2D/2.5D/3D approaches to generating three-dimensional scenes with expedient solutions/techniques

Christopher Whyte

March, 2009

Abstract

To maximize the efficiency of a computer graphics project, the manner in we approach the creation of its environment must be considered. The purpose of my Innovations project is to provide information about the expedience of a few possible solutions to this problem. It is my intention to aid the producer in selecting the most suitable approach for the creation of their environment, taking into consideration the requirements and constraints of their project.

Introduction

The purpose of this project is to investigate different solutions to CG environment creation and evaluate their expedience. I will compare the selected approaches on the grounds of quality and speed of execution.

Different productions require different degrees of camera movement. Whilst a certain method of environment creation may quickly produce an acceptable or perfect quality result for a production with minimal camera movement, its quality may be unacceptable for another production with a wider degree of camera motion.

Therefore, to aid the producer, I intend to categorise my results based on how the solutions stand up to three varying degrees of camera motion (none to minimal, moderate and heavy). Through experimenting with various camera moves it will be possible to evaluate the quality and speed of execution of each approach for each category.

The solutions I explore in this project had either been proved a success on a time saving basis or had the potential to be. I attempted to not only determine the speed of execution of these approaches but to also push the limits of each approach in order to heighten the quality and the range of camera motion available when utilizing them.

I aim to produce a table of results detailing the expedience of each approach according to the range of camera movement required. It is also my objective to submit tests of these approaches to make clear the limit of their use as a basis for environment creation.

Environment Approach 1 – "Theatre Set"

This approach is for front view and very minimal angle cameras, almost like it's a matte painting technique. The environment is created as a collaged set as if it was physically a model of a theatre set, which is where the people who pioneered the solution, Nexus, got the idea from initially.

This technique is illustrated many times throughout Nexus' commercial for Ambi-Pur, created by Jim Lefevre. Props and scenery are extensively photographed, and a forced perspective 3D room (where the far wall is smaller than the front one) is built. Alpha channels are then created for the photographs. Once the images are matted onto planes and the alphas applied to their opacity channels the planes appear to be cut to the shape of the image.

Using this approach I took tons of photographs of the walls of a room in my house and the "props" inside it. It was at this point where little things came into play and doing some research into theatre sets provided me with some good ideas. For instance, if you were making one in real life you would soon realise that it is helpful to take the photographs that you'll be making the set from at a front on slightly low-ish angle, even if your camera is going to be at a different angle. This is because it's more important that the set objects connect with the floor than anything else.

In Maya I created 'planes' that are the same size as my source photos, and dragged their anchor/pivot point to the lowest side and then placed and angled them to the camera. After that it was a case of introducing some lights to the scene and getting the shadows right. I also tended to push the pictures that were making the set so that the lighting felt good, so I painted extra shadows on some of them, again with the theatrical set in mind.

A few important issues were identified in execution of this approach. These included lighting problems with the ceiling – Tom Dawson recommended editing the ceiling connection with the light source in the connection editor. The second issue was the shadows of the set's "props" were plane shaped instead of shaped like the object. With these issues resolved I began to testing the environment with varying degrees of camera moves. I found that the approach works well for pans along a front view of the scene as well as dollying in and out. However, introducing anything over minimal camera movements destroyed the illusion of the approach. The tests in the video file 'Approach 1 - Theatre Set Tests.m2v' illustrate the range of successful camera motion and also examples of where the technique begins to break.

In conclusion, for animations which require no or minimal camera moves the quality of the approach is perfect. However for moderate and heavy camera moves the quality is unacceptable. As for the speed of execution, setting up the forced perspective scene in 3D is simple and fairly quick. However, taking into consideration the time spent collecting photographs, creating their alpha channels in Photoshop, and lighting the scene, I would say that although this is far quicker than modeling and texturing a fully 3D scene, it is not the quickest form of environment creation.

Environment Approach 2 – Camera Projection

Camera Mapping is an extremely effective technique that is used extensively in movies to turn still photos or matte paintings into realistic 2.5D environments. Today matte paintings can be more than just static flat paintings in the background. They can now be projected onto 3D geometry which allows for much more complex camera moves and perspective changes.

In a 3D application such as Maya we can use one camera projection to "glue" an image to some matching geometry as it is projected from the camera at a specific angle. We can then animate using a separate camera.

The concept of camera projection is older than CG itself, and its principles can be found in many real world examples. The technique was used widely in the theater world for decades. When the photograph matches the foreground objects, and when the projection can be made to "stick" to the objects then camera projection can be used as the basis for creating view-dependent texture maps.

Another real world example is the sidewalk chalk paintings of Kurt Wenner and Julian Beever where the chalk painting, the pavement and the architectural surroundings all become part of an illusion. This form of perspective is known as anamorphism which was used by the great European Masters to give the illusion of soaring architecture and floating figures in ceiling frescoes.

In anamorphic perspective, painted forms appear as three-dimensional when viewed from one point in space. Wenner adjusted this geometry used by the great European Masters to create compositions that seemed to rise from and fall into the ground.

An example of Kurt Wenner's anamorphic street painting

The system of anamorphic projection can also be seen on text written at a very flat angle on roads which then becomes much easier for drivers to read as they approach the text; when the vehicle is nearly above the text, its true abnormally elongated shape can be seen. Similarly, in many sporting stadiums, the same technique is used to promote brand names. The adverts are painted onto the playing field in a manner so that when viewed from the angle of the camera, the writing appears to be standing vertically.

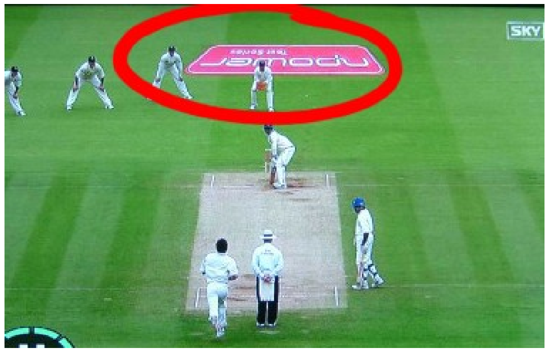

An example of anamorphic perspective used in advertising

In anamorphic perspective, painted forms appear as three-dimensional when viewed from one point in space. The technique of camera mapping was intriguing to me as it presents a way around this limitation (although to what extent I was not sure). Sticking to my objectives, it was my aim to see how just far this technique could be pushed in terms of quality and degree camera motion.

Another reason for selecting camera mapping was that not only is its quality high enough for big budget motion pictures, but it is also an economical production technique that can be used in just about any film, no matter how small the budget. For example, suppose you were producing a film and the director requires an establishing shot of a street full of buildings and the only available footage is of the insides of the buildings so it has to be done using computer graphics. One option would be to hire a camera crew – a very expensive and complicated situation even for something as simple as this. Other options would be to build a miniature of the scene or model the scene in Maya. However these are also very expensive processes. This is the type of situation in which camera mapping can save a lot of money because all that is needed for implementation is a single photo of the environment.

The process usually involves building simple 3D geometry, picking a viewpoint and rendering what you see. Textures are then painted for that camera's perspective in Photoshop, then the image is taken back into Maya, and projected back onto the corresponding geometry.

In an attempt to improve the expedience of this technique I decided to work in reverse order. This involved simply taking a photograph of a scene, projecting that photograph from a camera in Maya, then building geometry that lines up or 'catches' the projected image. Then a second camera is created and used to animate with.

I created a simple scene containing just a book, a stack of DVD cases and a box. This allowed me to assess the merits of the technique without having to do one or more time consuming matte paintings.

Figure 1. Scene with no geometry to match an object in the projection

I encountered a few issues when using this technique. For instance, when the photograph contains an object and the scene contains no matching 3D geometry, (as it does in the Figure 1) the object can appear normal from one angle, and distorted or cast like a shadow when viewed from another place.

I came across this problem many times (see Figue2 and Figure 3). The most time consuming part of this process is taken to avoid this problem, by accurately matching your camera's angle and position with the perspective lines in your photo and moving, rotating and scaling the separate objects to make them match the photo.

Figure 2. No geometry to match projection

Figure 3. Geometry matching projection

The beauty of using Camera Mapping is that all of the lighting information is already in the photograph being used, as is the specular, diffuse, bump and nearly every other attribute we would otherwise have to consider. This factor is a huge time-saving aspect of this technique.

Another advantage of this technique is that if a change to the lighting is required it is possible to add extra lighting and shadows to the scene effect using your 3D application. This provides more flexibility and I discovered in my tests that this can add significantly to the over all look of a scene as well as helping to anchor objects to the surrounding environment.

Therefore, not only is this an extremely quick method of environment creation, but a high standard of quality can be achieved with little effort. Additionally, if a matte painting is projected instead of a photograph, it can be of any style you desire and the lighting information can be edited solely in that image (though admittedly, this would slow down the speed of the approach).

To find the breaking point of the approach I tested my scene with various camera movements. For the most part, with well placed geometry matching the projection correctly, moderate camera moves were achieved. See the video file 'Approach 2 – Camera Projection Tests.m2v'. The illusion was shattered as soon as the camera movement revealed anything not visible to the camera in the projected photograph, for example the back side of the red book. An example of this is also shown.

Contemporary practice predominantly employs single point camera projection. However, at this stage I decided to explore two point camera projection (and beyond if necessary) to see how fluidly the camera might move around a scene without breaking the illusion.

Several articles on computer graphics have mentioned one particular shot in Star War: Episode 1: The Phantom Menace, in which the artist used several photographs of a model taken from different positions and projected them from multiple cameras onto roughly modelled geometry. The camera which was actually rendering the scene flew through these simple objects that were being texture mapped from the different cameras to create a fly though of the entire model.

I began to explore using a second camera projection to simply cover up the smeared textures from the example of unsuccessful camera motion. In Maya I rendered out a still from the perspective that was showing the problems and corrected all the problem areas in Photoshop using a photograph taken from roughly the same angle. After that it was a case of making a new material with the new painting, setting up a second projection camera from this perspective and assigning the new material to those models that had shown the problem. To avoid seams I painted an alpha maps to blend between the two paintings and only make the second painting show up in specific areas.

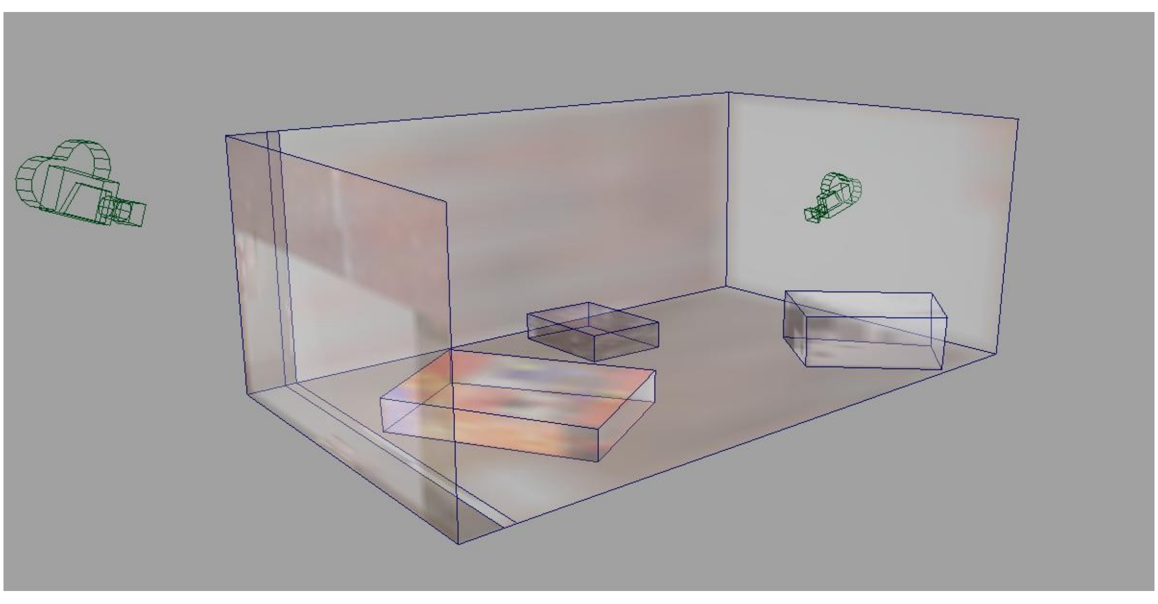

Figure 3. Approach 2 Projection Receiving Geometry

After employing this technique I rendered a few heavy camera moves of the scene. With the added projection from the second camera, the scene allows for a significantly greater degree of camera motion. Also, the quality was further improved by the addition of CG lighting to the scene. With further projections, I am sure that the scene would allow for fully comprehensive camera motion.

Finally, I added lighting to the scene so that the analysis can focus purely on the effectiveness of 2/2.5D techniques. The fake lighting and shadows serve as something subtle to anchor the objects to the surface they lie upon. The result can be seen on the video file 'Approach 2- Camera Projection tests.m2v'.

It was relatively quick to add in the second projection. For any additional projections the projected image is already lined up so that it matches the geometry. Thus, the most time consuming part is taken away. The most time is spent painting over problem areas in Photoshop which is relatively simple given photographs from a roughly similar angle.

If a more stylistic or imaginative look is required, matte paintings might be preferred over photographs. Although this would slow the speed of the technique, it would remain a much faster solution than modeling and texturing a fully 3D scene.

Overall, the quality of the approach is high even for heavy degrees of camera motion. The speed of execution of the approach was fast, especially in comparison to modeling, texturing and lighting a matching scene. However, the heavier the camera motion required the more this time is increased due to the requirement of additional projections.

Approach 3 - Full 3D Environments

It could be argued that in a project such as this it is imperative that I also produce a fully modeled and textured 3D environment; especially for the purpose of comparing its expediency with that of the alternative methods. Although this would have been useful, I decided only to create such an environment using this technique if time allowed towards the end of the project.

It is already established in Computer Graphics that not only is this approach the most flexible and produces the highest quality environment, but it is also the slowest to execute. Therefore, I felt that time would be better spent experimenting with, and pushing the limits of faster and lesser established approaches to provide a more interesting and useful report. For this reason I was not fearful of being left without sufficient time to incorporate it into my artefact.

The main objective of this project is to investigate effective methods to avoid having to use this time consuming solution. Granted, for bigger companies with significant man power this may not be an issue. However, it is my intention for the project to be of great use to degree level students and employees of lower budget productions looking for ways to avoid this expensive and time consuming approach.

Approach 4 - Exploration Into Texturing Approaches

Bump Mapping

Within this project I have explored various texturing techniques to see if they can be extended to be used as a basis for environment creation. Ultimately this was done with the aim of discovering different and more expedient solutions to this problem. However I have also been able to improve my knowledge of texturing through this research.

I decided to explore texturing techniques for my project as they are a useful way of saving time, not only when adding detail to your scene, but also saving on time during computation. Therefore, if I was able to harness an approach for environment creation using texturing techniques as a vehicle, it would most likely be a much faster process compared to increasing the polygon count in a scene. The reason for this is that texturing only adds linear interpolation over a polygon and one look-up for each part of the texture, both of which can be done very efficiently in hardware. Whereas if the polygon count is increased by modelling, each new vertex must pass through all three transformations, as well as clipping and lighting.

Bump mapping is a technique that attempts to make a surface look more realistic by modelling the interaction of a bumpy surface texture with the lights in the scene. This is achieved by altering the brightness of the pixels on the surface according to a specified height map.

Bump mapping is known to be good for small and inward-facing bumps (like those on a golf ball or an orange). From the first test (see Figures 4 and 5) I discovered that this technique is also effective in adding an illusion of three dimensionality to a brick wall.

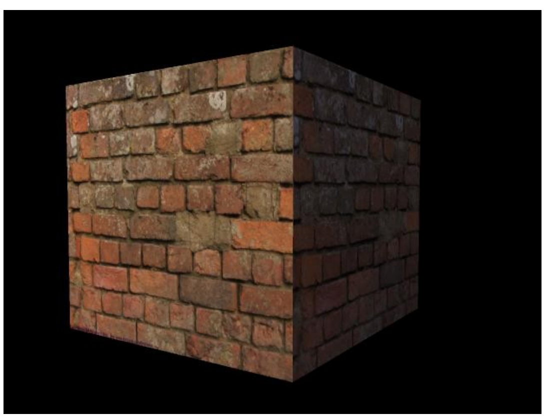

Figure 4. Scene without bump mapping

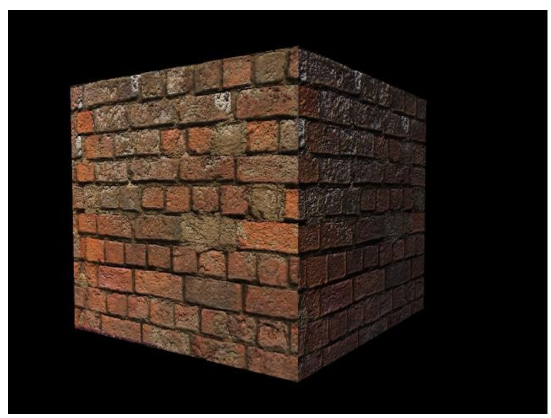

Figure 5. Scene with bump mapping applied

Figure 6. Scene without bump mapping

Figure 7. Scene with bump mapping applied

The second test (see Figures 6 and 7) was relatively quick to set up. The process was far less time consuming than modelling the detailed bricks and windows of the same scene to the nth degree. It revealed that the effect on geometry such as protruding ledges and archways is not desirable and consequently they would have to be modeled.
Another downside to bump mapping is that only the effect of lighting on a surface is changed, not the shape of the surface. Therefore, bump-mapped objects generally look realistically bumpy, except that they have smooth silhouettes. This attribute reveals the illusion and causes the technique to lack significant realism.

I have discovered that despite working well for quickly applying an illusion of minor detail to certain surfaces, bump mapping is not useful in terms of establishing a basis for environment creation.

Displacement Mapping

Displacement mapping is a technique commonly used for adding surface details to an object. A height map is used to cause an effect where the actual geometric position of points over the textured surface are displaced according to the value that the texture function evaluates to at each point on the surface.

I was particularly intrigued by displacement mapping for this reason. It seemed to me that altering geometry using a single image could be utilized as a quicker alternative to environment modelling. Additionally, having an actual three dimensional model as a result would allow for comprehensive camera moves. After observing the effects on displacement mapping from many different sources, I started to wonder how far this technique could be pushed in order to save time when modelling an environment

I decided to see whether displacement mapping could be used to tackle the windows and protruding ledges in my previous test. Subsequently I wasbe able to bring out the edges of the ledges in by modifying the map in Photoshop. With bump mapping then applied to the brick walls, perhaps a combination of texturing techniques would reveal an expedient method of environment creation.

The map I applied to the wall was a grey scale image I had created from the colour map texture of the same side of the building. The raised ledges were defined as white amongst a grey plane. The result was poor. It seemed as though the geometry did not have a high enough resolution to convert the detail of the map. After upping the resolution it started to look better but the scene became heavy and slow. This indicated that any further increases would render the technique impractical for creating even a small environment.

Figure 8. Displacement Mapping Experiment

Seemingly the protruding edges in straight lines (and thus correspondent with the geometry) achieved a far better result than the displacement of the archways but were ruined when they came to the detailed corners (see images above). From a front view they also lacked realistically sharp deformations.

From this test I can conclude that displacement mapping is not advantageous as a basis for environment creation. This technique is able to give surfaces a great sense of depth and detail, meaning it is perfect for making terrains, mountains and even bathroom tiles. However this test alone proved that the technique was worthless in terms of being used as a basis for creating a 3D scene seeing as the only way to improve the effect of this technique is to increase the resolution of the geometry (as discussed above), which would render the scene much too heavy to be manageable.

If an approach to environment creation using texturing as a basis could be harnessed it would be much faster process, However it seems from my tests that texturing, despite having its uses, is not yet advanced enough for this. Displacement mapping seemed by far the most promising but still uses a lot of power at the moment, and it is unlikely to be a useful feature for some time. In my opinion a useful report of this type should not only highlight the techniques which it has found to be beneficial but also warn of which ones to avoid. These approaches are not included in the results.

Table of Results

Camera Motion	Approach 1 'Theatre Set'	Approach 2 Camera Mapping	Approach 3 Full 3D
None	Quality: Perfect Speed: Moderate	Quality: High Speed: Fast	Quality: Perfect Speed: Slow
Moderate	Quality: Unacceptable Speed: Moderate	Quality: High Speed: Fast/Medium	Quality: perfect Speed: Slow
Heavy	Quality: Unacceptable Speed: Moderate	Quality: High Speed: Medium	Quality: Perfect Speed: Slow

Conclusion

One of the challenges of my project was to determine the speed of execution of each approach for comparison. For the results to be of more use to the industry these would have to be according to industry standards. Unfortunately it was not possible to accurately provide this information seeing as I was learning the techniques for the first time. However I am sure that the findings of the project are useful to the industry as a rough guide.

I was able to assess the rough time taken for someone at my level of knowledge to complete the tasks. Thus, I feel that the findings of this project also serve as a useful guide to those with a level of specialist subject knowledge similar to that of degree level students. The project has also served well for improving my knowledge on skilled areas such as texturing and camera mapping and lighting.

The most challenging aspect of the project was to push the limits of each approach in order to heighten their quality and the range of camera motion available. I feel that I did not deal with this aspect of the project particularly well. For instance, regarding the 'theatre set' approach, although a range of basic camera moves were achieved, I feel that I could have pushed the merits of the technique further. For instance, I could have experimented more with the idea of forced perspective and experimenting with the positioning and shape of the 2D planes to trick the viewer. I could have also tried roughly modelling the geometry more to see how far the camera might move around a scene without breaking the illusion.

In moderation I feel that I succeeded in pushing the limits of camera mapping by using two camera projections to heighten the degree of camera motion available, especially seeing that in contemporary practice single point camera projection is most commonly used. However, I could have pushed this idea much further by including many more camera projections in my tests. This is a promising aspect of the project which I intend to pursue in the future.

In my opinion the project would have been more useful if many more methods of environment creation had been investigated, but unfortunately due to time constraints this was not possible. The solutions I selected were ones that had either been proved a success on a time saving basis or had the potential to be. Therefore, I felt that I had chosem the best possible methods to experiment with in the given time.

It was also my objective to submit tests of these approaches that clearly express the limit of their use. I feel that the tests in the videos submitted convey this aspect very well. It is possible to see exactly what degree of camera motion is available when using the techniques.

In conclusion I am happy with the outcome of the project. Although it is not a complete guide, it provides useful information to any body looking for assistance in choosing an expedient solution for the creation of 3D scenes. It also provides a detailed analysis of the merits and pitfalls of both successful and unsuccessful approaches and provides visual confirmation to back up its conclusions.

I have learnt a great deal through completing this project. There is still much more to be under-stood, but that can only come about through further experimentation and development of my skills.

The wealth of knowledge I have accumulated through completing this project will be invaluable in the future. I will be able to approach the creation of a scene in many different approaches and a much wider base of knowledge.

References

Books

ARONSON Arnold (1990): American Set Design. Theatre Communications Group Inc.
HOLT Michael (1988): Stage Design and Properties. Phaidon Press Ltd.
HORNE Gary (1999): Stage Design: A Practical Guide. The Crowood Press Ltd
MEADE Tom & ARIMA Shinsaku (2004): Maya 6: The Complete Reference. McGraw Hill Professional.
DOBBERT Timothy (2005): Matchmoving: The Invisible Art of Camera Tracking. John Wiley & Sons.

Internet

http:://www.nexusproductions.com
http://qvolabs.com/otis_cam_proj_mp_01.html
http:://www.artbydaniel.co.uk
http://www.3dtotal.com/
http://www.digitalartform.com/archives/2004/11/camera_projecti_1.html
http:://www.renderosity.com/news.php?viewStory=13286 - 64k

Acknowledgements

I would like to thank:
Ben Cowell and Jim LeFevre at Nexus, for their support, Phill Allen, for his guidance and ideas, my housemates and Harriet for all the motivation I could ever need.