resolvable. The third approach, used in Table 1 below, is to simply count the number of barrier lines, since there is exactly one eye channel per line, and the channels are so narrow that there is no more than one resolvable unit per eye channel. Fortunately, the three methods converge to within about 10-15% in most cases. In all cases, the minimum resolvable unit in the horizontal direction equals one unit of effective resolution. In the vertical direction, one unit of effective resolution simply equals one screen pixel. The product of the horizontal and vertical resolutions equals the net effective system resolution. To be absolutely correct, this value is reported "per eye," although it is not correct to claim that twice the resolution exists for both eyes together.

4.3 Tracking

The camera based face tracking system consists of an artificial neural network technique to recognize and track faces. Three cameras are utilized along with infrared illumination to track users at 120 Hz to tolerances of several millimeters, comparable to commercial tethered tracking systems. The advantage over these other systems is the elimination of wires, headbands, sensors, and markers, and this freedom complements the ideals of AS. The camera based face tracker represents the state of the art in Varrier tracking.

4.4 Construction and calibration

The actual parallax barrier is high resolution Kodalith™ film printed on a Dolev™ 450 or Dolev 800 image setter, a machine used in the Pre-Press industry. Resolution of the image setter is 140 lines per mm, or 3560 dpi, and image files are generated in binary TIFF format @ 355.6 dpi resolution. The film is laminated to a glass pane, usually 1/8 inch or 3/16 inch thick, and the assembly is mounted in front of the LCD panel offset by some spacer thickness. The glass may be oriented with the film facing inward or out, depending on the desired optical thickness.

The optical thickness, or barrier offset corrected for refraction, determines the minimum, maximum, and optimal view distances of the system. The minimum and maximum distance formulae were derived in [3] for the pixel-scale parallax barrier (280 lpf), and can be converted to the sub-pixel scale barrier (560 lpf) by replacing the pixel width by the channel width. The channel width, or width of the transparent portion of a cycle, is generally ¼ of the total cycle, and substituting yields the following:

$$z_{min} = 4t * (e - p) / (3p) \qquad (1)$$

$$z_{max} = 4t * e / p \qquad (2)$$

Where z_{min} = minimum view distance, z_{max} = maximum view distance, t = optical thickness, e = interocular distance, and p = barrier period (ie, the reciprocal of frequency). Since the distance equations are non-linear, the optimal view distance does not occur at the arithmetic mean of the minimum and maximum; rather, it is defined in [7] as the distance where channels and guard bands are symmetrically spaced:

$$z_{opt} = 2t * (e - p) / p \qquad (3)$$

The distance values in Table 1 are computed using equations 1,2, 3, indicating optical parameters of the systems. Actual working values may be smaller depending on tracker limitations.

The ability to calibrate the system in software remains as a strength of the Varrier method. Calibration of shift, pitch, angle, optical thickness, and duty cycle is software adjustable through both an on-screen control panel as well as with keyboard shortcuts. An automated camera calibration method has been used [3], as well as manual methods. In practice, the optical calibration of a multi-panel system such as the CV-65 requires several hours to one day, depending on other circumstances such as tracker calibration. Maintaining calibration is relatively easy; occasional adjustments due to tracker drift can be accomplished in fifteen minutes for a large system.

5. FUTURE WORK

Over the past six years, Varrier has evolved from a prototype research project in parallax barrier autostereo into a reproducible system available in a number of sizes and shapes. The growth from prototype to production has fleshed out problems in performance, quality, tracking, construction and calibration, and other practical engineering details that are overlooked in early research stages.

Table 1: Varrier systems are quantified according to a number of criteria.

System	Number of panels	Panel diagonal size (in)	Panel native resolution	Net effective system resolution	Minimum, optimal, maximum view distance (in)	Tracking system
V-1	1	17	1600 x 1024	390x1024 .4 MP	41, 61, 125	IS-900
CV-4	4	20	1600 x 1200	740 x 2400 1.8 MP	33, 49, 100	IS-900, face tracker
CV-35	35	20	1600 x 1200	2600 x 6000 15.6 MP	33, 49, 100	IS-900, face tracker
PV-1	1	30	2560 x 1600	600 x 1600 1.0 MP	15, 22, 45	face tracker
FV-6	6	20	1600 x 1200	1110 x 2400 2.6 MP	33, 49, 100	face tracker
CV-65	65	20	1600 x 1200	9230x6000 55.4 MP	41, 61, 123	ART, face tracker
PV-3	2 varrier 1 console	30	2560 x 1600	2270x1600 3.6 MP	14, 21, 43	face tracker

Solutions to some of these problems have been found, and others are still being sought. EVL continues to focus on the following problem areas in both static and dynamic barrier AS research:

- complexity, cost, and performance of camera based face tracking
- reduction of system latency and AS sensitivity to latency
- more than one viewer
- software corrections to spherical aberrations at wide fields of view
- new form factors such as combinations of AS and haptic devices

Thus far, all of the Varrier systems relied on the same optical principles of tracked 2-channel static parallax barrier autostereoscopy. Optimizations have been discovered through the years, but as in all AS techniques, there are pros and cons inherent to this method. Some of the advantages are simplicity of construction and tracked first person VR perspective. The drawbacks are: fixed view distance range, sensitivity to latency during head movements, and the single tracked user restriction. All three of these disadvantages are due to the static nature of the barrier, and can be mitigated with a fully addressable, dynamic barrier. This research will be presented later this year at the IEEE VR'07 conference in Charlotte, NC, and will be published [7] after March 2007.

6. ACKNOWLEDGEMENTS

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago specializes in the design and development of high-resolution visualization and virtual-reality display systems, collaboration software for use on multi-gigabit networks, and advanced networking infrastructure. These projects are made possible by major funding from the National Science Foundation (NSF), awards CNS-0115809, CNS-0224306, CNS-0420477, SCI-9980480, SCI-0229642, SCI-9730202, SCI-0123399, ANI 0129527 and EAR-0218918, as well as the NSF Information Technology Research (ITR) cooperative agreement (SCI-0225642) to the University of California San Diego (UCSD) for "The OptIPuter" and the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement (SCI 9619019) to the National Computational Science Alliance. EVL also receives funding from the State of Illinois, General Motors Research, the Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC), and Pacific Interface Inc. on behalf of NTT Optical Network Systems Laboratory in Japan. Varrier and CAVELib are trademarks of the Board of Trustees of the University of Illinois.

REFERENCES

1. R. Kooima, Varrier Combiner, http://www.evl.uic.edu/rlk/varrier_combiner/varrier_combiner.html
2. D. Sandin, T. Margolis, G. Dawe, J. Leigh, T. DeFanti, The Varrier Autostereographic Display, *Proceedings of SPIE*, vol. 4297, San Jose, California, 2001.
3. D. Sandin, T. Margolis, J. Ge, J. Girado, T. Peterka, T. DeFanti, The Varrier ™ Autostereoscopic Virtual Reality Display, *ACM Transactions on Graphics, Proceedings of ACM SIGGRAPH 2005* 24, no. 3, p. 894-903, 2005.
4. J. Ge, D. Sandin, T. Peterka, T. Margolis, T. DeFanti, Camera Based Automatic Calibration for the Varrier™ System, *Proceedings of ProCams 2005, IEEE International Workshop on Projector-Camera Systems*, 2005.
5. J. Ge, D. Sandin, A. Johnson, T. Peterka, R. Kooima, J. Girado, T. DeFanti, Point-based VR Visualization for Large-Scale Scientific Datasets by Real-Time Remote Computation, *Proceedings of ACM VRCIA 2006 Conference*, 2006.
6. T. Peterka, D. Sandin, J. Ge, J. Girado, R. Kooima, J. Leigh, A. Johnson, M. Thiebaux, and T. DeFanti, Personal Varrier: Autostereoscopic Virtual Reality for Distributed Scientific Visualization, *Future Generation Computing Systems*, 22, 2006.
7. T. Peterka, R. Kooima, J. Girado, J. Ge, D. Sandin, A. Johnson, J. Leigh, J. Schulze, T. DeFanti, Dynallax: Solid State Dynamic Barrier Autostereoscopic VR Display, *IEEE VR 2007 Conference Proceedings*, (accepted for publication), 2007.
8. R. Kooima, T. Peterka, J. Girado, J. Ge, D. Sandin, T. DeFanti, A GPU Sub-pixel Algorithm for Autostereoscopic Virtual Reality, *IEEE VR 2007 Conference Proceedings*, (accepted for publication), 2007.

9. C. Van Berkel, Image Preparation for 3D-LCD, *Proceedings of SPIE Vol. 3639 Stereoscopic Displays and Virtual Reality Systems VI,* San Jose, California, 1999.
10. D.F. Winnek, U.S. patent number 3,409,351, 1968.
11. T. Okoshi, Three Dimensional Imaging Techniques, Academic Press, N.Y., 1976.
12. A. Schmidt, A. Grasnick, Multi-viewpoint Autostereoscopic Displays from 4D-Vision, *Proceedings of SPIE Photonics West 2002: Electronic Imaging*, San Jose, California, 2002.
13. K. Perlin, S. Paxia, J. Kollin, An Autostereoscopic Display, *Proceedings of ACM SIGGRAPH 2000*, ACM Press / ACM SIGGRAPH, New York. Computer Graphics Proceedings, Annual Conference Series, ACM, 319-326, 2000.
14. K. Perlin, C. Poultney, J. Kollin, D. Kristjansson, S. Paxia, Recent Advances in the NYU Autostereoscopic Display, *Proceedings of SPIE*, vol. 4297 San Jose, California, 2001.
15. J.-Y. Son, S.A. Shestak, S.-S. Kim, Y.-J. Choi, Desktop Autostereoscopic Display with Head Tracking Capability, *Proceedings of SPIE* Vol. 4297, Stereoscopic Displays and Virtual Reality Systems VIII, San Jose, California, 2001.
16. L. Lipton, M. Feldman, A New Autostereoscopic Display Technology: The SynthaGram, *Proceedings of SPIE Photonics West 2002: Electronic Imaging,* San Jose, California, 2002.
17. D. Sandin, E. Sandor, W. Cunnally, M. Resch, T. DeFanti, M. Brown, Computer-Generated Barrier-Strip Autostereography, *Proceedings of SPIE, Three-Dimensional Visualization and Display Technologies*, vol. 1083, pp. 65-75, 1989.
18. C. Cruz-Neira, D. Sandin, T. DeFanti, Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE, *Proceedings of ACM SIGGRAPH 1993*, ACM Press / ACM SIGGRAPH, New York. Computer Graphics Proceedings, Annual Conference Series, ACM, 135-142, 1993.
19. InterSense IS-900 Precision Motion Tracker. http://www.isense.com/products/prec/is900
20. Flock of Birds. http://www.ascension-tech.com/products/flockofbirds.php
21. A.R.T GmbH. http://www.ar-tracking.de
22. Vicon. http://www.vicon.com
23. J. Girado, D. Sandin, T. DeFanti, L. Wolf, Real-time Camera-based Face Detection using a Modified LAMSTAR Neural Network System, *Proceedings of IS&T/SPIE's 15th Annual Symposium Electronic Imaging 2003*, Applications of Artificial Neural Networks in Image Processing VIII, San Jose, California, 2003, pp. 20-24.
24. T. Peterka, Scientific Visualization of N-Dimensional Attainable Regions. Master's thesis, University of Illinois at Chicago, Chicago, 2003.
25. CAVELib. http://www.vrco.com/CAVELib/OverviewCAVELib.html
26. MPICH2. http://www-unix.mcs.anl.gov/mpi/mpich/
27. R. Kooima, Electro. http://www.evl.uic.edu/core.php?mod=4&type=2&indi=296
28. J. Ge, D. Sandin, A. Johnson, T. Peterka, R. Kooima, J. Girado, T. DeFanti, Point-based VR Visualization for large-scale scientific datasets by real-time remote computation. *Proceedings of ACM VRCIA 2006 Conference*, 2006.
29. J. Girado, Real-Time 3d Head Position Tracker System With Stereo Cameras Using A Face Recognition Neural Network, PhD thesis, University of Illinois at Chicago, 2004.