

ARP: A Specialized Computer Architecture Applicable to Virtual Reality, Telerobotics and Interactive Television

Ronald Pose

Department of Computer Science,
Monash University,
Clayton, Victoria 3168,
AUSTRALIA
E-mail: rdp@cs.monash.edu.au

Abstract. Virtual Reality is attracting a great deal of attention. Its use depends critically on having extremely low response times of the virtual environment to user interaction. The problem of reducing the lag between the user's head position or orientation changing and the updating of the displayed view has been described as one of the grand challenges of computing. The specialized computer architecture and other techniques described here form the basis of a solution for this problem. A radical new approach to the problem is presented. The architecture and methodology also have applicability in areas as diverse as telerobotics and interactive television.

1 Introduction

Researchers and practitioners have long been grappling with the serious problem of latency in responding to movement of a user in a virtual environment. This problem is most difficult to handle in the visual domain since the sheer quantity of data and processing required to update the view delivered to the user is beyond the capabilities of even high-end computer graphics systems. The problem, while simply described as the lag between a user's head position or orientation changing and the updating of the displayed virtual view to reflect that change, has rather severe consequences, and has eluded solution by many practitioners. While it is well known how to do the computation required to generate computer graphics images appropriate to the orientation of the observer in the virtual world, the problem is computationally extremely expensive. This has meant that even the fastest of conventional graphics computers have not been able to achieve the desired performance of over 60 correct frames per second.

It turns out that by solving this problem in the domain of virtual reality systems one has also provided a tool with much more general application. The problem of latency or lag in virtual reality systems has an analogue in systems involving telerobotics. In such systems a remote robotic device is controlled by a human guided by video input from the robot. As with virtual reality systems, any lag in response to

changes in viewing orientation can lead to disorientation and a consequent loss of ability to achieve fine control of the robotic device.

Another potential application domain for this new architecture is in a form of interactive television system. A conventional television displays the viewpoint and orientation determined by the director and cameraman. Watching such a system is a purely passive affair. This new architecture offers the possibility of viewers being able to decide and change their viewing orientation in real-time, independently of viewers using other receiving apparatus. An interactive television system is possible using this technology and it can even be implemented in a way that is compatible with current television systems.

In this paper I will first give an outline of the new way of thinking about the original problem that led to the revolutionary new architectural approach. The key to the success of this project has been to consider the system as a whole. Thus a new technique of exploiting the new architecture has also been developed, leading to a complete new system which offers orders of magnitude improvement in virtual reality systems performance, assists with a key difficulty with current telerobotics systems, and opens up new possibilities for interactive television and computer graphics systems.

Having set the scene by changing the way one thinks about the problem, I go on to develop the specialized computer architecture which enables new algorithms to be used to update the computer graphics images with dramatic performance gains.

Then I will outline how this new technology can be applied to telerobotics and to interactive television. Finally I will conclude by indicating the current state of the project.

It should be mentioned at the outset that the technology described herein has not only been shown through simulation and analysis to be of great benefit, but has also been developed to the stage of a working prototype Address Recalculation Pipeline (ARP). Thus the ideas have been demonstrated to be practical.

2 Conventional Approaches to the Problems

One approach to building Virtual Reality systems has been to slow the update rate of the display so as to enable the computer system to cope with the processing load. This leads to a rather jerky displayed world with latency being at least the time between updates of the display. Experience with motion pictures and with television has shown that the higher the frame rate the better.

An alternative approach is to reduce the processing load by reducing the quality and resolution of the images displayed. While the display update rate is increased and latency somewhat reduced with this approach, the realism of the virtual world is sacrificed leaving one with cartoon-like imagery rather than anything approaching photo realism.

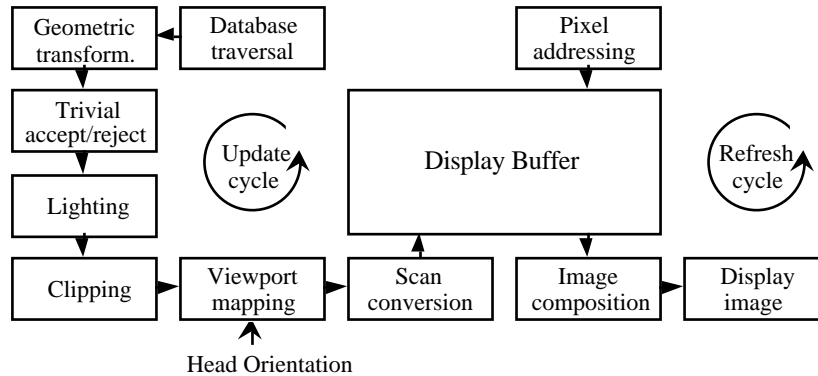


Fig. 1. Conventional Virtual Reality Display Architecture.

The state-of-the-art in conventional Virtual Reality systems rely on using the fastest available graphics processing and rendering equipment as depicted in Figure 1. This tends to be very expensive with the very best systems costing over a million dollars. In effect the problem is tackled indirectly through the use of massive amounts of processing power. Even so it can be observed that the trade-off between latency and image quality still has to be made, and that with reasonably realistic looking scenes the latency is still a significant problem. An extremely successful Virtual Reality computing system is the PixelPlanes system developed at the University of North Carolina [Fuchs et al., 1989]. This uses thousands of processors to give extremely impressive rendering performance and very good latency characteristics, however PixelPlanes is rather expensive technology which most developers could not afford.

3 Re-Examining the Problem

In order to deal with this problem of latency it is helpful to forget for a moment the conventional ways of computer graphics and instead look at the overall environment in which we are working.

The problem as stated earlier is one of mismatch between what is displayed to the user and what should be displayed according to the virtual world model in which the user is situated. This mismatch usually is due to the latency involved in updating the display to reflect the virtual reality. The easy solution, involving reducing image quality so one can keep up with the virtual world, creates a mismatch of a different kind, but still spoils the illusion of Virtual Reality.

Let us examine what is happening in the virtual world itself. It could be that some part of the virtual environment is changing or moving in some way. While indeed this does occur, it is also true that in general only a small part of the virtual world is changing at any time, and typically the virtual world contains many static objects or scenes. Thus we can deal with such activity without recomputing

everything, and since the objects involved are generally independent of the user, a slight delay in reflecting their changes does not cause the serious problems mentioned above.

The other event, which is much more significant, is the movement of the user within the virtual environment. Any slight movement of the user's head leads to a change in everything the user sees, hence it looks to the user as though the observed environment has moved completely. Thus one can see that it is the relative movement between the user and the objects in the virtual world which is observable to the user, and it is the user's own movement which has the most dramatic effect, since it affects the display of even static parts of the virtual world.

Now we should look more closely at the various types of user movement, to see what the consequences are. Consider the obvious movement, translation, for example walking forward. When translating, objects in the world appear to move relative to you. Close objects appear to move more than distant ones. Very distant objects may not appear to move much at all. Thus for translations we have to concentrate especially on close objects, and can largely ignore the distant background. In typical scenes the majority of the complexity of the scene is in the background, hence the problem appears tractable. We are also fortunate in that people do not tend to translate so very quickly so changes tend to happen at a manageable rate.

Another user movement is that of rotation of the head. Here we have a rather different effect. Even a small rotation immediately causes everything in the observed view to appear to move, even the background and static objects. There is also the added problem that rotations are much faster and more frequent than translations; one tends to look around and observe one's environment by rotating one's head or eyes. It looks as though everything has to be recomputed and re-rendered.

However, let us not forget the idea of stepping back and looking at the overall environment. It seems that the environment is staying relatively static but the user's view of the environment is changing quite rapidly and somewhat unpredictably, mainly due to rotation. The key point here is that the environment itself is not changing dramatically, so one needs to find a way in which the representation of the displayed environment also stays relatively static, and hence can be computed without the serious problems with which we are concerned.

4 A New Model

The ancient Greeks found a similar problem in trying to describe the motions of the planets and stars. They had a concept of planets moving relative to stars and to the Earth, and stars moving relative to the Earth. A model involving *crystal spheres* was developed. In this model the heavenly bodies were *painted* onto various layers of crystal spheres which were centred on the Earth and could move relative to one another. Objects on close spheres move more compared with objects on distant spheres. Various complex models of the movements of heavenly bodies were

formulated in terms of these spherical shells surrounding the Earth. We now know that much of the observed motion of the stars is due to the rotation of the Earth, but since the observed motion is relative, a successful model centred on the Earth is possible.

We can use a similar model for a virtual world. Concentric spheres centred on the user would have the objects of the virtual world painted on them. Close objects on inner spheres and distant objects on outer spheres. In essence all possible views from the user's position are already rendered onto the spheres, and it only remains to display appropriate views as seen by the user. A rotation merely involves displaying a different portion of the sphere. A translation will require some updating of the spheres, however outer spheres will change much less than inner spheres, and hence may require little or no updating. Of course changes within the virtual environment will have to be reflected in the spheres, but typically only a subset of the spheres will be involved.

5 Implementation of Viewport Independent Display Memory

We have devised a model in which we render the images of the virtual world onto a surface or surfaces surrounding the user. However it may appear that in so doing we have actually created a larger processing bottleneck, both in generating the images and rendering them, and in selecting the appropriate view to display. On first examination it may appear we have significantly greater rendering overheads such as scan conversion, clipping etc. than a conventional system, however this is rarely the case, and is only found if the scene has polygons evenly spread out in all three dimensional directions. With a viewport independent display memory one must scan convert all polygons received. This is the worst case scenario for a conventional system, but for guaranteed interactive response one must allow for the worst case. Many conventional rendering systems are designed to cope with situations approaching the worst case scenario [Molnar & Fuchs, 1990]. The rendering overheads for a conventional system may be reduced if the user is not looking at a complex part of the scene, however as the system has no control over the user's choice of direction for viewing, it is fair to assume the user is looking at the most polygonally dense section of the world. The viewport mapping is indeed computationally expensive however this task has been offloaded into relatively simple and cheap dedicated hardware, the *Address Recalculation Pipeline* [Regan & Pose, 1993].

The latency is determined by how long it takes to select the required portion of the encapsulating surface and write the image to the display device. This is handled by the Address Recalculation Pipeline which runs at the speed of the display device and introduces negligible latency. The new improved approach is depicted in Figure 2 and a comparison of the latency components with those of conventional systems is shown in Figure 3.

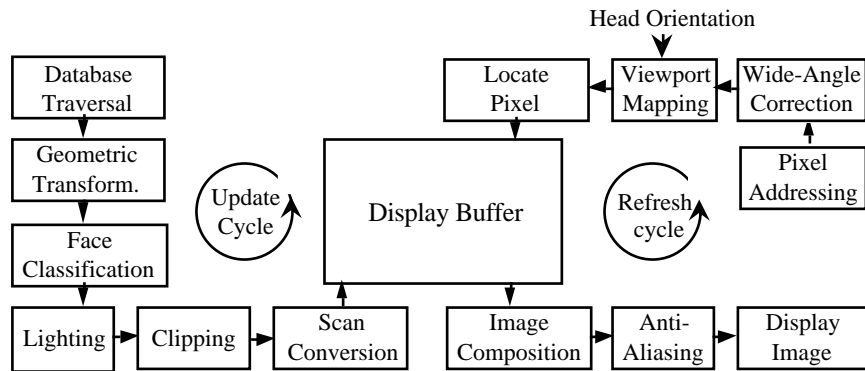


Fig. 2. Delayed Viewport Mapping.

The selection of the appropriate region of the sphere to display can be modelled as projecting rays from the eye onto the surface of the sphere, so defining the area to be seen. Alternatively one can view the process as one of changing coordinate systems from real-world to spherical world. This indeed requires a great deal of computation for each pixel displayed and would be infeasible in software. However the approach is ideal for a pipelined hardware implementation, which while using fairly fast hardware has a fairly simple structure and can be built economically [Regan & Pose, 1993].

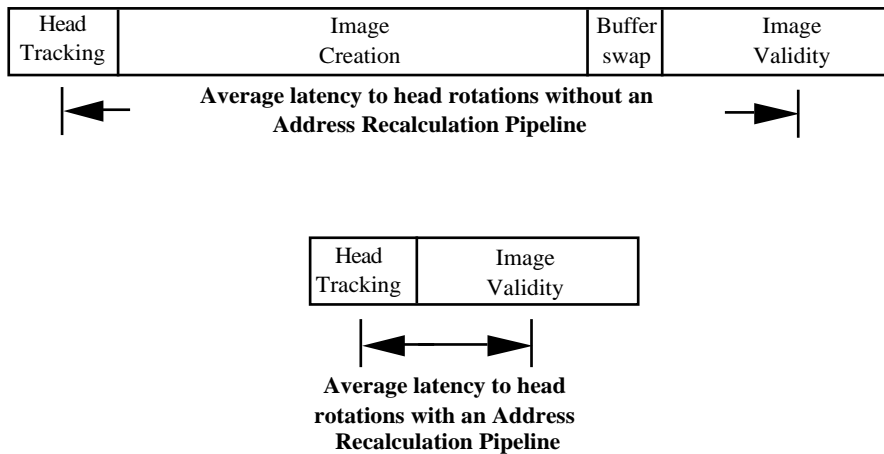


Fig. 3. Comparison of rotational latency components.

We did investigate a hardware design based on a spherical system but found some annoying characteristics. First, in mapping a sphere onto a memory array one tends to waste a lot of memory, or else have a very complex mapping scheme. Second, the apparent sizes of pixels in a sphere vary greatly as one moves from equator to the poles. Third, the coordinate transformations involve trigonometric functions which are moderately expensive to implement. By far the dominant issue is the first, in that we are using much more display memory than a conventional graphics system, and since we are not accessing it in a nice regular pattern, we are forced to use quite fast memory, hence we do not want to waste any of it.

Essentially the basic model will work for any surface surrounding the user. A sphere is intuitively obvious in that everything is equidistant, however one can go to the other extreme and look at a tetrahedron. What we eventually chose was a cube. It has very nice properties in that no memory is wasted since its surfaces are the same shape as memory chips. The coordinate transformations are also much simpler in that they are linear and can be implemented essentially with a matrix multiplication. Pixel sizes do not vary as much as for a sphere. While there may be some concern about strange effects occurring in corners and edges these have been shown not to be significant.

A prototype display memory system based on the cube has been implemented and functions well. We call this method using viewport independent rendering, *Delayed Viewport Mapping*, and the hardware realization, an *Address Recalculation Pipeline* [Regan & Pose, 1993].

6 Image Composition

A technique which has been shown to work well in accelerating computer graphics systems is image composition [Molnar, 1991]. Instead of treating the scene as a single entity, one can break it down into component objects and combine them at the end to produce the scene. One can render these component objects separately, and if one has multiple renderers available one can get parallelism in the rendering process. The big gain is that objects which do not change can be left, and not rendered again if you want to change the scene. Thus one can achieve better than linear speed-up. Unfortunately this does not help much in Virtual Reality applications since, as we have seen, a simple rotation will cause every object in a scene to change.




Fig. 4. Image Composition.

We can get the benefits of image composition by employing our viewport independent approach in Virtual Reality applications. Since the surface encapsulating the user does not change when the user rotates, the objects already rendered can remain, and once again we get better than linear speed-up with multiple renderers. Conceptually we have gone back to the *crystal sphere* model, and our image composition involves combining images of objects painted on various concentric cubes. We use the Z-values of pixels to enable us to choose the closest pixel to the user as the one to be displayed. Hence we do not force objects onto particular cubes in relation to their apparent depth, but instead allow the hardware to resolve dynamically which pixels to display. We also provide separate rendering engines for each cube. Thus it makes sense to distribute objects over cubes in such a way that those that need to be updated at similar times are grouped together. In this way the rendering capacity is not wasted on redundantly re-rendering objects which have not changed. A technique called *Prioritized Rendering* has been developed to handle display updates efficiently [Regan & Pose, 1994].

The technique of image composition is depicted in Figure 4. Image composition occurs as pixels are fetched from the display memories to be displayed on the output device. All of the pixels from the display memories which correspond to the same screen location are compared to find the closest pixel, which is then displayed.

7 Prioritized Rendering

Objects in the virtual world will appear to change at different rates. The apparent movement of the object due to a rotation of the user is handled automatically by the *Address Recalculation Pipeline* which implements *Delayed Viewport Mapping*. Other changes in the way the object is perceived are due to the object's own animation and due to translational movement of the user in the virtual world. It is possible to classify further the kinds of changes in the appearance of the object into those involving its displayed size and those involving its displayed position. Objects will appear larger when one is closer to them and smaller when one is farther away. Relative side-to-side movement will cause an apparent change in the position in which the object is displayed. Of course any animation of the object itself can lead to both these effects. The effects of user translation can be quantified using simple geometric analysis [Regan & Pose, 1994]. Using this approach one can calculate the amount of time for which the current displayed image of the object will remain valid, and hence determine the required rate of update of the object's displayed image.

Prioritized Rendering takes advantage of the fact that not all objects need to be updated at the same rate. By sorting the objects into a priority order based on their required update rates, one can render the most important changes in object appearance first, thus guaranteeing that if perchance there is insufficient rendering capacity, one has achieved the best possible result. Even in such an unfortunate case the latency of response to head orientation changes is not affected, since that is handled by the *Address Recalculation Pipeline*.

This however is not the main advantage of *Prioritized Rendering*. Its most dramatic property is that the overall rendering required has been reduced significantly. We have shown that an order of magnitude reduction in required rendering is not unreasonable due to this approach [Regan & Pose, 1994]. This occurs because the average required update rate for all objects in the virtual world is much less than the display update rate. By balancing the rendering load among multiple renderers updating multiple display memories at various update rates and composing the final image for display, the combination of *Prioritized Rendering*, image composition, and the *Delayed Viewport Mapping* provided by the *Address Recalculation Pipeline* can outperform conventional Virtual Reality implementations.

8 Implications for Building Virtual Reality Applications

Mechanisms have been described which enable the efficient implementation of Virtual Reality applications without the serious tradeoff between latency and image resolution which has plagued most existing systems. However the mechanisms will not necessarily be exploited to their fullest unless one designs the Virtual Reality applications with them in mind.

The conventional approach has been to run a tight loop involving the tracking equipment which monitors the user's position and orientation, the virtual world model, and the rendering of views of that model into the display device. What tended to suffer was frame rate, and hence latency, when the computer system could not keep up with the processing load. In such systems the frame rate would drop as the changes to the scene increased. More sophisticated systems would drop the accuracy of the virtual world model to minimize changes which require re-rendering. Other systems would reduce the accuracy of their rendering to try to maintain a minimum frame rate. All such systems degrade rapidly when stressed by changes in the view of the virtual environment.

In our system, latency to rotation is independent of the virtual world; it is a function of our hardware and is effectively limited by the tracking equipment. Translational latency is reduced indirectly in that distant objects in the background will not be observed to move and need not be re-rendered. Changes in the environment itself have still to be handled, but by grouping them according to their observed speeds one can minimize the number of cubes which need to be updated. In the same way, changes due to translations and due to preserving stereoscopy can be grouped according to the required update rates, and allocated to the display memories appropriately.

We do not claim that our system can display perfect renditions of a fast changing view of a virtual world without latency. Certainly the extremely important and difficult case of rotational latency is effectively solved, however translations and actual virtual environmental changes can conceivably overload the system. The important thing to note is that such problems would result in out of date representations of less important parts of the virtual world as seen by the user. It would not affect the smooth motion by lowering the frame rate and would not let the image quality suffer. Such an overload condition can only occur during very rapid translation or when there are fast moving objects in the virtual environment. In such cases the ability of the human eye to perceive much detail is limited since fast moving objects tend to blur. Hence it is unlikely that such overload would even be noticed by the users unless the system was grossly under powered.

In order to exploit the system one must change the paradigm used in conventional Virtual Reality systems. One must, in effect, ignore the rendering and instead concentrate on a much higher level, that of objects in the virtual environment. Mechanisms are already in place to handle rendering of objects in viewport independent display memories, but the management of these objects and the allocation of them to particular display memories is very important. One must arrange for objects which are the focus of attention to be rendered on time at all costs. Similarly one must identify background features and fairly static objects and group them together as objects that rarely need to be re-rendered.

Indeed one can consider the required approach to be somewhat object-oriented, and one should try to concentrate on gross object behaviour rather than fine object

detail. Given enough processing power the detail will be handled. In those cases where processing power is a limiting factor, it is likely that the view is changing so rapidly that fine details would not be missed anyway.

9 Computer Architectural Difficulties which were Overcome

The system as described thus far takes the form of a pipeline, which we term the Address Recalculation Pipeline. In it a rotation matrix derived from tracking or other devices is used together with data describing the optical characteristics of the display system, and of course the actual image data, to re-orient the images viewed by the user, in real time. This involves a matrix multiplication as each pixel is displayed on the screen, that is millions of matrix multiplications per second. This as has been shown is quite manageable, but one must consider the enormous quantities of data flowing through the system. Given that the overall structure is that of a pipeline, much of the data traffic only flows from one stage to the next, so is not too difficult to handle, however there is an aspect which did require an innovative architectural solution.

Because the images inside the system do not correspond directly to any images seen by the user, it is impossible to do the usual anti-aliasing operations which help minimize the *jaggies* and other effects caused by the discrete sampling nature of a raster graphics system. Thus we cannot rely on the image rendering software to do the anti-aliasing. Instead we need to run an anti-aliasing filter as the images are being displayed, in real-time. Even a fairly simple anti-aliasing filter, such as bilinear interpolation filtering, requires four adjacent pixels to be averaged to create a pixel value to be displayed. In our system we do not even know which pixel is going to be displayed until just before it appears on the screen, hence we needed to find a way to simultaneously get the four adjacent pixels from our display memory.

An ingenious solution was found which involves dividing the display memory into six independent segments [Regan and Pose 94A], and allocating pixels to the segments such that it is guaranteed that the four pixels required for the bilinear interpolation are each in different segments. It then remains to choose which four of six pixel values are required. Each segment is of course assigned to different memory banks so that they may all be accessed simultaneously. The technique may be likened in some ways to the way one can colour a map with only four colours, although six were more convenient for our topology.

10 Telerobotics Applications for the Technology

One of the major problems with telerobotic operation of remote devices is the lag in the operator's view of the remote site. This causes over-reaction to remote stimulus since the operator causes some action, doesn't immediately see the effect so increases

or repeats the operation, until when the updated images do appear, they show a misleading view. The problem of over-reaction and over-compensation is well known, and forces operators to proceed very cautiously, and makes it rather difficult to learn how to operate telerobotic machinery.

The solution, once again, is to think back to the underlying problem. In a sense, the problem is not just the lag, but really the temporal mismatch between the user initiating some action and his seeing its effect. Most of this difficulty is due to delays in the communication link which carries the video images. In real life we get visual feedback at the speed of light from a very small distance, hence the image we see appears instantaneous with our actions. If we put a communications link in the way, especially one going via satellite, then we incur a significant delay with its attendant problems.

However let us consider carefully the domain in which we are operating. At the remote site is a robot. Most current robotic devices comprise an arm rotating about various axes with some camera and implements attached. Hence the main change in viewpoint is one of orientation. We have a device which can change effective viewpoint orientation in real time. We also have a device generating the images which is controlled very precisely although with some apparent lag. Hence we know fairly precisely the orientation the camera should be in, even before the video images have been received reflecting the view from that orientation. Thus given a suitable wide-field camera and the orientation we assume the robot has attained, we can generate using the Address Recalculation Pipeline, an image as it will appear, hence avoiding some of the lag which causes our problems [Regan and Pose 94B]. Of course when the actual image arrives we can correct any minor errors that crept in.

Hence we have applied our new technology to a different application domain, using it in a slightly different way, but with great benefit.

11 Interactive Television Applications

We can generalize the approach used for telerobotic applications to the ubiquitous home television. While, in a way, the easiest of the application areas, it is perhaps the one with enormous market potential.

A conventional television receiver displays exactly what is broadcast from the television station. Let us pretend that the studio TV camera is equipped with a fish-eye lens which gives say a 180 degree field of view. On the conventional TV channel images corrected via an Address Recalculation Pipeline are transmitted and appear normally on conventional television receivers. On another channel the uncorrected fish-eye view is transmitted. Those viewers with television receivers equipped with Address Recalculation Pipelines will have this distorted image mapped straight into their display memory and can use their television set's Address Recalculation Pipeline to navigate around the 180 degree scene. Of course it could be extended to an even wider field of view.

One can reduce bandwidth requirements by having the new television receivers use both transmitted channels avoiding having to transmit redundant information. This could add a new dimension to television viewing and appears far more interesting than current video-on-demand proposals. It would also appear to be a candidate for using the excess cable television bandwidth currently being installed in Australia, as well as giving people an incentive to connect to a cable television network.

12 Influence on Future Virtual Reality Applications

The technology of the *Address Recalculation Pipeline* employing the method of *Delayed Viewport Mapping* combined with *Prioritized Rendering* is clearly superior to existing Virtual Reality implementation techniques. It is also much more cost effective than existing systems using extremely powerful and expensive computer graphics equipment. With its clear advantages in latency, reduced rendering costs and graceful overload management, we expect that these techniques will become dominant for Virtual Reality display systems.

With the improved performance available using these techniques it is expected that future Virtual Reality applications will be optimized for this technology. Furthermore, the virtual worlds created for such systems will be characterized by complex scenes with highly realistic backgrounds and large numbers of independent objects. Conventional Virtual Reality systems suffer from unacceptable performance degradation with this kind of virtual world.

The technology also is directly applicable to telerobotic applications such as those found in controlling remote robotic manipulators in space or under sea. It is also ideal for Augmented Reality applications where virtual world images are superimposed on the real world. There is also much scope for our techniques to be used in future interactive television systems, and elsewhere in the entertainment industry.

13 Current State of the Project

The project has gone from the conceptual stage through to an initial prototype ARP implementation in the Computer Science Department of Monash University. A company has been formed to commercialize this exciting technology and patenting is underway. It is expected that further research into this and related technologies will continue.

Conclusion

By changing dramatically the way one views the virtual world, one can take advantage of techniques such as *Delayed Viewport Mapping*, image composition and *Prioritized Rendering* to increase the performance of Virtual Reality computer systems. A novel but affordable hardware architecture, *Address Recalculation Pipeline*, was employed to implement an extremely low latency Virtual Reality display system. In itself this can eliminate rotational latency, however to exploit fully the technology, the structure of Virtual Reality applications needs to be changed.

In this paper an outline of the new display architecture was presented. This set the scene for an outline of how to structure Virtual Reality applications so as to take advantage of this architecture, and the method of *Prioritized Rendering* which can exploit it effectively. The techniques described can enable Virtual Reality applications to be run on economical hardware, and in fact outperform very expensive conventional graphics computers.

The extension of the architecture to domains such as telerobotics and interactive television show the flexibility and generality of these techniques. Prototype Address Recalculation hardware has been constructed and the technology described herein is being patented around the world.

Acknowledgement

The prototype hardware was funded by an Australian Research Council small research grant. Patenting and commercialization of this technology is being supported by a Monash Reserach Fund grant.

References

- Fuchs, Henry et al. (1989) Pixel-Planes 5: A Heterogeneous Multiprocessor Graphics System Using Processor Enhanced Memories. Proceedings of SIGGRAPH 89. In *Computer Graphics*, Annual Conference Series 1989. pp. 79-88.
- Hettinger, Lawrence and Riccio, Gary (1991) Visually Induced Motion Sickness in Virtual Reality Systems: Implications for training and Mission Rehearsal, Presented at a DoD sponsored Interagency Tech Simulation, 1-3 Oct, 1991.
- Molnar, Steven and Fuchs, Henry. (1990) Advanced Raster Graphics Architectures. Chapter 18, *Computer Graphics*, Foley and van Dam, pp. 872-873.
- Molnar, Steven, (1991) Image Composition Architectures for Real-Time image Generation. Ph.D. dissertation, University of North Carolina, 1991.
- Regan, Matthew and Pose, Ronald (1993) An Interactive Graphics Display Architecture. Proceedings of IEEE Virtual Reality Annual International Symposium. (18-22 September 1993, Seattle USA), pp. 293-299.
- Regan, Matthew and Pose, Ronald (1994) Priority Rendering with a Virtual Reality Address Recalculation Pipeline. Proceedings of SIGGRAPH 94. In *Computer Graphics*, Annual Conference Series 1994. pp. 155-162.

- Regan, Matthew and Pose, Ronald (1994A) Display Memory Access Issues and Anti-Aliasing with a Virtual Reality Address Recalculation Pipeline. Proceedings of the 9th Eurographics Workshop on Graphics Hardware, SINTEF, Oslo, Norway, Sept. 12-13, 1994, pp. 23-27.
- Regan, Matthew and Pose, Ronald (1994B) Virtual Reality and Telerobotics Applications of an Address Recalculation Pipeline. ISMCR'94: Topical Workshop on Virtual Reality, Proceedings of the Fourth International Symposium on Measurement and Control in Robotics, Houston, Texas, USA, Nov. 30-Dec. 3, 1994, pp. 31-36.