

Cosmic Cookery: Making a Stereoscopic 3D Animated Movie

Nick Holliman^a, Carlton Baugh^b, Carlos Frenk^b, Adrian Jenkins^b, Barbara Froner^a, Djamel Hassaine^a, John Helly^b, Nigel Metcalfe^b, Takashi Okamoto^b

^aDepartment of Computer Science; ^bDepartment of Physics,
University of Durham, United Kingdom

ABSTRACT

This paper describes our experience making a short stereoscopic movie visualizing the development of structure in the universe during the 13.7 billion years from the Big Bang to the present day. Aimed at a general audience for the Royal Society's 2005 Summer Science Exhibition, the movie illustrates how the latest cosmological theories based on dark matter and dark energy are capable of producing structures as complex as spiral galaxies and allows the viewer to directly compare observations from the real universe with theoretical results. 3D is an inherent feature of the cosmology data sets and stereoscopic visualization provides a natural way to present the images to the viewer, in addition to allowing researchers to visualize these vast, complex data sets.

The presentation of the movie used passive, linearly polarized projection onto a 2m wide screen but it was also required to playback on a Sharp RD3D display and in anaglyph projection at venues without dedicated stereoscopic display equipment. Additionally lenticular prints were made from key images in the movie. We discuss the following technical challenges during the stereoscopic production process; 1) Controlling the depth presentation, 2) Editing the stereoscopic sequences, 3) Generating compressed movies in display specific formats.

We conclude that the generation of high quality stereoscopic movie content using desktop tools and equipment is feasible. This does require careful quality control and manual intervention but we believe these overheads are worthwhile when presenting inherently 3D data as the result is significantly increased impact and better understanding of complex 3D scenes.

Keywords: Stereoscopic Animation, Scientific Visualization, Cosmology, Astronomy, 3D Video Editing, 3D Video Coding, 3D Display, Human Factors, Rendering

1. INTRODUCTION

Our understanding of the development of structure in the universe is increasing through a combination of improved observations, theoretical modeling and simulation. These methods produce huge data sets that on their own are difficult to present to general audiences. The aim of producing the Cosmic Cookery movie was to present this information to a general audience in a comprehensible way using advanced 3D visualization.

The requirements for the Cosmic Cookery exhibit were to create stereoscopic video and multiscopic images for large scale stereoscopic projection, medium sized multiscopic lenticular prints and small scale auto-stereoscopic display. This became a significant test of the level of support for stereoscopic image creation in rendering software, editing tools and compression standards.

As we have recently produced new stereo image rendering tools based on novel algorithms first reported at this conference¹⁻³ we were interested to compare these with traditional approaches to stereo image generation using verging and parallel cameras. The wide range of rendering software we needed required all these methods to be adopted and enabled a practical comparison of the different approaches.

In the following section we introduce technical background regarding the target displays and the stereoscopic camera control methods adopted. Details relating to the cosmological simulations and astronomical observations we are visualizing are presented when we describe the production of the individual sections of the movie.

Further author information: Send correspondence to Nick Holliman.

Email: n.s.holliman@durham.ac.uk, Telephone: +44 191 334 4287, Web: <http://www.durham.ac.uk/n.s.holliman/>

2. BACKGROUND

2.1. Displays

The choice of 3D display technologies was driven in part by availability of the displays but also by availability of the software tools to render, edit and distribute the visualizations in the required formats. While large-scale auto-multiscopic displays would have been an attractive option for public presentation their current scarcity and lack of software support ruled this option out.

The primary target display for the movie was an Inition Duality S2 system, a portable 2m wide stereoscopic projection screen using twin projectors and linear polarization. This was driven at a stereo resolution of 1024x768x2 from standard PC hardware using a dual channel nVidia Quadro card.⁴

A second, interactive, display was placed on a pedestal allowing sections of the movie to be played back individually from a web page interface. To support this feature a Sharp LL151D auto-stereoscopic display was used. This was driven with pre-interlaced images at a stereo resolution of 512x768x2 using a single channel from an nVidia Quadro card.⁴

Finally two large 1mx0.9m auto-multiscopic lenticular prints were used to show 3D views from the finale of the movie. Each print required the generation of 27 source images, each individual view at a resolution of approximately 1024x768, giving a multiscopic resolution of 1024x768x27.

2.2. Stereo Camera Settings

One of the most discussed issues regarding stereoscopic image capture is the stereo camera settings required, specifically what camera separation to use. Two common approaches are the verging camera model and the parallel camera model; these are widely recommended in stereoscopic graphics and photography books.⁵




Figure 1. Stereoscopic camera models; left diagram shows a verging camera arrangement; right diagram shows a parallel camera arrangement.

The verging camera arrangement is shown in Figure 1. This is often a preferred approach as it is easy to setup and requires little post-processing of images for synthetic and photographic cameras using symmetrical frustum. If the cameras are rotated to verge at the zero disparity plane Z' then objects in the scene in this plane will appear to the viewer to be in the plane of the display. However, this rotation generates opposite keystone distortion⁵ in the left and right images which in turn produces false vertical disparities in the stereo pair. The result is images that are subjectively harder to view and quantitatively distort stereoscopic space.⁶

A parallel camera arrangement, also shown in Figure 1, avoids the problem of keystone distortion by capturing the left and right images using coplanar image(film) planes. To ensure infinity is behind the screen surface the camera frustum should be asymmetrical such that the cameras have coincident field width at the zero disparity plane Z' . This can be achieved in systems without asymmetric frustum by capturing a symmetric view and cropping to the required region of the image. Using the parallel camera model avoids vertical disparity but, in

common with the verging camera model, it does relatively little to help the image producer control total apparent depth. The result is for both approaches there is a need to engage in repeated trial and error, adjusting camera separation until a comfortable image is created.

A confusion regarding the utility of verging vs parallel camera setup is often brought about as the eyes verge when viewing a scene. However, the aim of image generation for most 3D displays is not to generate images for the retina but for a display surface that is then viewed by the eyes from a distance. Hence the requirement is to capture each view for the same image plane since they are displayed in the same display plane.

A second confusion is that it is often assumed that camera separation should be the same as average eye separation. Again this would be true if we were generating images to write directly onto the retina. But in almost every 3D display, including head mounted displays, we are generating images to view on a remote display plane. The human factors of viewing 3D displays like this are such that it is rarely the case that camera separation should equal eye separation; more typically camera separation should be less than eye separation.

A growing number of independent studies⁷⁻¹⁰ investigating the depth range that 3D displays can reproduce have shown there are limits regarding how much depth can be viewed in-front and behind the display plane. The limits vary between displays and they are thought to arise in part because of the mismatch between focus and vergence requirements compared to viewing the real world. On a 3D display wherever the eyes verge to fuse depth behind and in-front of the screen, the eye's accommodation system must work to keep the image at the physical display screen distance in focus. This is thought to be a situation unlike the real world where vergence and accommodation systems work together in a flexible but inter-linked arrangement.

As a result of the lack of direct control over perceived depth in existing stereo camera control methods and the growing weight of evidence from human factors research, we have over the last five years developed a range of new camera models for stereo image generation. These are based on two assumptions; first that the content creator wants to frame a view before taking a stereo picture of that view; secondly that the content creator wants the resulting picture to be guaranteed to be comfortable to view on a target 3D display. The result has been methods that, from a given viewpoint, map depth from scene space to a display space in a controlled manner.^{1, 2, 9}




Figure 2. New stereoscopic camera models allow the content creator to position a camera view and define the depth range they wish to capture as shown in the diagram on the left. The stereoscopic cameras shown in the diagram on the right are then calculated automatically to capture the required scene depth range.

This new depth region mapping approach to stereo image generation is illustrated in Figure 2. The content creator simply positions a camera to frame the desired view of the scene and measures (usually automatically) the total range of depth in the scene as seen by the camera. The stereo camera separation is then calculated automatically with a guarantee that the viewer looking at the image on the target 3D display will not see perceived depth outside defined limits. There is then no need for trial and error capture to adjust camera separation.

We have implemented a recent depth region mapping algorithm¹ in the inThreeD³ plug-in for 3ds Max¹¹ which supports animated scenes. One of our goals in this paper is to test this approach and compare it with existing methods. In the following sections we discuss in turn each of the individual sections of the Cosmic Cookery movie; in particular we highlight the rendering tools we used and discuss the different issues the varying stereo support in these tools raised during the movie production process.

3. FLIGHT OUT OF THE SOLAR SYSTEM



Figure 3. Our journey into space and time begins at home on Earth.

The movie is constructed to begin and end at the Earth providing a reference point for the viewer as we fly far out in space and time. In this first sequence, Figure 3, we see the Earth and then fly out passing Jupiter and its giant red spot before sweeping on past Saturn where the rings and moons become clearly visible in 3D.

After hunting for software supporting suitable 3D models of the Solar System we chose Celestia.¹² Celestia can generate an animation and capture a sequence of still images along an animation path. Unfortunately Celestia's animation output is not stereo capable and therefore we had to capture left and right viewpoints separately. We used trial and error to set the separation for a pair of verging cameras and then kept the same fixed camera separation for the entire sequence. This caused artificial vertical disparity and does not allow us to adjust the available depth to track the key object in the scene. However, since the main object of interest is in the screen center the vertical disparity is minimized.

As the left and right images in this sequence were generated as separate screen captures, we had to process the entire sequence to crop images to the required size and then produce the side by side format required for the movie. We used automated scripts in Adobe Photoshop¹³ to size and crop images individually before merging the separate left and right images together into the side-by-side format using the batch processing mode of StereoPhoto Maker.¹⁴

4. FROM THE MILKY WAY TO VIRGO

The Sun is one of the hundred thousand million stars that make up the Milky Way and as we leave the Solar system we see a spectacular view of our own spiral galaxy. Then traveling across intergalactic space, we encounter the large and small Magellanic clouds, two miniature galaxies attached to the Milky Way. Eventually we arrive at the Virgo cluster, at the center of Virgo lies a large galaxy called M87 which is home to a massive black hole weighing more than three thousand million times the weight of the Sun.

This sequence was rendered by the Visualization group at NCSA, UIUC, using their PartiView¹⁵ software tool. PartiView allows the stereo camera separation to be modified along a flight path on a frame by frame basis. A verging camera model is available providing control of focal point.

Trans-Atlantic communication of the movie was achieved using ftp and JPEG compressed still sequences. Although using lossy compression degrades image quality and is not ideal for production, the substantial reduction

in file size made this worthwhile when communicating by ftp. Stills captured from an MPEG preview movie proved a useful tool for feedback on depth adjustments. During this process the left and right image sequences were merged using StereoPhoto Maker working in batch mode to create side-by-side image pairs. The result was coded and previewed in stereo using QuickTime Pro¹⁶ and the nVidia Quadro horizontal span mode so that left and right views appeared on the left and right channels driving the projectors.

5. FLIGHT THROUGH THE 2DF SURVEY



Figure 4. The 2dF Galaxy Redshift Survey, this fly through contains some of the highest disparities in the animation as galaxies rapidly fly past the viewer.

The 2dF Galaxy Redshift Survey (2dFGRS)¹⁷ was designed to measure redshifts for approximately 250,000 galaxies and maps our cosmic neighborhood far beyond the Virgo cluster. This sequence begins with a fly through of the galaxies, see Figure 4, and then pans out to show the whole survey. As the camera zooms out the view shows how galaxies clump together in a rich variety of structures, revealing an intricate cosmic web of filaments and clusters.

The sequence was generated at Durham from the 2dFGRS data and a series of galaxy images using the NCSA PartiView¹⁵ software. This involved converting the observation data to a 3D model with galaxy images and then planning an animation path through the sequence. The stereoscopic camera separation was then defined and tuned throughout the sequence to give the best depth effect within a reasonable comfortable range. Because PartiView presents an interactive graphics interface to the user it is possible to preview flight paths in real time and this allows faster iteration when making creative decisions and tuning the stereo.

In this sequence certain fast moving objects briefly exceed what we usually consider to be a comfortable depth range when they are flying rapidly off the screen. This produces a dramatic effect but is used sparingly as we suspect were we to repeatedly over-drive the vergence muscles in this way it would result in eye fatigue.

PartiView output images are saved as separate left and right files using the ppm file format. This required batch conversion, using Adobe PhotoShop, to PNG and then merging into the side-by-side format for the final production using StereoPhoto Maker.

6. THE MILLENNIUM SIMULATION

This section of the movie, shown in Figure 5, illustrates the simulated dark matter distribution in the universe at the present time, using data from the Millennium Simulation.¹⁸ The structure of the dark matter looks like the observed cosmic web seen in the previous sequence reflecting the statistically similar nature of the real and simulated universes. These vast concentrations of dark matter are believed to surround real galaxies; the gravity of the dark matter holding the galaxies together. Dark matter is thought to consist of particles significantly different from those that make up the atoms of stars, planets and people.

The Millennium Simulation represents a landmark in computational cosmology following the gravitational interaction between more than 10^{10} particles in a volume larger than that mapped by the 2dFGRS which featured




Figure 5. An image of the Millennium simulation illustrating the distribution of dark matter at the present day.

in Section 5. This sequence shows approximately 5×10^6 particles or 0.05% of the total. The analysis of such a massive simulation has posed a new set of challenges for cosmologists, resulting in shifts in their research methodology. Visualizations such as the one described here are becoming a valuable research tool in addition to providing striking images for outreach programmes.

In order to render these frames a computing cluster of 16 Sun SPARC processors was used running custom rendering software¹⁸ over several days. The software supported verging stereoscopic cameras and again required several trial runs to establish a reasonable perceived depth range. With camera control limited to a single setting of camera separation a compositional balance was required to illustrate the dark matter halo at the end in sufficient 3D depth while not producing excessive depth in the smaller clusters flying out of the screen towards the viewer early in the sequence.

7. THE BIG BANG AND INFLATION




Figure 6. A still from the middle of the Big Bang sequence illustrating the creation of the universe as it grew from a pinpoint to the size of an orange, before stereoscopically cross-fading to the show the CMB radiation; the evidence still measurable today from the Big Bang.

Here we visualize the Big Bang; when a feature the size of a pin-point formed our universe in a process known as vacuum fluctuation. This was followed by inflation; a period of rapid expansion, lasting a tiny fraction of a second. Following which the universe was approximately the size of a large orange and the forces and particles of matter we are familiar with began to form. The sequence ends showing the fluctuations in the Cosmic Microwave Background (CMB), measurable today as the fossil record of the Big Bang illustrated in Figure 6.

In this sequence we planned a pinpoint to travel towards the viewer as the commentary began to explain the origin of cosmic structures. We then planned an explosion to illustrate the Big Bang and inflation followed by

the appearance of the illustrative orange cross fading to show the CMB. If we had rendered the sequence with a standard stereo depth mapping most of the stereoscopic depth would have been used in the pinpoint section and the explosion and subsequent action would have been left stereoscopically flat. Instead we used our inThreeD plug-in for 3dsMax allowing us to manipulate the depth mapping ensuring most of the available stereoscopic depth was allocated to the explosion and the remaining part of the sequence.

The new control provided by the three region depth mapping implemented in the inThreeD plug-in provided a clear benefit in composing the depth presentation of this sequence. In addition it required significantly fewer trial renderings to adjust the depth effect since we knew the mapping by design. An unexpected benefit of the plug-in was the ability to directly generate correctly cropped side-by-side image sequences as output. This then required no time to batch process the images to change format, cropping or image size and they could be directly imported into QuickTime Pro and Adobe Premiere for movie creation.

Overall the time savings using the inThreeD plug-in to render this sequence allowed us to spend significantly more time considering creative rather than technical aspects of production. This was in contrast to the other tools we used where we spent significant time tuning detailed technical aspects of the stereo production.

8. DARK MATTER HALOS




Figure 7. In this sequence we see dark matter condensing due to gravitational attraction to form dark matter halos.

This simulation shows the formation of dark matter structure in the universe around one dark matter halo,¹⁹ illustrated in Figure 7. These halos form due to the small fluctuations in the early universe and the action of gravity over time. As we see in the next sequence within dark matter halos galaxies form as matter is sucked in over billions of years.

Output from the halo simulation software consisted of 900 files of volumetric data each containing 256^3 values with a total data size of 57 Gbytes. We chose Kitware's VolView²⁰ as the rendering software for this sequence as it was able to use floating point voxel data and produces high quality volume renderings. However, VolView had two drawbacks; it was designed for single file manipulation rather than batch processing and it was not able to save stereo images to file.

Fortunately KitWare were able to supply scripts enabling VolView to process the large sequences of data files and output left and right stereo images to file. These allowed us to batch process the animation with a total rendering time for the sequence of the order of two days. Without this support manual loading and rendering of the 900 individual files would have been an infeasible task in the time available.

Stereoscopic camera control in VolView is provided simply by setting a vergence angle and therefore repeated trial and error renderings were needed to ensure the stereoscopic depth was acceptable throughout the sequence. Therefore vertical disparity is a potential issue, however as the sequence was primarily centered in view vertical disparity is minimized. Since VolView rendering was to the on-screen buffer the image resolution was defined by window size. This combined with the output of separate left and right files required post-processing to crop

and re-size images and merge them to form combined left/right side-by-side pairs. Both Adobe Photoshop and StereoPhoto Maker batch mode processing tools were used to do this.

9. SPIRAL GALAXY FORMATION




Figure 8. Simulated behaviour of gas and stars as a spiral galaxy forms.

In this sequence we again turn back the clock 13 billion years to the beginning of the universe and now examine what happens in the inner region of a dark matter halo. Over time matter collects drawn together by gravity, it cools and as the density increases myriad stars form. In particular this simulation illustrates how the gas and stars can merge to form a spiral galaxy,²¹ as shown in Figure 8.

As with the previous sequence there were a series of 256³ simulation output files; 850 in total. Like the last sequence we used VolView to render these and relied on the scripts supplied by Kitware to handle batch rendering and stereo image output. Post-processing using Adobe Photoshop and StereoPhoto Maker was required to obtain the final side-by-side stereo pair. With rendering times for this sequence exceeding two days the stereo camera control available in VolView and consequent trial and error approach to adjusting camera separation to obtain a good depth presentation proved time consuming.

10. FINALE




Figure 9. The simulated galaxy, left, and real galaxy, right, are shown together to allow a direct visual comparison.

The life of a galaxy, from its microscopic beginnings in the early universe to its present day splendor can be recreated in a computer. This final sequence, Figure 9, shows both a simulated galaxy and a volumetric rendering of a real galaxy side-by-side. The sequence zooms into the two galaxies and then rotates them to provide a visual comparison between simulation and reality.

Rendering this sequence using VolView involved rendering both the rotation sequence for the real galaxy and a separate sequence for the simulated galaxy. Each sequence was then scaled individually in both left and right views to achieve the zoom-in. This required loading the four image sequences individually and then scaling each image sequence up to full size using built in Adobe Premiere animation tools. Ensuring the scaling was consistent for each of the four sub-images required some care.

Overall the editing and post-processing for this sequence was surprisingly time-consuming, particularly given the relative simplicity of the individual tasks required. The renaming software tool MRename²² became of crucial practical significance to ensure file names were compatible as we transferred them between the different tools used.

11. WORK FLOW AND TOOLS




Figure 10. A sketch of the work flow for the dark matter halo sequence on the left and the Big Bang sequence on the right.

The work-flow for each sequence varied depending on the degree to which the rendering systems supported stereoscopic rendering and output. In addition edits to a sequence typically required manual intervention and batch processing.

The work flow in two different sequences is shown in Figure 10. For the dark matter halo the volume renderer, while supporting stereo, rendered to screen resolution and produced as output separate left and right image files. These images required re-naming, cropping and re-sizing before merging into a side-by-side left/right pair. This then allowed preview movies to be created from the PNG format image files using QuickTime Pro. As the stereoscopic camera separation was only controlled approximately using verging cameras this process was repeated more than once to tune perceived depth before the final version was produced.

For the Big Bang sequence we used the InThreeD plug-in which allowed direct control over the stereoscopic effect and produced as output the correct size side-by-side pairs. As a result not only was the stereo effect easier to control there was no need for any image processing before the preview movie could be created. The extra time this saved allowed us to spend more time on modeling and planning the composition of depth in this sequence.

12. COMPRESSION AND PLAYBACK

The main goal both in production and playback was to have the best possible image quality and where the rendering tools allowed we used lossless image coding formats; either PNG or TIFF. When these were not available JPEG was used at the highest quality setting. Our aim was to avoid multiple lossy compression steps as files were edited and re-edited several times in some sequences.

Each sequence after rendering and editing was stored as a set of numbered stereo images in side-by-side format, using either lossless PNG or high quality JPEG when this was not possible. For preview movies we used

QuickTime Pro to read in the numbered sequence and playback the resulting preview movie in realtime. This ran on an nVidia Quadro graphics card running in horizontal span mode across dual DVI outputs driving our BARCO Gemini passive stereo dual projector and preview monitors. QuickTime worked well for early previews of individual sequences, when the preview was small enough to fit in system memory (2Gbytes), but the full movie was too large to playback this way on a standard PC.

For the production version of the movie we switched to use an Inition 3DVidBox, containing a high speed RAID array and dual DVI nVidia Quadro cards again driving passive dual channel projection. This allowed us to playback the entire movie (47Gbytes) as a single uncompressed AVI file at 2x1024x768 at 25fps. The final edit to create this version of the movie was undertaken using Adobe Premiere, this brought together the individual sequences, titles and sound track. Real time playback of the full production version of the movie was achieved using the BSplayer software tool.²³

Subsequently we have used Premiere and StereoPhoto Maker to generate anaglyph versions of the movie for playback on 2D displays. These have required some care with compression standards and we chose QuickTime using MPEG4 coding and also a lossy AVI format. It has however not been straight-forward to find a single compression standard that works well on the range of platforms we have ported the movie to (MAC, PC, DVD-Video). Monoscopic versions of the movie have been simple to generate using Premiere to crop the sequence to remove the right channel images.

12.1. Editing

Tools for editing stereoscopic images are not widely available and often lack the batch processing technologies suitable for movie generation. A range of tools were therefore used in post-rendering edits and movie production.

Each of the rendering systems were used to generate sequences of still stereoscopic images. Some of the rendering tools were able to generate side-by-side pairs at the required resolution 1024x768x2 but many generated non-standard resolutions and separate files for the left and right views. This required tools capable of batch cropping or extending images to the required size and often this also required splitting and/or merging of stereo pairs. To achieve this Adobe Photoshop and Stereo PhotoMaker were both important in the production process.

Once the images were formatted correctly we used QuickTimePro to generate rushes for previews and then finally Adobe Premiere to merge the sequences into a single movie file.

A crucial tool was file renaming software used to rename sequences so that the various automated file readers read them in the correct order. For this we used MRename²² which had a flexible and powerful interface based on a regular grammar.

13. SUPPORTING ADDITIONAL DISPLAY TYPES

We have already described support for passive stereoscopic projection, anaglyph displays and monoscopic (2D) sequences. In addition several clips from the movie were formatted using custom software in the colour column interleaved format required for the Sharp RD3D display. This interleaved each pair in the required format and saved them using lossless PNG files. These could then be read in as a numbered sequence and played back in QuickTime Pro, the PNG lossless compression retained the interleaving pattern.

In all cases we assumed the largest display used was the 2m projection screen and the scaling of the images to fit smaller displays would reduce disparity to remain within comfortable limits. The images were not recalculated for each smaller display and would have benefited from this had sufficient time been available.

We also created two 1mx0.9m lenticular prints of the simulated and real spiral galaxies for the exhibition. The prints were produced by David Burder at 3D Images Ltd. but the images were rendered at Durham again using VolView. A series of twenty seven images were created at one degree rotation intervals and merged into a single layered Photoshop file. As these had to use simple camera rotation around a point to generate the views vertical disparity is present in the stereoscopic image. However, as the content is centered in the image this artifact is minimized.

14. CONCLUSIONS

Our first aim was to produce a 3D movie that would visualize current theories and observations describing the creation of cosmic structures in the universe and in particular the formation of matter into galaxies. The movie was well received by the general public at its first showing at the Royal Society Summer Science Exhibition in July 2005. It has subsequently been successful in attracting attention to and generating substantial discussion about research into the universe and its origins at many national and international venues, including the American Astronomical Society meeting in January 2006.

Technically we used a wide range of 3D techniques and tools to produce the movie and associated material. Underlying this we had a constant requirement to retain the best possible stereo quality throughout the production for the viewer; three specific issues concerned us:

Camera control for comfortable depth In most of the sequences we had to use traditional stereoscopic camera methods to set camera parameters; requiring a trial and error approach to depth composition. We also used our new methods; based in human factors research and allowing precise control over the depth presentation in an image. Both approaches worked well but the new methods gave more control and were much faster to use; freeing more time to consider scene design and the creative use of the stereo effect.

Image editing Very few image and movie editing tools are stereo aware and this resulted in a need to use a range of tools to edit and format the image sequences. Of these four tools became essential; Adobe Photoshop, Stereo PhotoMaker, MRename and Adobe Premiere. We would particularly highlight Stereo PhotoMaker as pointing the way to general purpose cross-platform stereo image editing.

Compression During production we worked whenever possible with lossless image formats such as PNG and TIFF and always used side-by-side full resolution images. When JPEG was the only available output format from a renderer we saved images at the highest quality and converted to a lossless format for editing. For presentation of the movies our main platform is a 3DVidBox allowing us to playback uncompressed AVI video at full 2x1024x768 resolution. We have also successfully created and distributed lossy QuickTime and AVI versions of the movie in side-by-side, anaglyph and monoscopic formats.

We were careful to control the final depth effect in the movie and tuned the effect to a 2m diagonal screen viewed from about 2m. It is very clear when we view this enlarged on our 3m screen that the depth starts to push the limits of acceptable comfort. This confirms human factors work in the literature and predictions from the viewing geometry. In addition the reduced sizes of the movie for desktop display work but show less depth than they ideally could. We believe it remains an open question how portability can be successfully managed if 3D movies become a widespread consumer commodity

It is worth noting that without nVidia introducing comprehensive driver support and multi-screen windows management technology we would find stereo displays very much more difficult to use. We would encourage industry to enhance this support in three specific ways:

Capture Enable real time capture of the pixel stereo format fed to the display. This would allow direct capture of real time graphics in a display specific format, e.g. colour column interleaved, for subsequent playback as a movie, or:

Movie support Enable the side-by-side format we used to play back to the left and right stereo buffers (instead of the left and right screens). Then the stereo pixel format required for a specific display device could be generated in real time by the graphics driver and the side-by-side movie format could become standard across all stereoscopic platforms.

Left/Right standards In our recent experience left/right standards for stereo are urgently needed. It is still feasible to show these reversed and still hard for non-experts (and experts) to detect this problem.

The Cosmic Cookery project was a significant production effort allowing us to exercise a wide range of standards and tools for animated stereoscopic movie generation. We believe we have shown that this is a feasible process using current desktop tools. We have been encouraged that the anticipated benefits of our research into stereoscopic camera control have been demonstrated in this project.

ACKNOWLEDGMENTS

We would like to acknowledge the support and help of a wide range of people, funding agencies and companies in the production of the Cosmic Cookery movie and Royal Society exhibits. In addition to those mentioned in the text we would like to thank; for the real galaxy volume data and colour schemes Richard Bower and Simon Morris at Durham University, for rendering the Milky Way to Virgo sequence Bob Patterson and Stuart Levy at NCSA UIUC, for his software and advice on rendering the Millennium Simulation data Volker Springel at MPA Garching, for advice on conveying science to real people Pete Edwards, for systems support Lydia Heck, for anaglyph conversions summer students of the Institute for Computational Cosmology including Diego Hartasanchez, for substantial equipment loans and advice for the exhibition Initium Limited particularly Stuart Cupit, for Premiere advice and creating the final edit of the movie Phil Brown at Media55, for laboratory space and equipment access the eScience Research Institute Durham University, for additional financial support the Dean of the Science Faculty Durham University.

REFERENCES

1. N. Holliman, "Mapping perceived depth to regions of interest in stereoscopic images," in *Stereoscopic Displays and Applications XV, Proceedings of SPIE* **5291**, 2004.
2. N. Holliman, "Smoothing region boundaries in variable depth mapping for real time stereoscopic images," in *Stereoscopic Displays and Applications XVI, Proceedings of SPIE* **5664A**, 2005.
3. B. Froner and N. Holliman, "Implementating an improved stereoscopic camera model," in *Eurographics Theory and Practice of Computer Graphics, ISBN: 3-905673-56-8*, 2005.
4. nVidia Corporation, "nVidia Quadro FX." <http://www.nvidia.com/>.
5. D. McAllister, ed., *Stereo computer graphics and other true 3D technologies*, Princeton University Press, 2003. ISBN 0-691-08741-5.
6. D. Diner and D. Fender, *Human engineering in stereoscopic viewing devices*, Plenum Press, 1993. ISBN 0-306-44667-7.
7. Y. Yeh and L. Silverstein, "Limits of fusion and depth judgements in stereoscopic color displays," *Human Factors* **1**(32), 1990.
8. A. Woods, T. Docherty, and R. Koch, "Image distortions in stereoscopic video systems," *Proceedings of SPIE* **1915**, 1993.
9. G. Jones, D. Lee, N. Holliman, and D. Ezra, "Controlling perceived depth in stereoscopic images," in *Stereoscopic Displays and Virtual Reality Systems VIII, Proceedings of SPIE* **4297A**, 2001.
10. E. W. Jin, M. E. Miller, S. Endrikhovski, and C. D. Cerosaletti, "Creating a comfortable stereoscopic viewing experience: effects of viewing (accommodative) distance and field of view on fusional range," in *Stereoscopic Displays and Applications XVI, Proceedings of SPIE* **5664A**, 2005.
11. Autodesk Corporation, "Autodesk 3ds Max." <http://usa.autodesk.com/>.
12. C. Laurel, "Celestia." <http://www.shatters.net/celestia/>.
13. Adobe Corporation, "Adobe Photoshop." <http://www.adobe.com/products/photoshop/main.html>.
14. M. Suto, "StereoPhoto Maker." <http://stereo.jpn.org/eng/>.
15. S. Levy, "PartiView." <http://niri.ncsa.uiuc.edu/partiview/>.
16. Apple Corporation, "QuickTime Pro." <http://www.apple.com/quicktime/download/win.html>.
17. M. Colless, G.B. Dalton, S. Maddox, W. Sutherland, P. Norberg, S. Cole, J. Bland-Hawthorn, T. Bridges, R. Cannon, C. Collins, W. Couch, N. Cross, K. Deeley, R. D. Propriis, S. Driver, G. Efstathiou, R. Ellis, C. Frenk, K. Glazebrook, C. Jackson, O. Lahav, I. Lewis, S. Lumsden, D. Madgwick, J. Peacock, B. Peterson, I. Price, M. Seaborne, and K. Taylor, "The 2df galaxy redshift survey: Spectra and redshifts," *Mon.Not.Roy.Astron.Soc* **328**, 2001.
18. V. Springel, S. White, A. Jenkins, C. Frenk, N. Yoshida, L. Gao, J. Navarro, R. Thacker, D. Croton, J. Helly, J. Peacock, S. Cole, P. Thomas, H. Couchman, A. Evrard, J. Colberg, and F. Pearce, "Simulations of the formation, evolution and clustering of galaxies and quasars," *Nature* **435**(7042).
19. J. Navarro, E. Hayashi, C. Power, A. Jenkins, C. Frenk, S. White, V. Springel, J. Stadel, and T. Quinn, "The inner structure of Λ CDM haloes - iii. universality and asymptotic slopes," *Mon.Not.Roy.Astron.Soc* **349**, 2004.
20. Kitware Corporation, "VolView." <http://www.kitware.com/products/volvview.html>.
21. T. Okamoto, V. Eke, C. Frenk, and A. Jenkins, "Effects of feedback on the morphology of galaxy discs," *Mon.Not.Roy.Astron.Soc* **363**, 2005.
22. Rajko, "MRename." <http://fly.cc.fer.hr/~rajko/MRename/>.
23. Boris, "BSplayer." <http://www.bsplayer.org/>.