Conference Proceedings 2003

Theory and Practice of Computer Graphics

The University of Birmingham
Birmingham, UK
3-5 June 2003

Edited by Mark W Jones

www.eguk.org.uk
'brought to you by Eurographics UK'

Proceedings

Theory and Practice of Computer Graphics

University of Birmingham, UK
3-5 June 2003

Sponsored and supported by
Blackwells Publishing
Elsevier
Hewlett Packard LtD
Reachin Technologies AB

IEEE
COMPUTER
SOCIETY
http://computer.org

Los Alamitos, California
Washington • Brussels • Tokyo

IEEE Computer Society Order Number PR01942
ISBN 0-7695-1942-3
Library of Congress Number 2003105394

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1-714-821-8380
Fax: + 1-714-821-4641
E-mail: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1-732-981-0060
Fax: + 1-732-981-9667
http://shop.ieee.org/store/
customer-service@ieee.org

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Individual paper REPRINTS may be ordered at: reprints@computer.org

Editorial production by Bob Werner
Cover art production by Joe Daigle/Studio Productions
Printed in the United States of America by The Printing House

IEEE COMPUTER SOCIETY

IEEE

Proceedings

Theory and Practice of Computer Graphics

TP•CG•03

IEEE COMPUTER SOCIETY

IEEE

Conference Committee

Conference Chair
Paul Lever

Program Chair
Mark W. Jones

Local Organization Chair
Paul Hatton

The International Program Committee
Ken Brodlie (Leeds, UK)
Min Chen (Swansea, UK)
Andy Day (East Anglia, UK)
Yiorgos Chrysanthou (Cyprus)
Neil Dodgson (Cambridge, UK)
David Duce (Oxford Brookes, UK)
Rob Fletcher (York, UK)
Vladimir Geroimenko (Plymouth, UK)
Joerg Haber (Max Planck, Saarbrucken, Germany)
Peter Hall (Bath, UK)
Nicholas Higgett (De Montfort, Leicester, UK)
Roger Hubbold (Manchester, UK)
Nigel John (Manchester Visualisation Centre, UK)
Huw Jones (Middlesex, UK)
Mark Jones (Swansea, UK)
Leif Kobbelt (Aachen, Germany)
Argy Krikelis (Aspex Technology, UK)
Alf Linney (UCL, UK)
Ralph Martin (Cardiff, UK)
Theresa-Marie Rhyne (North Caroline State, US)
Jonathan C Roberts (Kent, UK)
Bob Stone (Virtual Presence, UK)
Martin Turner (De Montfort, Leicester, UK)
Gordon Watson (Edinburgh, UK)
Phil Willis (Bath, UK)
Derek Wills (Hull, UK)
Guang-Zhong Yang (Imperial, London, UK)
Jian Jun Zhang (Bournemouth, UK)

Invited Speakers

The Need for Realism: Graphics and Archaeology

Alan Chalmers, Department of Computer Science, Bristol, UK

Recent developments in computer graphics are providing powerful tools to help archaeologists investigate the multi-dimensional aspects of data they have gathered. Computer graphics techniques can be used to reconstruct and visualise features of ancient environments which may otherwise be difficult to appreciate. However, if we are to avoid the very real danger of misrepresenting the past, then the computer generated images should not only look "real", they must also simulate very accurately all the physical evidence for the the site being considered.

This paper discusses the application of high fidelity computer graphics with respect to the "real" archaeological scenes they are intended to depict.

Biographical Sketch

Alan Chalmers is a Reader in the Department of Computer Science at the University of Bristol. He has published over 90 papers in journals and international conferences on high fidelity graphics. He is a former Vice President of ACM SIGGRAPH. His research is investigating the use of very realistic graphics in the accurate visualisation of archaeological site reconstructions and techniques which can be used to reduce overall computation time of high quality images without reducing the perceptual quality of the images.

Freeform Shape Representations for Efficient Geometry Processing

Leif Kobbelt, Department of Computer Science, Aachen, Germany

The most important concepts for the handling and storage of freeform shapes in geometry processing applications are parametric representations and volumetric representations. Both have their specific advantages and drawbacks. While the algebraic complexity of volumetric representations $S = \{(x,y,z) \mid f(x,y,z) = 0\}$ is independent from the shape complexity, the domain D of a parametric representation $f : D \rightarrow S$ usually has to have the same structure as the surface S itself (which sometimes makes is necessary to update the domain when the surface is modified). On the other hand, the topology of a parametrically defined surface can be controlled explicitly while in a volumetric representation, the surface topology can change accidentally during deformation. A volumetric representation reduces distance queries or inside/outside tests to mere function evaluations but the geodesic neighborhood relation between surface points is difficult to resolve. As a consequence, it seems promising to combine parametric and volumetric representations to effectively exploit both advantages.

In this talk, a number of projects is presented and discussed where such a combination leads to efficient and numerically stable algorithms for the solution of various geometry processing tasks. Applications include global error control for mesh decimation and smoothing, topology control for level-set surfaces, and multiresolution editing without local self-intersections.

Biographical Sketch

Leif P. Kobbelt is a full professor and the head of the Computer Graphics group at the Aachen University of Technology, Germany. His research interests include all areas of Computer Graphics and Geometry Processing with a focus on multiresolution and free-form modeling as well as the efficient handling of polygonal mesh data. He was a senior researcher at the Max-Planck-Institute for Computer Sciences in Saarbrüken, Germany from 1999 to 2000 and received his Habilitation degree from the University of Erlangen, Germany where he worked from 1996 to 1999. In 1995/96 he spent a post-doc year at the University of Wisconsin, Madison. He received his master's (1992) and Ph.D. (1994) degrees from the University of Karlsruhe, Germany.

Image of Archaeological Geophysics from the Small Screen

Chris Gaffney, GSB Prospection, UK

There are few more rewarding experiences than confidently predicting to over 3 million people what is buried under the ground, watching the 'ground truth' appear,.and being thrust back in front of the cameras to spontaneous applause. Anyway, less of my dreams, the reality of TV archaeology is usually very different and does not always end to great acclaim.

In this talk I shall chart the use, and occasional abuse, of geophysical techniques in archaeological investigations on television. I hope to answer the questions:

- Does my image look big (on the small screen)?
- Can we say anything other than the blindingly obvious on TV?
- Why do we put ourselves through mental torment and professional suicide on a regular basis?

Biographical Sketch

Chris has worked in geophysics since 1983, including extensive site-based experience in the UK, Greece and the former Yugoslavia. In 1989, he formed a partnership with John Gater at GSB Prospection. He, too, is an associate editor of the Journal of Archaeological Prospection.

Chris has worked on the Time Team digs. Athelney was probably his favourite Time Team dig ('a cracker'), but he also remembers Tockenham — the site of a Roman villa in Wiltshire — fondly: 'It was huge. The scale of the results was never really captured on the programme. Everything was really clear, and for once on Time Team, the gradiometry worked well — usually it's only resistance that does, unlike 90 percent of our other work.' The excavation in Maryland was also memorable: 'It was amazing how the remains of the first brick-built building in Maryland just popped out. This happened on the first morning of the first day, which made the rest of the dig fairly anti-climactic for us but got everyone else off to a running start.'

Chris' ideal site is a monastery: 'Monastic sites conform to certain patterns, and are nice and simple and very clear.' However, he admits that simplicity and clarity are not the words that he would use to describe the geophysics team's experience.

Table of Contents

Theory and Practice of Computer Graphics 2003

Rendering and Modelling

Virtual Worlds

Volumes

Work In Progress

IMAGES AND DISPLAY

Designing Projective Environments

Apostolos Dountsis and Paul F. Lister
Centre for VLSI and Computer Graphics,
University of Sussex, Brighton, UK
A.Dountsis@sussex.ac.uk P.F.Lister@sussex.ac.uk

Abstract

A number of projection systems have been constructed that create an immersive experience for the viewer. However, little consideration has been given to the development of projection systems that can use the surfaces in our everyday environment as their target.

This paper discusses the progress that has been achieved in the area of projection systems and the technologies involved. A system is presented that facilitates the creation of a multi-projection immersive environment using display surfaces such as walls, desks, bookcases and other furniture that may exist in our living environment.

1. Introduction

The conventional way to view visual information is using a CRT/LCD display. Nowadays, such screens vary from 12" to 28". The progress in increasing viewing dimensions is slow due to the considerably high cost. Several new types of displays have emerged such as the widescreen and plasma screens in an attempt to create a cinematic atmosphere. The latest plasma screens have increased their viewing dimensions up to 61". It is obvious that these devices are not in a position to create an immersive experience.

The user has to compromise by sitting as close as possible to the monitor since it is not possible to cover their field of view with such a device. Actually, a 24" monitor can subtend about 43° horizontal field-of-view (FOV) when viewed at a distance of 2 feet. The horizontal FOV for both eyes is approximately an ellipse that is 130° vertical by 200° horizontal. These values do not take into account any head motion, which can increase the angle to over 270° [2][3].

The viewing dimensions of TVs may increase in the future but they will not fulfil the immersive requirement. Even if they managed to do so they still need to produce a resolution far higher than the 72dpi that they currently produce [2].

Ivan Sutherland perceived the need for an alternative type of display in 1965 [1] that would offer total immersion. His vision led to the manufacturing of the Head-Mounded-Display (HMD) [10]. This is a solution that satisfies the task of delivering an immersive environment to the user. They have the disadvantages of considerable cost and the fact that each user needs to physically wear the device to experience the virtual environment. Technological advances enable us to develop relatively affordable, light and wireless HMDs but some constrains persist like the need to increase resolution [8].

Alternatively, we can attempt to generate an immersive feeling using the existing surrounding environment by displaying images on it.

An approach to this concept is to use a number of projectors as the display source. However, such a system is not straightforward to setup and operate and a number of issues have to be dealt with in order to make such a system operational. When more than one projector is used seams may be created due to projection overlaps, alternatively, perfect tiling is not a trivial task as projectors require additional adjustment due, for example, to heat expansion. Furthermore, buildings contain many different coloured surfaces which are not always flat. Therefore, such a system has to take into account these issues and compensate for them.

This paper presents a system that performs all the operations described above and allows multi-projection systems to be configured and operated in a virtual environment. Furthermore, it is envisaged that the described system would enable a group of people to experience total immersive virtual worlds in the comfort of their living room. The display surfaces intended for use in this system could be walls in offices [4], bedrooms or even sports centres, opera theatres [5], exhibition halls or conference rooms. The projection would physically surround the viewers and could create a truly immersive environment.

2. Background Research

This chapter concentrates on the work that has taken place in the creation of immersive projected display systems. Typical steps are identified for a pre-distorted multi-projection system. Existing implementations and their technologies are discussed.

0-7695-1942-3/03 $17.00 © 2003 IEEE

2.1. Fundamental Requirements

The goal of such a system is to produce a seamless immersive environment by placing several projectors in a room. In order for this to operate, it is vital that an association is established between the projectors and the room (projected) space. Therefore, a 3D representation of the room needs to be generated. This should include the shape of the display surface, the position and orientation of the projector(s) and the position of the user. Several registration and image-based rendering techniques are available that can assist with this process.

If more than one projector is used then possible overlaps can occur. This is a very important issue that needs to be resolved otherwise the produced seams are noticeable to the user. The ideal solution is to blend the individual parts so that a seamless result is achieved.

If the projected surface is not planar and the projectors optical axes are not perpendicular to the projection surface, then the projections will appear distorted. Therefore, some pre-distortion needs to be applied on the projection so that when the image source or video frame is projected, it should appear undistorted.

2.2. Projects

A number of projects have been designed and implemented in an attempt to create immersive environments. These projects vary from domed displays to non-planar multi-projection systems such as the "Office of the Future" [4][13].

CAVE™ is a project that was announced in 1992 [15] and since then has spawned a number of variations. Cruz-Neira describes it in [16] as "a theatre with dimensions 10'x10'x10' ". Three projection screens are used as the surrounding walls and one for the floor. Rear-projections are used; one per cube side. The user's head is tracked and stereoscopic glasses enable the user to gain a feeling of depth. Only a single user can operate it at a time.

There is no warping or blending involved since each projection is to a different plane. Moreover, since it uses one projector per surface, the edges and corners cannot be blended and therefore it does not create a seamless environment.

Another type of projected display is the dome [17]. It comprises a hemisphere that acts as the display surface. At the centre of the hemisphere, a single projector is placed that uses spherical projection to fill the dome. Such a system, allows a vast number of viewers to share the immersive experience. It is a solution that has achieved wide use in planetariums. However, domed displays cannot be used in private rooms such as domestic rooms or offices since the ceiling would need to be turned into a dome.

Ramesh Raskar in 1998 introduced the "Office of the Future" project [4][13]. It is the first description of a system, which uses the walls on an office as the display surface of a multi-projection system. Raskar attempts to provide a general solution to the problem of multi-projection in offices. A number of technologies are combined to achieve the goal and simplification to the display surfaces are required to speed-up the system.

Another project is the Infinity Wall [18]. This system uses one flat screen and four projectors which tile the display. It is designed around the same idea as CAVE and also uses the same software. Stereoscopic display is used with the system to give the illusion of depth to the scene. The main target of the Infinity Wall is to be used in classrooms.

A typical projected display system requires a number of technologies to be combined to enable the creation of a seamless immersive environment.

2.3. Registration

In order to evaluate the scene and create a pre-warped projected system, we need to capture the scene and its objects (display surface and projectors) in a computer model. By doing so, we will be able to calculate the necessary blending and warping algorithms for the input frame that is to be projected onto the real surfaces.

The area of image-based modelling is able to give an effective solution to the issue presented in this section. Cameras can be used to obtain the scene information and transfer it to a 3D model that can be used in a projection system. The required scene information includes depth and colour extraction.

Increased accuracy is crucial to such a system since the projections need to be correct. Moreover, if the scene changes then the registration process needs to identify the adjustment and compensate accordingly. Furthermore, image-based modelling techniques tend to produce models with an extremely high number of polygons. Intelligent polygon reduction may be required to make the model useful in a real-time rendering application such as projected displays.

Raskar [4] makes use of the imperceptible structured light technique. This is a technique for tracking and acquiring 3D scene geometry and surface characteristics. It is patented to Henry Fuchs and the University of North Carolina at Chapel Hill [9].

2.4. Blending

When we project a scene from multiple projector units then some of the projections may overlap. This can happen because the projectors do not have to be positioned in any particular order (i.e. parallel to each other). This could cause seams to appear in the overlapped regions (see Figure 1) that would have a negative affect on the immersive sensation portrayed to the viewer.

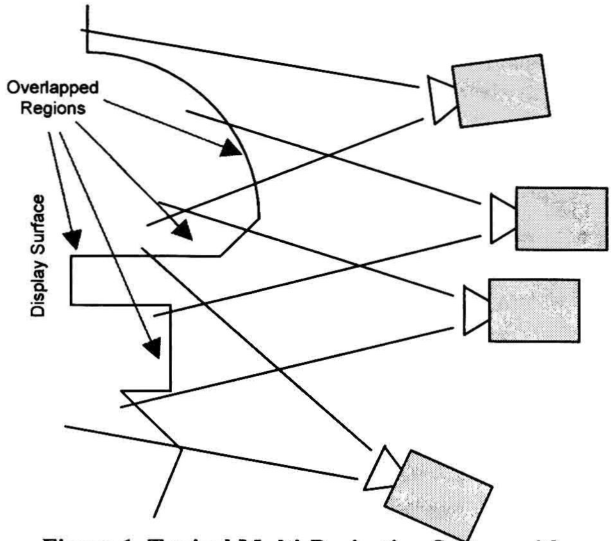

Figure 1. Typical Multi-Projection System with overlapping regions

Consequently, we need to modulate the intensity contribution of each pixel on every projector. When a uniform colour and intensity bitmap is input into the frame buffer of the computer, the resulting image, when projected across the system, will appear uniform for any particular brightness level.

P. Lyon [11] and [12] describe the problems and suggest manual solutions. Surati [7] builds upon Lyon's observations.

The seams can be eliminated by applying blending functions to the overlapped pixels of each projector. The resulting brightness of each pixel in the projected image should sum to unity and therefore make the seams invisible. Black level may need optical enhancement.

2.5. Warping

Carolina Dorsey [5] demonstrates an algorithm that pre-distorts an image in such way that when projected and viewed from the position of an ideal viewer, the projected image appears undistorted (see Figure 2).

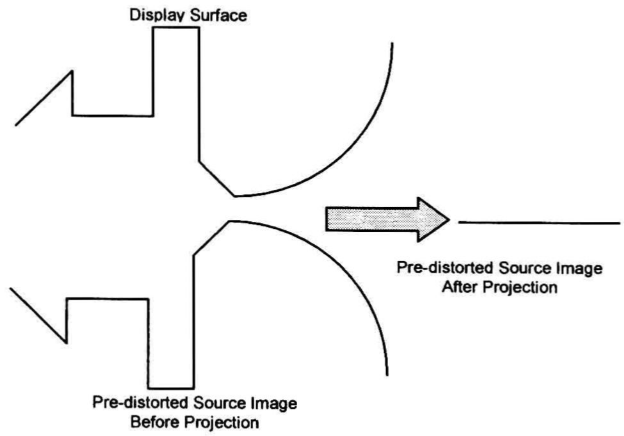

Figure 2. Warping an image based on the display surface

Dorsey explains that image warping techniques could simulate and resolve the distortion problem. She computes the pre-distorted slide based on transformations. The problem with the use of transformations is that for complex surfaces the analytical computation is non-practical. She avoids this computation by computing a small number of points and interpolating the rest. The degree of the interpolation increases with the complexity of the surface.

2.5. Projection

Several attempts have been made in the area of computer graphics to render an image as the result of a pin-hole projection.

July Dorsey published a paper [5] in SIGGRAPH '91 that focuses on the simulation of large environments like opera houses and facilitates in the stage design. She considers lighting coordination, projection systems and set designs. A novel projection approach was presented to simulate the optical effect of scenic projectors.

She simulates a real projection system where each pixel is evaluated based on the equation that calculates the radiosity at a vertex on the display surface.

Radiosity is a very important component in computer graphics but it lacks significant hardware acceleration. Therefore, the required computations are mainly taking place on the CPU making radiosity algorithms very slow and inappropriate in a real-time multi-projection system.

The following year Mark Segal [6] extended the Dorsey [5] projection technique and the Upstill [14] one. He describes a projective texture mapping technique to achieve realistic and real-time projection rendering.

The use of texture maps allows the projective textures to take advantage of hardware acceleration and operate on a real-time scale. His algorithm has now been included in the OpenGL architecture.

3. System Concept

Projection systems are certainly a promising way forward in the realisation of immersive environments. Currently such systems operate in a well-defined space such as a CAVE or in a simplified area where the projection occurs on planar surfaces. For more advanced systems, a specialist, usually an engineer needs to be employed to calibrate the system. This cannot be considered as a one-off duty because minor corrections may continuously be required due to thermal effects (heat displaces the projectors) in a small period of time.

Our application delivers a novel approach to the creation and development of a multi-projection system.

The system models the projection and pre-distorts the source image (either still or video frame) so that the resulted projection would appear undistorted from the camera's position and orientation. The camera in the virtual environment represents the viewer in the real setting. Moreover, it blends the individual projections so that a seamless projection materialises.

Such a system could be used to bring projected immersive environments into the home or office. The space that is required for the display surface can be modelled. For instance, we could design our living room and furniture and then import this model into the virtual environment. The system should be able to handle possible overlaps and pre-warp the source image in such a way that when projected in the modelled space, it does not appear distorted by the scene's objects.

4. System Design

The system is based on the notion of offering a complete pre-warped multi-projection system that projects onto any physical space. It is designed to be a low cost solution. PCs are becoming more powerful and less expensive day by day. Moreover, OpenGL acceleration is considered as standard in all the graphics cards. Thus, our system is designed as a PC application that takes advantage of OpenGL to gain access to hardware acceleration.

Its operation has been divided into three phases (see Figure 3), design, blending and execution or projection phase.

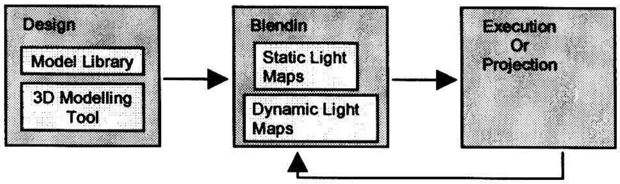

Figure 3. The Three phases of our system

During the design phase, the user can setup the virtual environment from a set of predefined objects such as projectors and a number of surfaces. Alternatively, a modelling tool or image-based modelling can be used to create any object to act as the display surface. This could be considered necessary when the replication of an existing setting is attempted. The produced model can then be imported and positioned on the virtual canvas. Finally, a still image or video needs to be provided as the media that we want to project onto the display surface.

When all the components are in place in the design mode then the system can initiate the blending phase. This enables the necessary blending that produces a seamless projection within the virtual environment. The system is able to handle single and multi projection scenarios (see Figure 4). The blending step is required only in the case where more than one projector is used.

The last phase is execution or projection. The warping is performed on the source image that would make the projection appear undistorted on the display surface to the viewer. As the viewer moves within the virtual environment, the system is able to compensate the produced distortion by dynamically modifying the pre-distortion of the source image. Furthermore, it allows the projection either to be simulated in software or drive a set of real projectors.

A modular approach has been followed that treats the warping and blending as separate components. Hence, different warping or blending algorithms may be used in the form of modules.

The display surface can be changed during the execution or projection phase. If the change signifies the appearance of additional seams then the system initiates the blending phase to correct them. The warping is being calculated on every execution cycle and therefore if the display surface is altered that would have no effect on the system. In case any of the projectors tilt or pan only the blending phase needs to be executed again.