DOES SIZE MATTER?
THE IMPACT OF SCREEN SIZE ON STEREOSCOPIC 3DTV

Laurent Chauvier, Kevin Murray, Simon Parnall, Ray Taylor, James Walker NDS

Increase the size of your television screen and the picture gets bigger - but if you have a 3D screen, should the picture also get deeper? From a single common signal, what is the effect of larger or smaller screen sizes on our 3D perception?

This paper looks at what happens to the S3D effect as screen sizes scale. It does so based on an analysis of the geometry of stereoscopic 3DTV and argues that there is an impact as screen sizes scale. The paper looks at the theoretical potential for compensation for different screen sizes under typical viewing distances, and the challenges in performing this in low-cost STBs and TV sets where the input is a stereoscopic image pair.

INTRODUCTION

Stereoscopic 3DTV (S3DTV) continues to maintain a high level of interest and enthusiasm from consumers, content creators and now broadcasters too [9,10,11]. We expect to see S3DTV content for the home come from a wide range of sources. Several broadcasters have made announcements about S3DTV trials and services, and by the end of 2010 many will be transmitting on-air S3DTV content [16,17,18]. Then there is packaged S3DTV content using extensions to the BluRay specification and the new 3D BluRay players to carry movies. And finally yet another reason for consumers to acquire S3D television sets are the various updates that are either already released or expected for various games platforms, together with new games that will provide a thrilling 3D gaming experience.

Matching this wide range of sources is the wide range of S3D television sets, primarily based on shutter technologies, many of which are already on the market place [12,13,14,15]. These televisions sets currently cover a relatively limited range of sizes between 40″ and 55″—the premium segment—but this size range will increase in the future, especially as the announced S3D capable projectors that are easily able to support 80″ images, or even larger, become available.

This paper explores the potential consequences of this wider range of sizes for televisions on the experience that the consumer will have when watching S3DTV content. Specifically, it asks the questions if and how the S3DTV experience is affected by changing the size of a television set. These questions are posed and answered examining how the perception of the depth of an object is altered as the size of a display changes. On the basis of this, it is argued that different display sizes will provide a perceivably different experience of the same S3DTV content.

It is clearly not realistic to expect broadcasters or content packagers[1] to produce different versions of the content for each display size, and so we can assume that all sets will be receiving exactly the same broadcast or packaged signal. Therefore this paper explores two potential main methods to provide corrections that could be made to the stereoscopic images to compensate for variations in display size. Each approach takes a different view on what the "correct" behaviour should be, and as such each has its own set of benefits and drawbacks, and these are discussed.

[1] Games content, at least where created within the games machine and not using pre-calculated sequences, has the potential to be created specifically for the display size to which the games machine is connected. Thus, at least in theory, 3D games can automatically adapt to the size of the display by slightly altering the parameters they use to render each frame.

THE GEOMETRY OF STEREOSCOPIC 3DTV

Stereoscopic 3DTV (S3DTV) provides the illusion of a 3D experience by providing two views, one for each eye. Whilst there are numerous technologies that provide the separation of the images to the eye[5] (shutters, passive polarisers, parallax barriers and lenticular lens to name the four best known technologies), all rely on the way that the two separate images are interpreted by the human visual system. The basic geometry of this interpretation, and how the depth of an object is perceived, is well known[1,3] and shown in ray diagram form in figure 1.

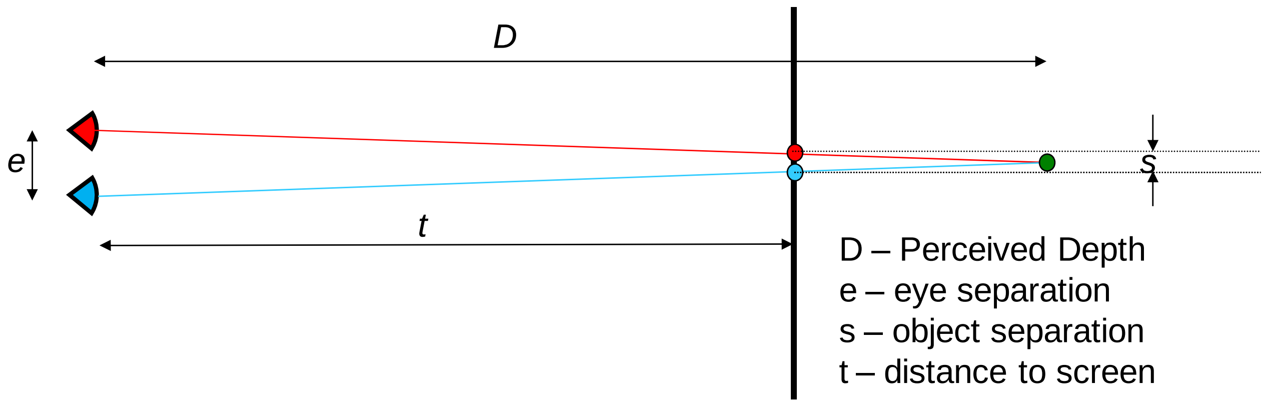

Figure 1: S3DTV Geometry Diagram

From this, we can see that the perceived depth of an object is determined by the following formula, using the symbols from figure 1:

$$D = \frac{e * t}{e - s}$$

There are numerous factors that affect the perceived depth of an object, including angle of view. For the remainder of the paper we shall ignore variation in the following factors:

- t the distance to the television is clearly a variable, but as with common broadcasting practice [4], we will assume that the viewing distance is fixed at three times the height of the display (3H). It is very rare for this to alter during viewing.
- e the eye separation, for any given viewer, is fixed, though every different person will have subtly different values for e [7].
- Angle of view[8]. For simplicity, it is assumed that the viewer is viewing objects as shown in figure 1, perpendicular to the screen. Clearly the angle will affect the view, but the angle will not change noticeably during viewing and so the impact of off-centre viewing is consistent.

Implications

The formula shows that the perceived depth is inversely proportional to the separation of the renditions. This in turn means that scaling the separation does not have a linear impact on the depth at which the object is perceived. For example, a set of objects at $4t$, $3t$ and $2t$ in an original screen will be perceived on a new screen of half the size at $1.6t_{new}$, $1.5t_{new}$ and $1.3t_{new}$ where t_{new} is the new distance to the new half sized screen. The following sections and diagrams explore the implications of this depth distortion in more detail.

THE DEPTH BUDGET AND DEPTH RANGE

The depth budget refers to the maximum range of depths in use at any given time, the distance between the nearest and furthest objects that the viewer perceives. The choice of depth budget may be artistic, but it has a significant impact on the comfort of viewing S3DTV [6]. As a gross simplification, the use of excessive depth budget or placing objects at the limits of the depth range (especially very close to the viewer and with typical TV viewing conditions) decreases the comfort of viewing content. Thus, in general, content uses a restrained depth budget [2].

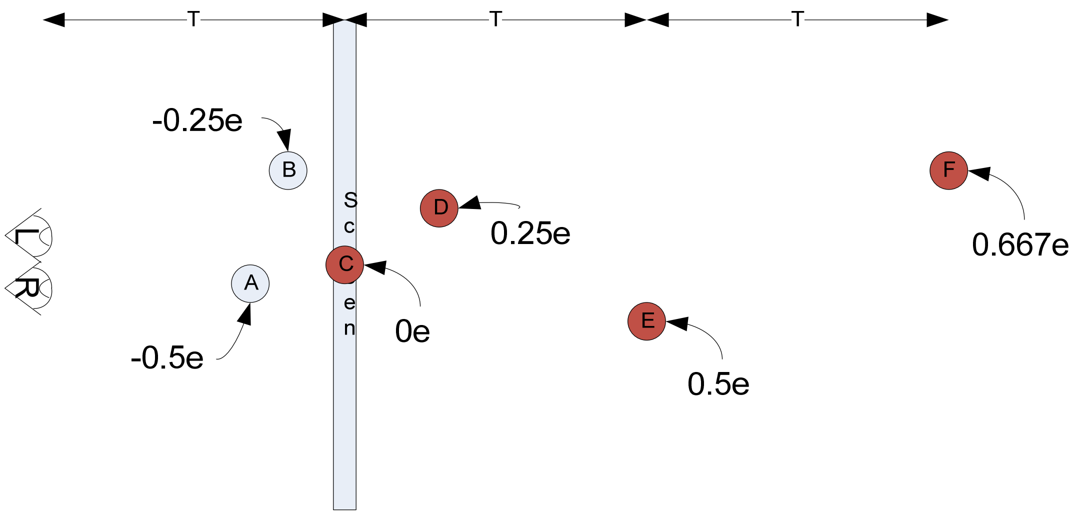

Figure 2: Depth placement of objects and the related left-right separation

Figure 2 shows the perceived depth placement for a range of objects based on the separation of the left and right rendition of the objects. Each object is also labelled with the separation of its left and right renditions on the screen as a multiple of eye separation. In the rest of this paper, we will show how this depth varies, or can be varied.

THE EFFECT OF DISPLAY SIZE CHANGES

Let us assume that we have a stereoscopic picture that provides for objects placed as shown in figure 2. Now let us consider taking this image and displaying it on two new screens of different sizes. Clearly, as the screen size scales, the separation scales and so the perceived depth scales. Figure 3 shows the impact on the depth range when displayed on a screen half of the size, showing the objects at their newly labelled positions. A result of this scaling is that objects that were previously at infinity are moved to a depth of 2t, and this is shown as ∞ in figure 3. In comparison figure 4 shows the impact on a screen twice the original intended size.

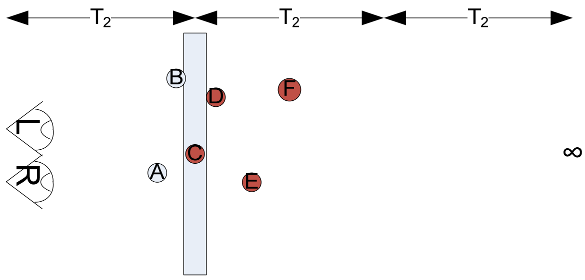

Figure 3: Scene from Figure 2, as perceived on a screen of half size

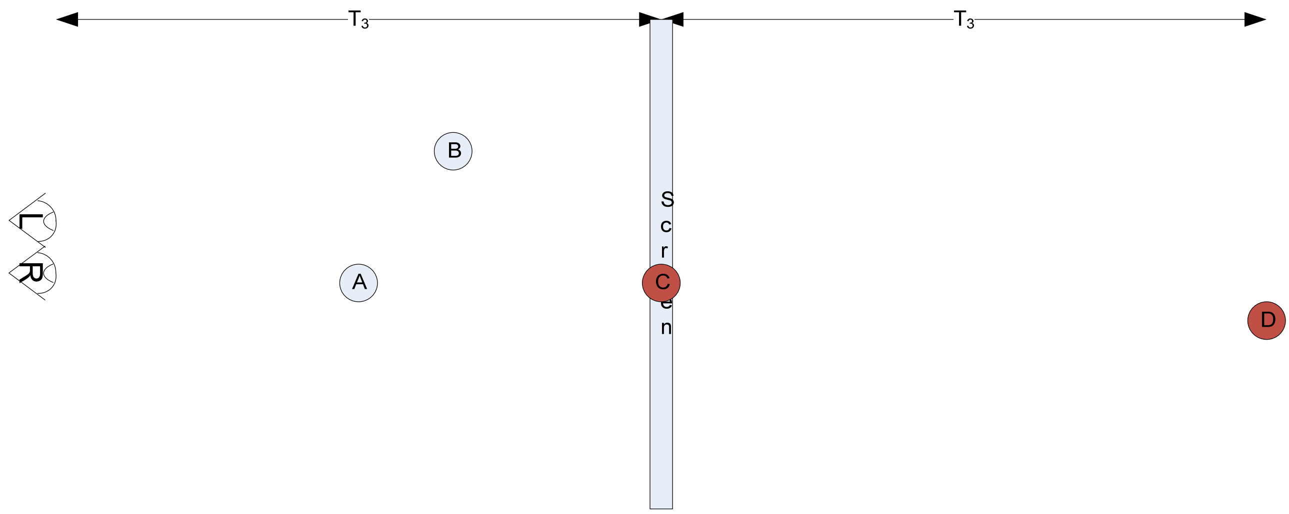

Figure 4: Scene from Figure 2, as perceived on a screen of double size

As may be seen from these two figures, shrinking the screen compresses the depth range whereas enlarging the screen expands the depth, but far more dramatically than just altering *t*. The perceived depths from the figures are shown in table 1, measured both in absolute terms compared to the original display size, and in relative terms.

Object		Depths in Original Units (t)			Display Size Relative Depths (t_2 and t_3)	
		Original	Half Size	Double Size	Half Size	Double Size
A	-e/2	2/3t	2/5t	t	4/5	½
B	-e/4	4/5t	4/9t	4/3 t	7/8	2/3
C	0	t	t/2	2t	t_2	t_3
D	+e/4	4/3t	4/7t	4t	8/7	2
E	+e/2	2t	2/3t	∞	4/3	∞
F	+2e/3	3t	3/2t	>∞	3/2	>∞
∞	e	∞	2t	>∞	4	>∞

Table 1: Variation in perceived depth as screen size scales

There are several implications from scaling the images as shown in figures 3 and 4. Firstly, consider an object that is moving towards the viewer or camera at a linear speed in the original image. When this object is viewed on a different sized display, the speed will no longer be perceived as linear towards the viewer—it will either be speeding up or slowing down as it approaches the viewer depending on the scaling and positioning of the object. In a very similar fashion, this will result in minor apparent size distortions.

Next, consider looking at an object. Cues such as binocular disparity allow the object to be accurately placed in depth by the human visual system. This depth placement is then also used by the human visual system to estimate the size of the object. We are already familiar with this effect, which is becoming known as the "subbuteo effect" in reference to the popular old football table game, describing how football players can appear as miniature people. This effect will be strengthened by shrinking the screen size as we now have the people appearing closer to the viewer. That is, by appearing closer the human visual system will interpret the football players as being even smaller.

Finally, if extremes of depth are used in the content, scaling the display introduces another set of problems. Where the display is larger than the original, then objects can appear "beyond infinity" (although for short term viewing, the human visual system seems to be able to accommodate this). For increased image sizes, objects close to the viewer are less of a problem since the increased (typical) viewing distance provides a degree of compensation, as shown in the table above. By contrast, where the display is smaller than the original, certain depth locations are no-longer achievable. Most obviously, it is not possible to place an object at mathematical infinity. Unlike enlarging a screen, close objects can now pose a problem as they are closer to the viewer, even though the reduced separation also places them closer to the screen.

DEPTH PERCEPTION CORRECTION

In the following sections, we outline two approaches to correcting the depth impact that we introduced above — relative and absolute.

Relative Depth Correction

The goal for relative depth correction is to alter the stereoscopic image pair so that depths scale consistently. An example of this is shown in figure 5, where the original screen and objects are shown dashed, and the new scaled screen is shown solid (in this case the screen is half the size). The depths that would be seen without this correction, as illustrated in figure 3 above, are shown in the bottom left of the diagram. This shows the desired relative depth corrections that are required to overcome the non-linear depth compression and maintain the correctly scaled relative depths.

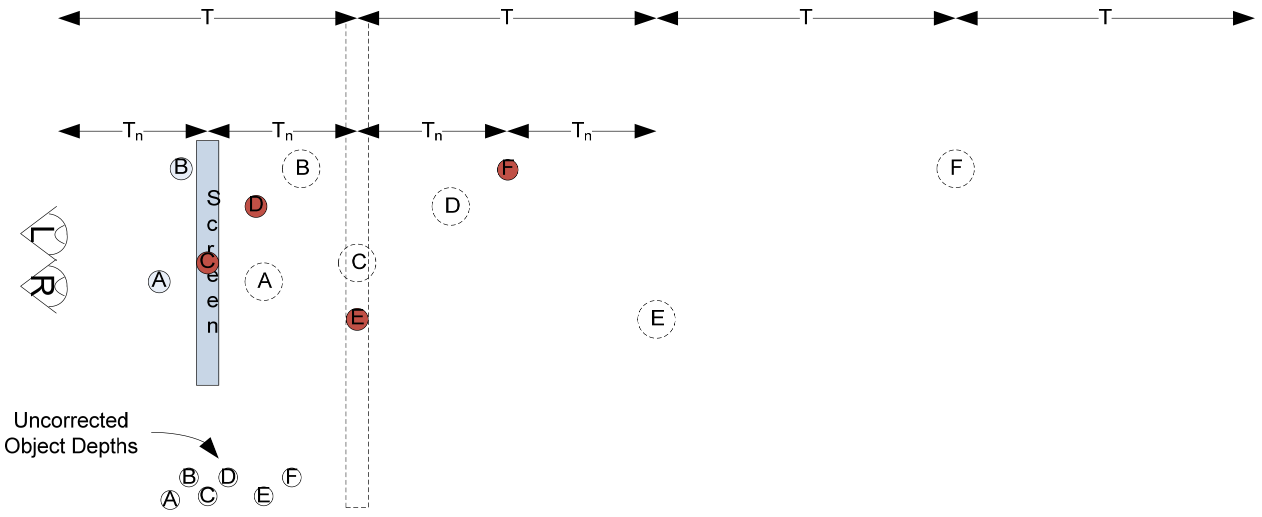

Figure 5: Relative Depth Corrected Example

Given the depth equation, the desire is to retain the same absolute object separation, s, on the objects on the new display, even once the display is scaled. However, for a display that is scaled relative to the original target display by a factor of *f*, the separation is also scaled, so that $s_{new} = s * f$. Hence, the correction we need to apply is

$$shift = s - s_{new}$$

which is

$$shift = s * (1 - f)$$

This means that the shift to be applied to an image varies across the image in relationship to the depth that is perceived for that part of the image. Therefore no single correction factor can be applied that is consistently correct.

Absolute Depth Correction (Window on the World)

Absolute depth correction is where the stereoscopic image is adapted so that the objects appear at the same absolute distance from the viewer regardless of the change in the size of the display. In this, objects can be moved from behind to in-front of the screen, or vice versa, depending on the size changes. This is shown in figure 6, where the original screen is shown dashed together with a new screen that is half the size of the original, and consequently half the distance from the viewer.

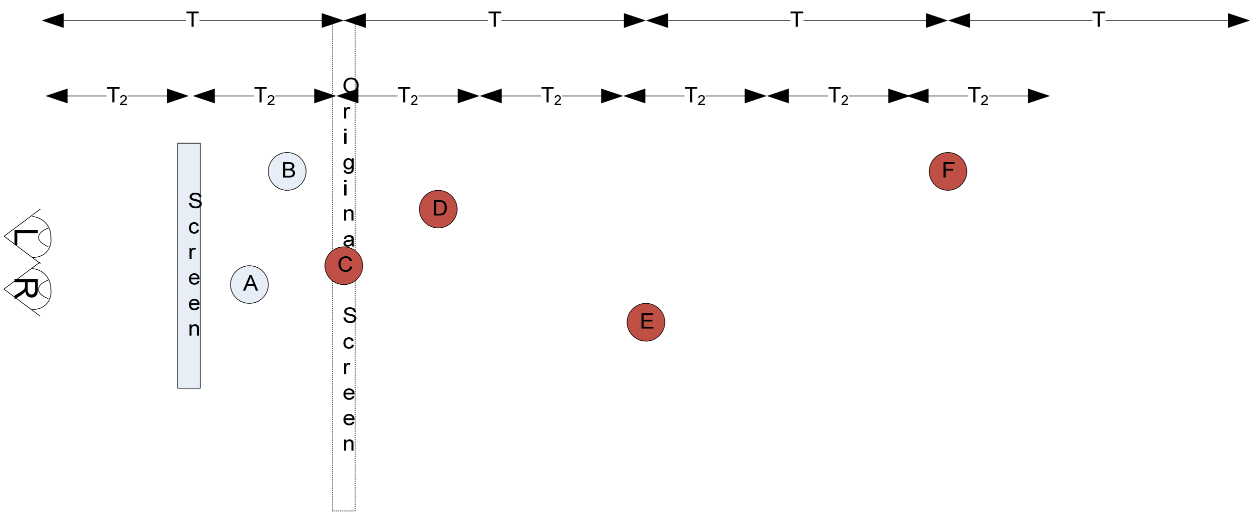

Figure 6: Absolute Depth

If we measure the separation as a fraction of the eye separation, i.e. s = i*e, then depth equation shown earlier becomes:

$$d = \frac{t}{1 - i}$$

The inverse, or the separation *i* required to place an object at depth d on a display at distance t, is:

$$i = 1 - \frac{t}{d}$$

For a screen that is scaled by a factor *f*, the new value for the distance to the screen, t_n, is $t_n = t * f$. For this scaled screen using the equations above, to represent an object at depth d requires a separation $i_{corrected}$:

$$i_{corrected} = 1 - (1 - i) * \frac{t_n}{t}$$

Which when t_n is replaced, simplifies to

$$i_{corrected} = 1 - (1 - i) * f$$

When the image is scaled, the separation is also scaled, so $i_n = i * f$, so the shift required is:

$$shift = e * i_{corrected} - e * i_n$$

Which, when substituted with the above becomes

$$shift = e * (1 - (1 - i) * f - i * f)$$

And which in turn simplifies to:

$$shift = e * (1 - f)$$

Thus a single shift across the entire image repositions all the objects back to their original depth placements as if they were viewed on the original screen at the original depth.

Comparison of Solutions

Absolute depth correction is clearly easier to apply, but it is important to consider the impact that this may have on the viewing experience, and the acceptability of this. For smaller displays, it means that the content will always be distant, and will occur further from screen depth. It is generally accepted that the closer to screen depth, the easier content is to view due to lower focus-convergence disparity. The absolute depth correction solution may reduce the amount of content that sits in this comfort zone. There is also a potential risk that making the action occur "further" from the screen may not be as acceptable or engaging to the viewer, somewhat like the differences between watching a play from a distance, rather than being seated closer to the action. Finally, the shift required by absolute correction will result in the loss of a small amount from the edges of the picture.

Although relative correction is a more complex solution requiring computationally intensive depth map calculation [19], the authors believe that this is an important alternative and so the next sections discuss some practical means of implementing an approximation of the relative depth correction. Clearly, however, comparative evaluations of these different solutions are required.

IMPLEMENTING APPROXIMATE RELATIVE DEPTH CORRECTION

In the previous sections we have seen how each separate object depth requires a different shift to be applied to compensate for the scaling of the display. Thus, an idealised solution requires an accurate depth map of the stereoscopic image and a processing stage that provides the required manipulation on a pixel by pixel basis. In theory, since an item of content could cover the entire range of depths, the shift required could vary from $e * (1 - f)$ (to correct for objects at infinity) through 0 (for objects at screen depth) through to negative values of perhaps $-e * (1 - f)$ or beyond.[2]

When an item of content is created, as we discussed earlier, it is rare for it to use the entire potential depth range. For any given item of content, the range of shifts to be applied will be defined by the range of depths that are represented within the stereoscopic content. Although there is still a learning experience underway to decide the most acceptable depth range to use, the current experience tends towards using a depth range from slightly behind the screen backwards, with a limit to the maximum depth. Translating this into fractions of eye separation, this currently appears to be between 0.6 e and about 0 or 0.1 *e* for most content (i.e. it is behind the screen in a relatively confined depth range), though clearly, this will be exceeded in certain dramatic conditions.

Figure 6 illustrates the impact of depth over this range, shown as fractions of eye separation, *e*, on the x axis, and the perceived depths over this range, in units of the distance to the display, on the y axis. The solid line shows the depth range for the original display. The dashed line shows the depth range for a screen of half the size, as measured in the units of the reduced distance to the display that is a result of the smaller screen. The final dotted line shows the depths with an arbitrary constant shift of 0.15e.

From this we can see that a shift can be used to minimise the overall total depth discrepancy that is introduced when a screen is scaled. The shift can be calculated in several ways, each with different properties. The simplest is to base this on the shift required to correct the mid-point of the shifts. In the case of the graph, the midpoint is 0.35, and would generate a shift of 0.175 e. However, the mid-point of the shifts is not the same as the midpoint of the depths, which is at a depth of 1.8 t, or a shift of about 0.44. This generates a correction of about 0.22 e.

[2] There is no theoretical limit to the separation that can be used for objects appearing in front of the screen. However, the convergence-focusing disparity makes perceiving objects far in front of the screen difficult. As such, except for very short durations, it is unlikely that a large negative separation will be used.

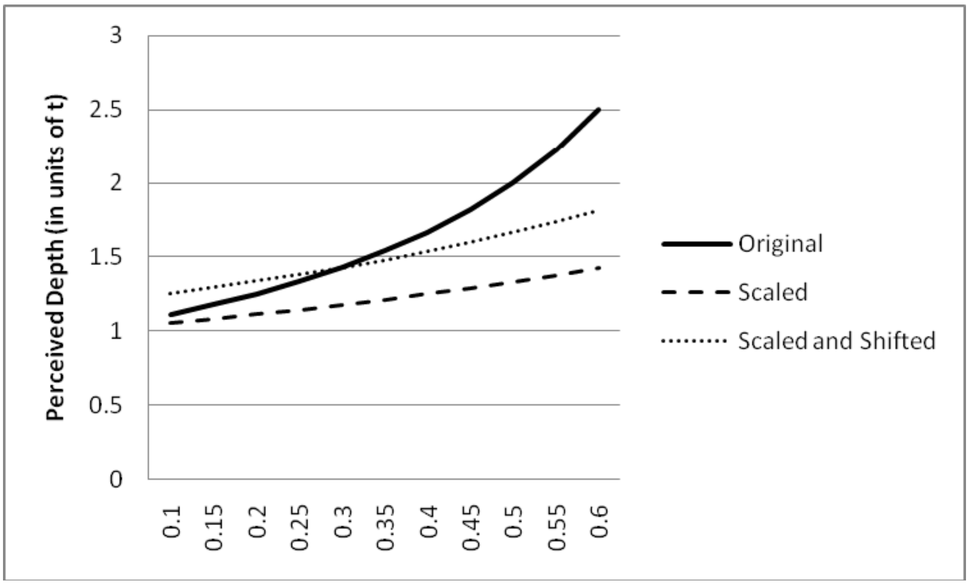

Figure 6: Graph showing impact of constant shift on perceived depth ranges

More complicated measures can take into account the use of each given perceived depth, i.e. it can be based on a statistical processing of the depth information. To perform this, however, one clearly needs a detailed depth map. Creating such a map is a well known problem, with solutions that operate at real time, but the current cost of such depth (or disparity) map creation is not realistic for STB hardware. More realistic is the generation, and potential processing, of such a map as part of the broadcasting process. Indeed, such processing can assist several other operations in the STB such as subtitle placement by providing details of the correct depth for objects to be placed at.

Clearly further improvements are possible, for instance by combining the depth processing with scene analysis to identify areas of changing depth and prioritising those areas, or by performing regional shifts on the image. The goal here would be to minimise the errors over the range that is seen as being of dramatic importance.

CONCLUDING REMARKS

The success of S3DTV will be in giving every viewer a high-quality experience without eyestrain - no matter what sized sceen is being used in the home. In this paper we have explored the effect of changing screen size for a given input signal and show that this will have an unintended impact on the perceived depth range of the resultant image. It would be possible, but clearly not viable, to transmit multiple streams for specific screen sizes. We show that it is also possible to provide local screen size compensation in the home effectively and efficiently from a common input signal, and propose two alternative approaches to this technique. Clearly, further study on the acceptability of these solutions is needed, but we assert that they do provide the basis by which 3D screen size compensation may be effected at a low cost.

In the longer term, it is likely that coding techniques for S3DTV and multiview 3DTV will be based on, or utilise, depth maps or something similar. This will then open up the possibility of a full and accurate relative depth correction solution.

REFERENCES

1. Matt DeJohn, Will Drees, David Seigle and Jim Susinno, *Stereoscopic Geometry of 3D Presentations*, In-Three, http://bit.ly/d4ENXq
2. Bernard Mendiburu, *3D Movie Making: Stereoscopic Digital Cinema from Script to Screen*, Focal Press, May 2009, ISBN 0240811372
3. David Wood, *The Truth about Stereoscopic Television*, Proceedings of the 2009 International Broadcasting Convention, September 2009.
4. ITU-R Recommendation BT.500-11 *Methodology for the subjective assessment of the quality of television pictures*, 2002
5. Mark Schubin, *What is 3D and why it matters*, NCTA 2010 Spring Technical Forum, May 2010.
6. David Hoffman, Ahna Girshick, Kurt Akeley, Martin Banks, *Vergence-Accomodation Conflicts Hinder Visual Performance and Cause Visual Fatigue*, Journal of Vision March 2008, (2008) 8(3):33, 1-30. http://bit.ly/aNbkc0
7. Neil Dodgson, *Variation and Extrema of Human Interpupillary Distance*, Proc. SPIE Vol 5291, Stereoscopic Displays and Virtual Reality Systems XI, January 2004, http://bit.ly/brHuaH
8. Martin Banks, *Stereoscopic Vision: How Do We See in 3D?*, NAB 2009, http://bit.ly/bcX2DN
9. BBC, *Broadcasting First with 3D Rugby*, BBC News Website, 10th March 2008, http://bit.ly/crCGsy
10. Richard Sandomir, *In 3-D, Masters Does have Extra Dimension*, The New York Times, 31st March 2010, http://nyti.ms/bjuru5
11. FIFA Press Release, *2010 FIFA World Cup to pioneer 3D Technology*, 3rd December 2009, http://bit.ly/bbOK2k
12. Samsung Press Release, *Samsung Announces True 3D Capability for Multiple Screen Technologies*, 26th January 2010, http://bit.ly/bdNHNN
13. Sony Press Release, *Complete 3D Experience Hits UK Shops – Sony Announces Lens to Living Room Offering*, 11th June 2010. http://bit.ly/aF2FLh
14. LG Press Release, *LG Brings the First Passive 3d-Ready TV To UK Consumers*, 31st March 2001, http://bit.ly/dq51RM
15. Panasonic Press Kit, *Panasonic 3D Full HD System Launch*, March, 2010, http://bit.ly/dcDACk
16. ESPN Press Release, *ESPN 3D to show soccer, football, more*, January 5th, 2010, http://es.pn/8zeFwG
17. BSkyB Press Release, *Sky 3D to Launch with Manchester United vs Chelsea on Saturday 3rd April*, 19th March, 2010, http://bit.ly/djRmje
18. DirectTV Press Release, *DirectTV is the First TV Provider to Launch 3D*, 6th January, 2010, http://bit.ly/7xKbhA
19. Daniel Scharstein and Richard Szeliski, *A Taxonomy and Evaluation of Dense Two-Frame Stereo Correspondence Algorithms*, International Journal of Computer Vision, April 2002, Vol 27, No 1-3, pp 7-42. http://bit.ly/clRvWQ

