

multi-view autostereoscopic 3D display

Neil A. Dodgson

Professor of Graphics & Imaging
The Computer Laboratory



©2011 Neil Dodgson

This presentation was given at the Stanford Workshop on 3D Imaging on 27 January 2011.

This document is Copyright ©2011 Neil A. Dodgson. You may make a single copy for your own use but you may not make derived works. In particular, you may not use the photographs, diagrams, or text without prior permission.

Contact information:

www.cl.cam.ac.uk/~nad

Neil.Dodgson@cl.cam.ac.uk

Good afternoon. I'm Neil Dodgson, from the Computer Laboratory at the University of Cambridge.

I'm going to talk about multi-view autostereoscopic 3D displays.

There is a lot of hype about 3DTV. If you were at CES, or even if you visit your local consumer electronics store, you see a lot of high-end TVs, 3D enabled, with glasses.



But what is wrong with this picture?

This is some workmates watching the 2009 SuperBowl in 3D. They all need a pair of glasses.

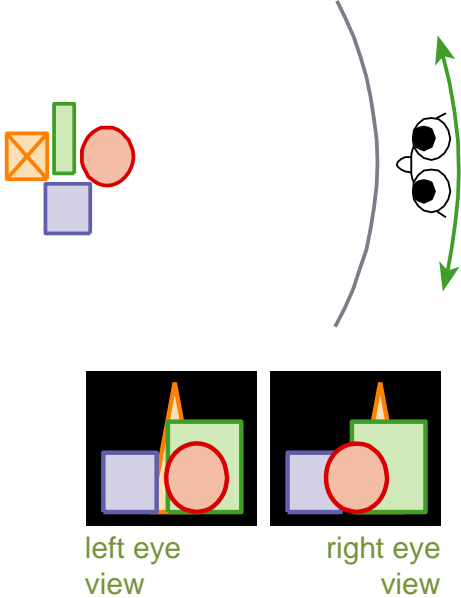
The glasses mean that watching 3D is something you do for a special occasion: watching a whole movie or watching a whole game. It is not for watching the news, or a soap opera, or anything where you want to do something else at the same time as you watch. If 3D is to be something for more than a special occasion, then we need to do without the glasses.



autostereoscopic means that we do 3D without the glasses.

stereo perception in the real world

- stereo parallax
 - each eye sees a different image of the world
- movement parallax
 - different images are seen when the head is moved
- there are an infinite number of different images





left eye view right eye view

OK. So how do we do stereoscopy without the need for glasses? How do we deliver a different picture to each eye if the viewer is not wearing glasses and we are not tracking her head?

Let's do a thought experiment. This is how we see in the real world. Each eye sees a different image of the world, so we get stereo parallax. And we can move our heads freely, so there are an infinite number of images that each eye could see and the two eyes always see just the right images to get stereopsis.

a thought experiment

- divide the viewing zone into a *finite* number of *windows*
- a single image is visible in each window
- still see a different image with each eye
 - stereo parallax
- still see different images when the head is moved
 - movement parallax



The thought experiment is to divide that continuous viewing zone into a *finite* number of *windows*.

In each window, a single image is visible across the whole window. No matter where the eye is in the window, it sees exactly the same image. As the eye moves from one window to the next, the image changes: there is a little jump from one image to the next.

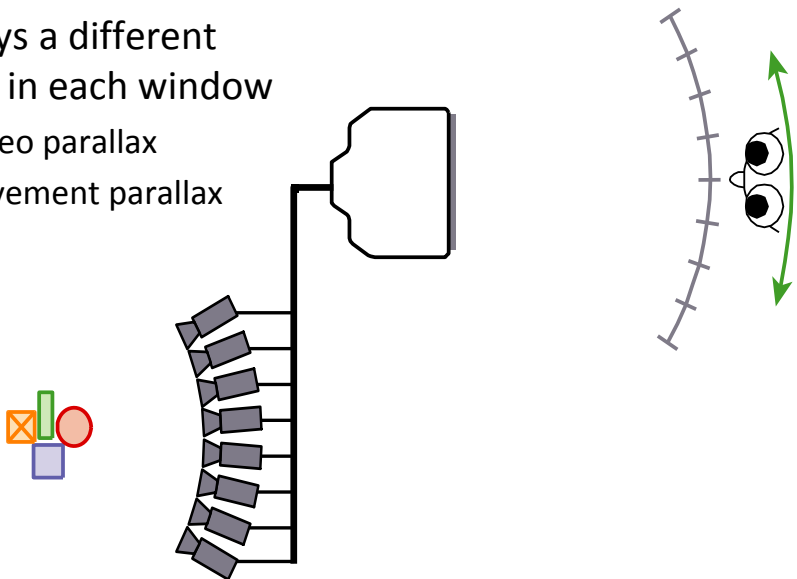
If we make the windows small enough, then each eye will still see a different image: stereo parallax.

And we still see different images when the head is moved: movement parallax

So we now have stereopsis with a finite number of distinct images. And multiple people can look at the scene from their own points of view, each getting stereo from their own point of view.

multi-view autostereoscopic display

- make a device which displays a different image in each window
 - stereo parallax
 - movement parallax



The diagram illustrates a multi-view autostereoscopic display. A central display unit is connected to a series of windows. A 3D scene with colored shapes (orange square, green rectangle, red circle, blue square) is shown behind the windows. A viewer's eyes are shown on the right, with a green arrow indicating movement parallax.

So all we need to do is build a display device that can present a different image on its screen to each of those windows.

autostereoscopic displays — many possible technologies—

- multiple projectors
- lenticular lenslets
- parallax barriers
- Fourier-plane shuttering
- retro-reflective mirrors
- half-silvered mirrors



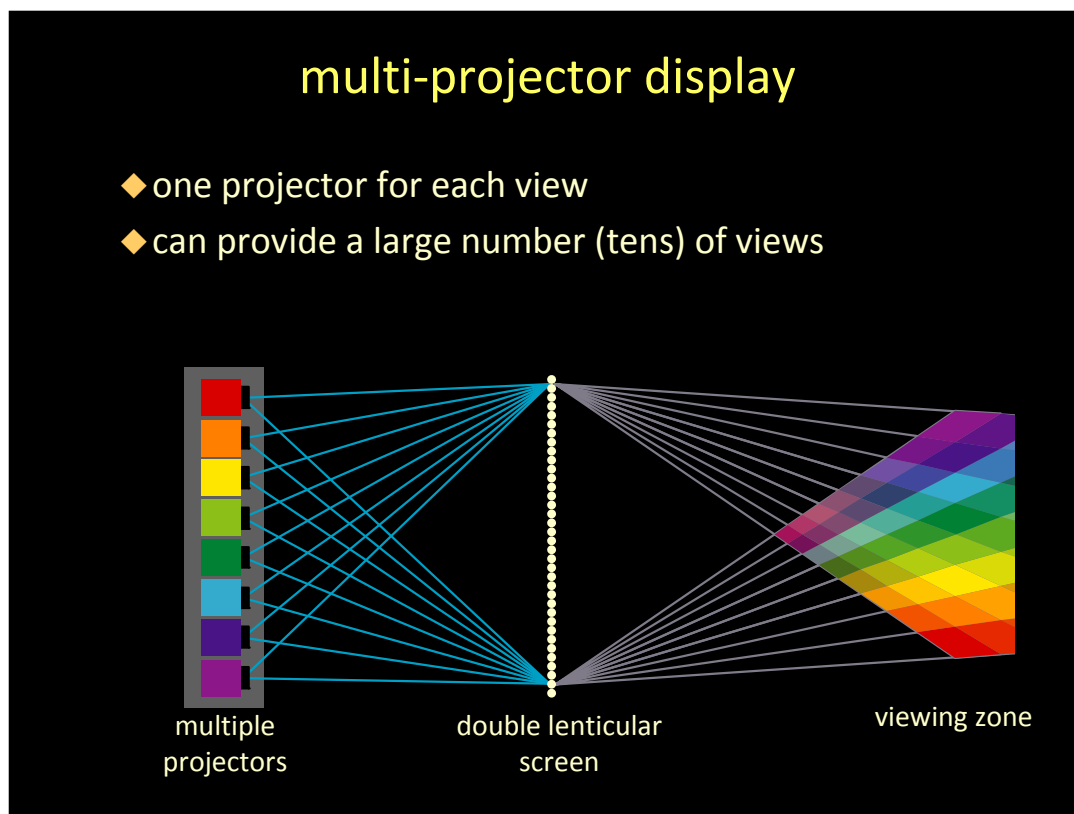
There are many ways to make a display that doesn't require glasses. All it needs to do is display a different picture to each eye, and there are a range of optical devices that can achieve that. I'm going to discuss the first four in the list, which are the three most common methods (left column) and the method which we used in the 1990s (Fourier-plane shuttering).

This is the 1999 display that I helped to build. It used Fourier plane shuttering. It produced a bright, clear image with a 50" diagonal.

This is a recently announced Newsight display based on parallax barrier technology. It has a 70" diagonal.

So we can build these displays, and you can even buy these displays

But there are fundamental limitations that mean it will be difficult for autostereoscopic technology to find its way into the home.

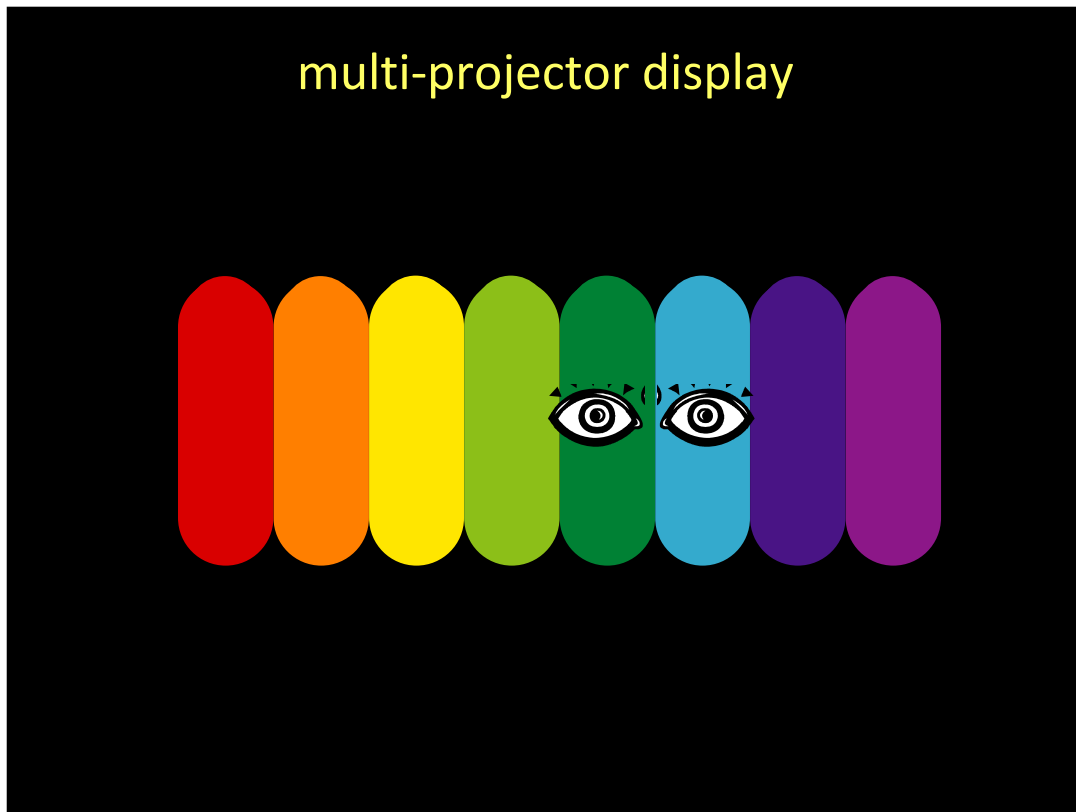


The simplest multi-view display to explain is the multi-projector display.

We take one projector per view and arrange them horizontally. Each projector images onto a double lenticular screen. That screen is the surface on which the viewer sees the image. The double lenticular screen focuses light back down into a little area in the viewing zone. If your eye is in the right place, it will see an image from exactly one projector.

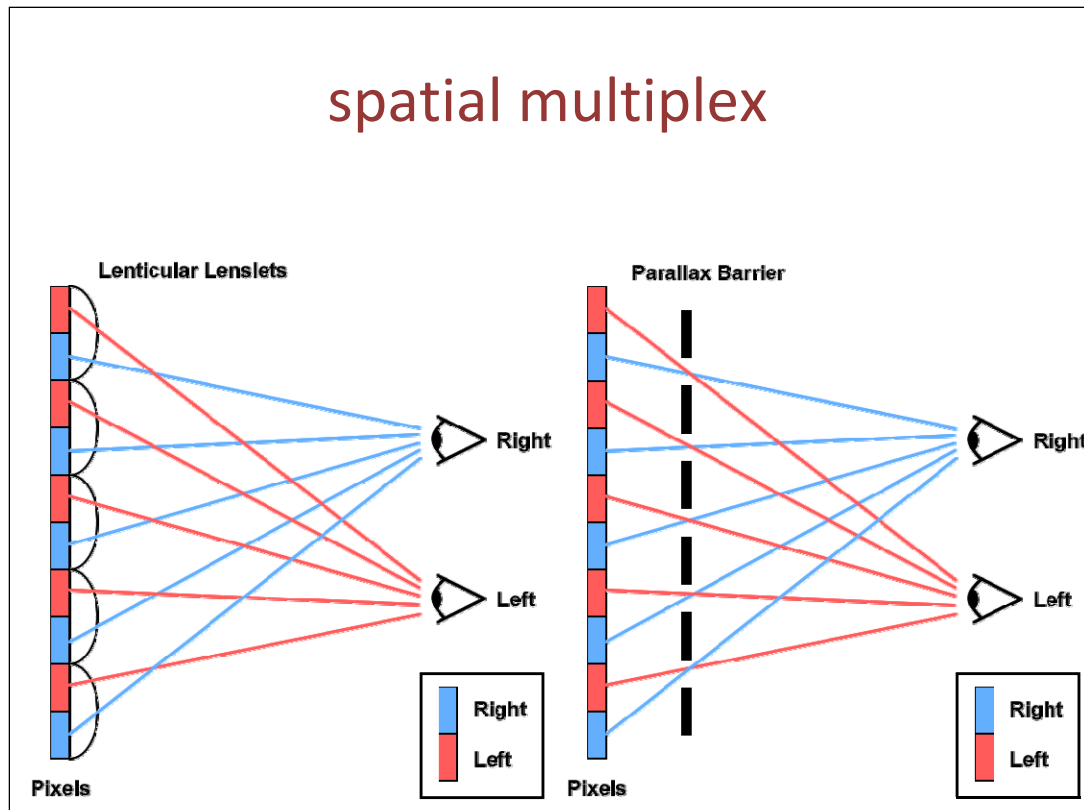
Later, we'll discuss the shape of the viewing zone, and what a viewer sees from each point within it.

Experimental multi-projector displays have been made with up to 128 views. That requires 128 projectors, which is expensive. Mounting and aligning those projectors is a time-consuming job. There have been attempts to commercialise this technology, with a smaller number of projectors, but I do not think there are any commercially-available multi-projector systems.



At the optimal distance the viewing zone comprises abutting windows, in each of which only one view is visible.

If we used a large lens, rather than a double lenticular sheet, each of the projectors' lenses would image to a circular region in space. That would mean that the eyes would have very limited positions in which they could see stereoscopically. The double lenticular sheet, by contrast, means that we get vertical stripes in space: allowing more freedom in the eye's positions. We need to add a little bit of horizontal diffusion to make those stripes abut to provide a continuous viewing space in which the eyes can move freely and view stereoscopically.




The most common way to make 3D displays is to use either a lenticular lenslet array or a parallax barrier array. These optically divide the columns of pixels into two or more sets, each visible from particular directions.

These divide the horizontal resolution of the display into two or more sets of pixels, each set visible in a particular window


conventional: vertical barriers/lenslets, up to four views

commercial: slanted barriers/lenslets, up to nine views




For two-view displays, we get these diamond-shaped zones in space, in each of which either the left or right image is visible. The nature of the optics means that a parallax barrier or lenticular display has multiple zones, alternating left, right.

So long as the viewer has her left eye in a left zone and her right eye in a right zone, she will see stereo (lower head). Unfortunately there is a 50% chance that the viewer's head will be in the wrong place (upper head): seeing the left image with the right eye and vice-versa. This gives a pseudoscopic image: inverted stereo.




We can make the lenticular lenses wider. For example, covering four columns of pixels. This produces a four view display, giving each viewer some ability to move their head (left three viewers). There is still the chance of having the head in the wrong place, seeing a strange image (rightmost viewer).

In the early 1990s, four views was thought to be the best you could do with lenticular or parallax barrier technology.



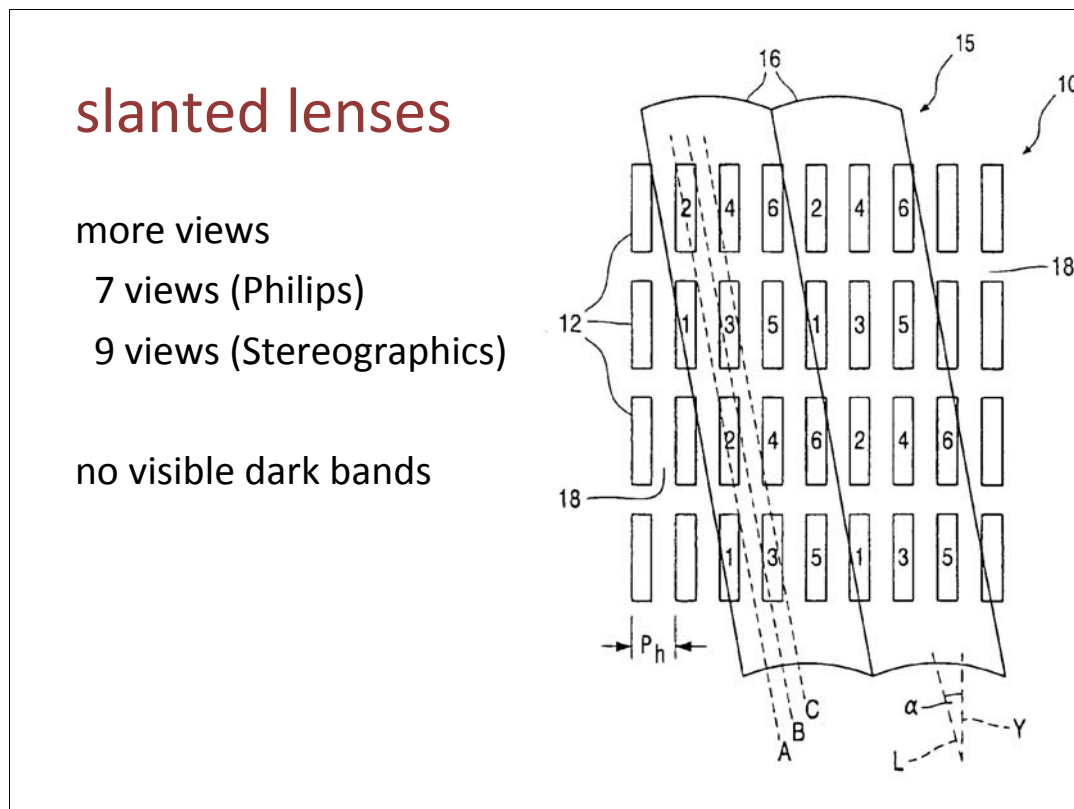
That was because the lenticular lenslets were vertical, so they split only the horizontal resolution of the underlying pixel array.

So a 1990s VGA display (640x480) became a 3D display with ridiculously low horizontal resolution.



Another problem was the fact that those lenslets magnified the sub-pixel structure out to the viewing zone.

In practice, this means that the viewer sees dark bands (the inter-pixel areas) sweeping across the screen when he moves his head. It made a lenticular display an interesting toy rather than a commercial proposition.




This all changed in the mid 1990s, when Cees van Berkel, at Philips research lab in the UK, discovered that slanting the lenticular array solved both problems. The slanted configuration meant that both horizontal and vertical resolution are divided amongst the multiple horizontally-spaced views. We thus get a reasonable resolution in both dimensions in all views. The slanted lenslets also smear out the dark bands so that they are no longer visible.

I remember the first time I saw the Philips display at their research lab. Because they had not yet filed a patent, Cees could not tell us how it worked. My team spent several fruitless hours trying to figure it out, but Cees didn't let on until the patent was in place.


Philips commercialised a 7-view display based on this idea. It was discontinued a couple of years ago, just before 3D movies really took off. They are considering re-introducing a product.

Stereographics produced a 9-view version. They avoided leasing Philips' patent by citing a much earlier patent that did much the same thing, demonstrating the well-known fact that some good ideas get invented before their time, get forgotten, and then get re-invented.



So let me now talk about our Cambridge display. I do this because our display is the basis of my analysis of multi-view displays and, while that analysis applies to all types of multi-view display, it is easiest to understand in terms of our display.

These are the original research prototypes: 10” colour display at left, 25” monochrome display at right. They have between 6 and 16 views, depending on the version. The image on the left hand screen is of John Moore, the genius engineer who did the detailed electronic design and built the displays. He is taking the photograph that you see, himself being videoed by the multi-view camera array (top centre) and imaged onto the screen, in multi-view 3D, in real time.

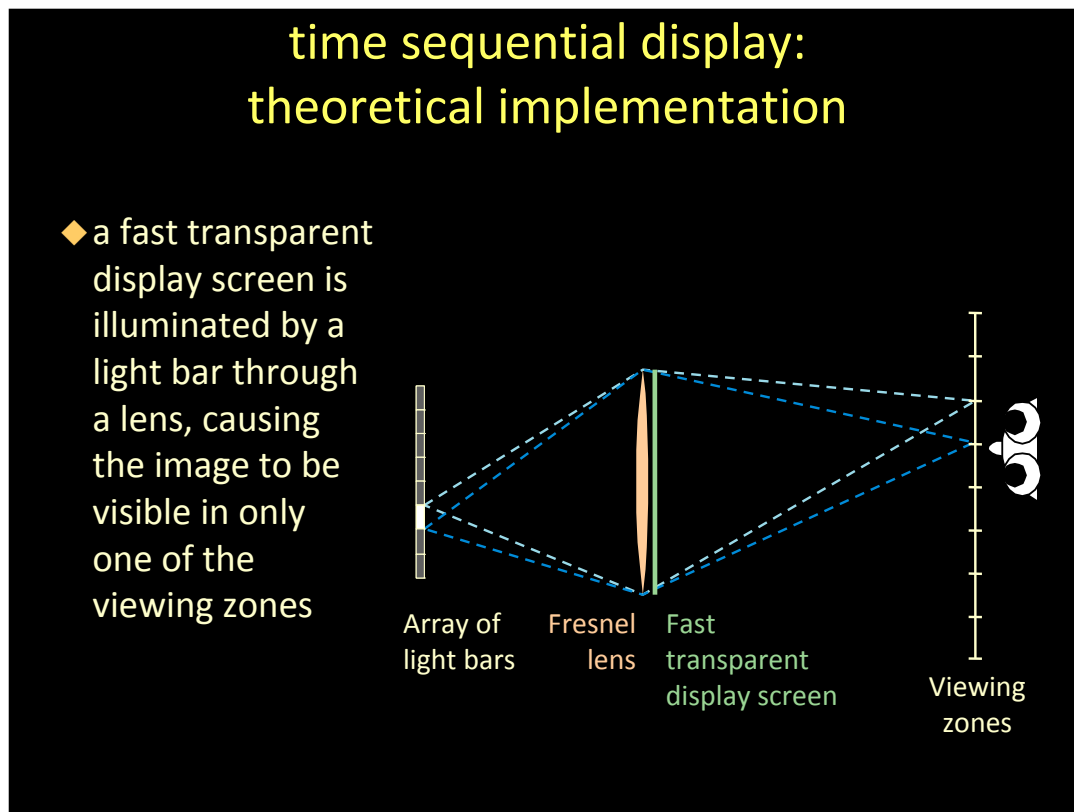


In 1995, we moved this out of the university and spent two years developing a large-screen prototype for video arcades.



And then developed further, to produce this 50" 3DTV. You won't have seen any of these on the market, because this is as far as development got before the money ran out.


This display produced a beautifully clear 50" diagonal stereoscopic image viewable under normal lighting conditions (notice that all the lights in the room are on, in contrast to many displays that require a darkened room). I have yet to see an autostereoscopic display that equals it in visual quality.



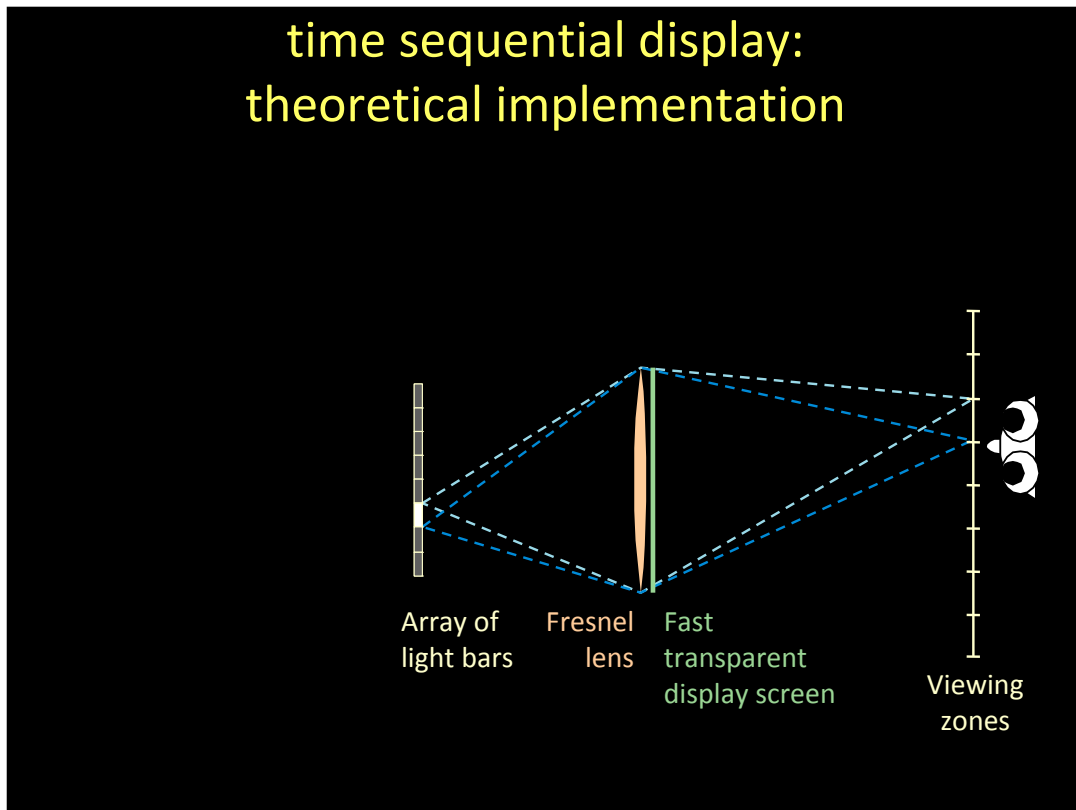
Adrian Travis' original concept is elegant. You display the multiple images sequentially on a fast transparent display screen. For 8 views, you need a refresh rate of at least 480Hz. Behind the screen is a lens. Some distance behind that is a set of vertical light bars. You turn one bar on. It illuminates the whole screen but the lens means that illumination is only visible in a certain zone in space. In our example, the viewer's right eye sees the image on the screen while the viewer's left eye sees a black screen.

time sequential display: theoretical implementation


- ◆ synchronising the display with the changing of the light bar causes a different image to be visible in each zone



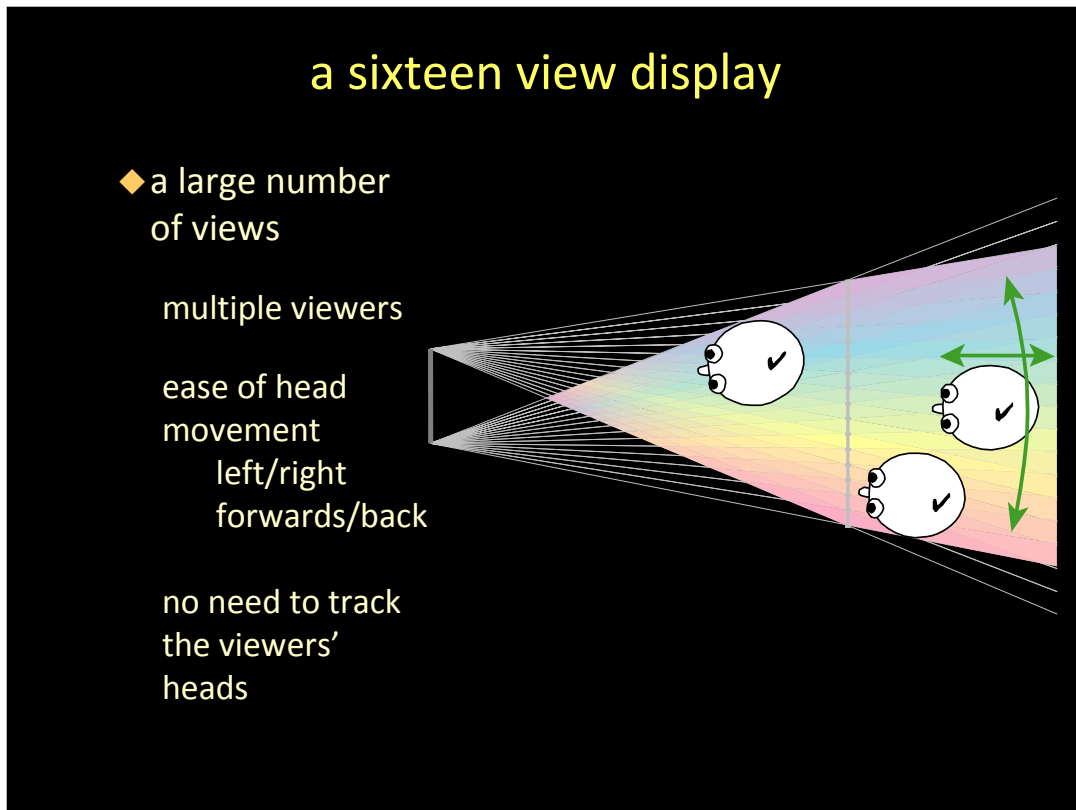
Synchronising the display with the changing of the light bar causes a different image to be visible in each zone, so each eye sees a different image and different images are visible when the head is moved. The display does not need to know where the viewer's head is, so long as the viewer's two eyes are in illuminated zones. Multiple viewers can look at the display, each seeing stereoscopic 3D from their own correct point of view.



Unfortunately, there is no technology that can produce a fast transparent display screen that can refresh at the required rate. Adrian thus shelved his idea for a couple of years.




But he then realised that he could build an equivalent system using a CRT. CRT refresh rates can be over 1kHz, so they are fast enough. The CRT image is projected onto the Fresnel lens. In the Fourier plane of the projection lens we put a set of shutters. At any given time, one shutter is transparent and the rest opaque. This gives an effect equivalent to the theoretical design. Today you would use a projector rather than a CRT.



This allowed us to build displays with a large number of views. We then began to ask: what is the optimal number of views? and what is the optimal density of views?

how many views?

- ◆ how many between the two eyes?
 - *if a view must be no more than a pupil wide*
3mm per view, 20 views between eyes
 - *practical compromise*
21mm per view, 3 views between eyes



If we talk to holographers, they tell us that each viewing zone must be no wider than the entry pupil of the eye. That makes each view only 3mm wide, giving 20 views between the eyes (lower set of bars). Of course, you need views outside the two eyes, so that you permit some head movement. This means that you need a display with at least 30 views, and probably more like 60 views. This is not practical with current technology.

However, if we go to the other extreme and have just one view between eyes, that is, each viewing zone being 63 mm wide, then we have a different problem. While each eye does see a different view, we get a dramatic change in view as we move from one zone to the next. Our experience is that this jump is so dramatic as to be disturbing to the viewer. Also, the viewer does not get much movement parallax effect when they move their head. Instead they strongly feel as if they are jumping from one view-point to the next.


As a compromise, my team found 21mm per view to be good (upper set of bars). This is three views between the eyes. The differences between adjacent views are sufficiently small that they are not disturbing as you move from zone to zone, and the views are close enough together that you get a fairly smooth effect as you move your head. Indeed, the combination of the two effects leads me to think that the quality of the effect improves quadratically with the number of views between the eyes.

Stereoscopic Displays & Applications

Reduced-view super multi-view display

Junya Nakamura, Yasuhiro Takaki

Tokyo U. of Agriculture & Technology



As a side note, on 25 January 2011, I saw a prototype super multi-view display at the annual Stereoscopic Displays & Applications conference (www.stereoscopic.org). This had viewing zones 2.6 mm wide. There were eight zones provided for each eye. That is: each eye had to be inside a region only $8 \times 2.6 = 19.2$ mm wide, so head movement was extremely restricted. However, it demonstrated to me a fact that I have long heard but never experienced: if the viewing zones are this narrow, then there are at least two views entering each pupil. This means that you can focus at distances other than the screen, thus solving the problem of accommodation/convergence conflict. The effect is similar to that you would get with a white-light hologram.

However, it will be some time (if ever) before such displays are a practical proposition.

The gentlemen in the photograph are, left to right, Andrew Woods (Curtin University, Australia, chair of Stereoscopic Displays & Applications), Junya Nakamura (TUA&T) and Yasuhiro Takaki (TUA&T).

For more information on the Stereoscopic Displays & Applications conference, visit www.stereoscopic.org

The next conference is planned for 23-25 January 2012, in San Francisco.

how many views?

- ◆ how many views total? (with 21mm per view)
 - *single viewer, 189mm (three times inter-ocular distance) lateral movement*

9 views


25" autostereoscopic display 1.5m optimal distance

If we assume 21mm wide views, then how many views do we need in total?

For a single viewer, we should allow some lateral movement. It is not clear how much should be allowed, but about three times inter-pupillary distance seems to be reasonable. This gives good space for head movement for a seated viewer. That calculates to nine views, which matches nicely with the seven and nine views produced by Philips and Stereographics lenticular displays. We appear to have agreement from three different sources that this number of views is some sort of optimum trade-off.

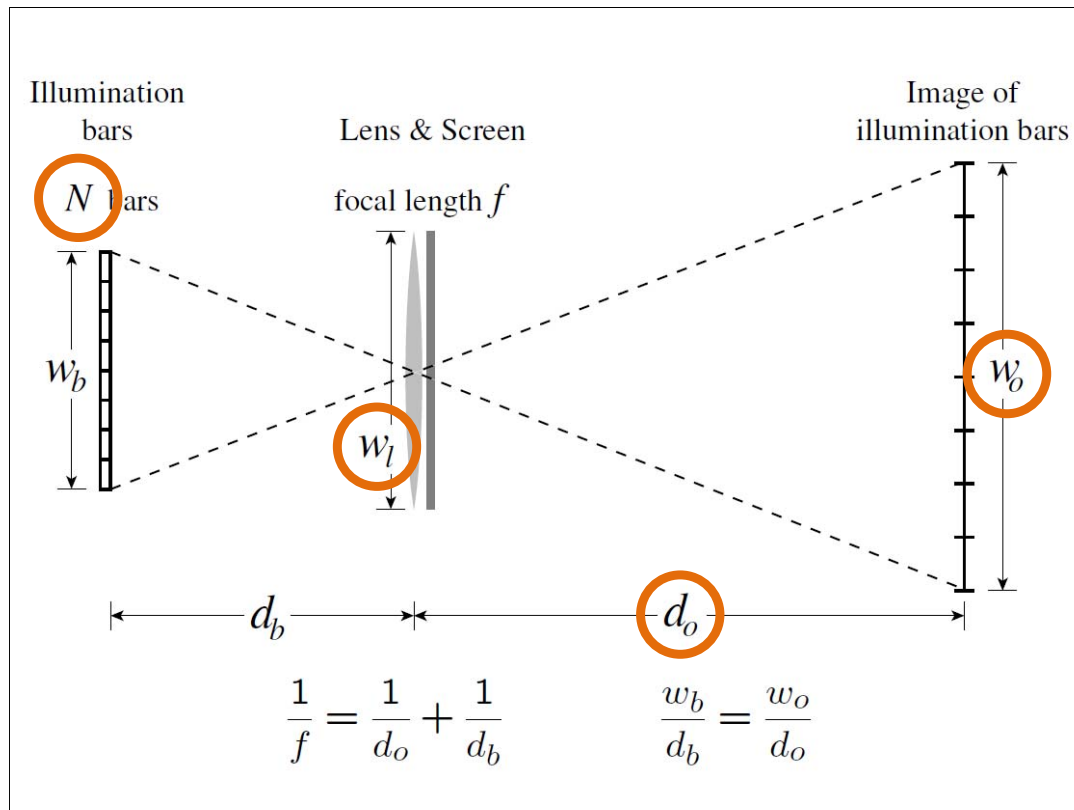
how many views?

- ◆ how many views total? (with 21mm per view)
 - *single viewer, 189mm lateral movement*
9 views
 - *two viewers, side by side*
28 views



If we want two viewers to sit side-by-side, as my team did for our 1997 25” display, then we need about 28 views. That gives each viewer some head movement without banging shoulders with their colleague. It leaves a lot of unused views in the middle, but those can provide stereoscopic viewing to a third person standing behind the other two.

The fact that the third person is not at the optimal distance leads us to ask what is visible on the screen to viewers who are at distances other than the optimal. This is also important for home use where viewers are unlikely to sit at exactly the optimum.



To answer this question, I analysed the Cambridge theoretical display design. This analysis generalises to all types of multi-view autostereoscopic display.

It turns out that only four parameters determine the characteristics of a display, as seen by the observers. These four parameters are all in “viewer space”: that is, in front of the display. One can specify a display by these parameters and then use whatever optical system is desired to produce the specified display.

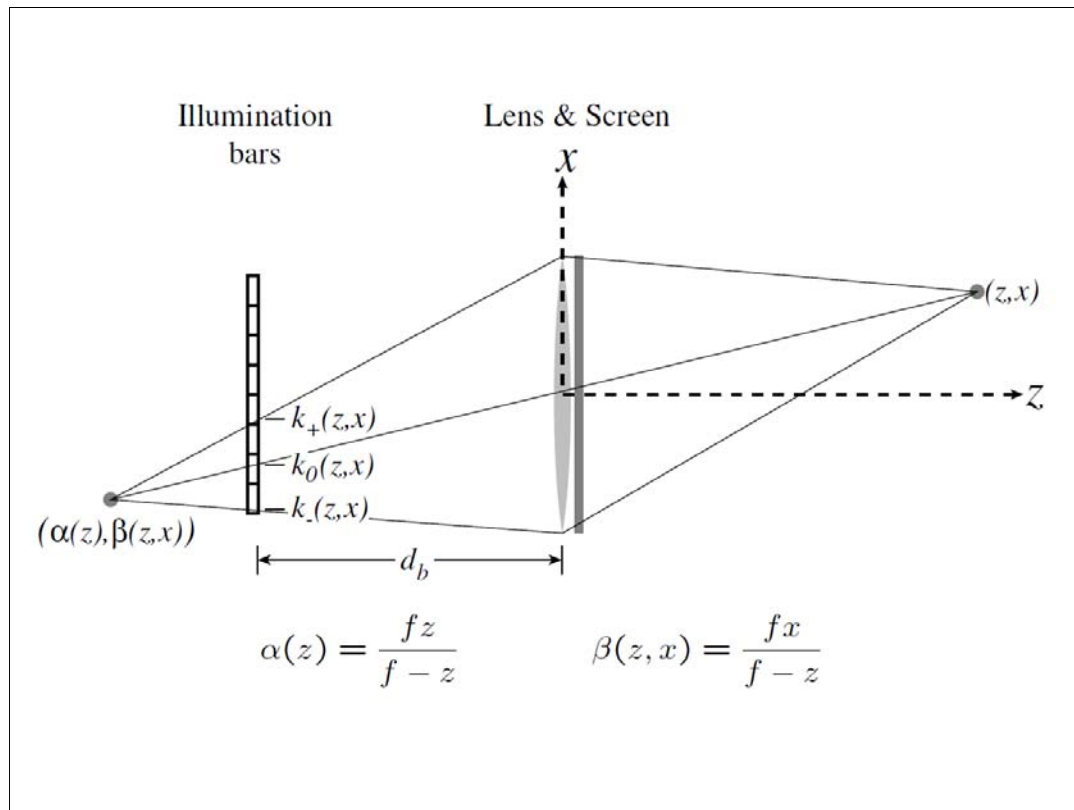
The parameters are:

The width of the screen, w_l

The optimal viewing distance, d_o

The width of the overall viewing zone at the optimal distance, w_o

The number of viewing zones within that overall width, N

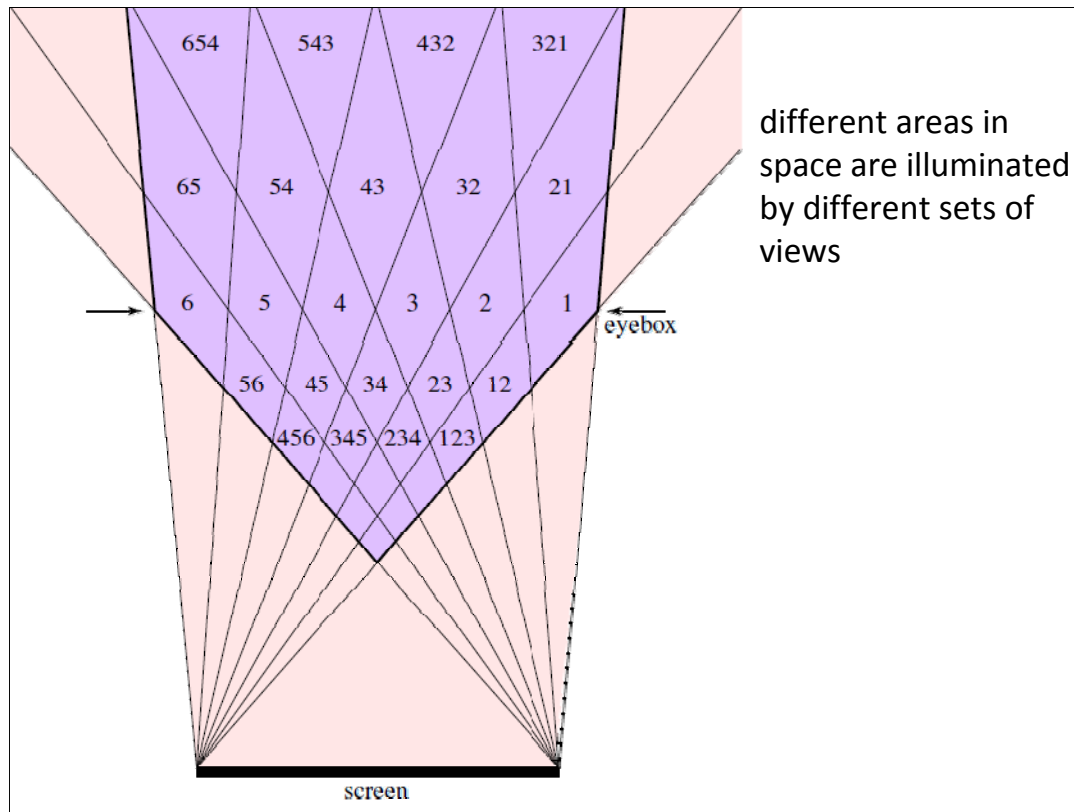


We can calculate which illumination bars illuminate which parts of the screen for an arbitrary position in front of the display. This allows us to work out what the viewer sees from anywhere in front of the display.

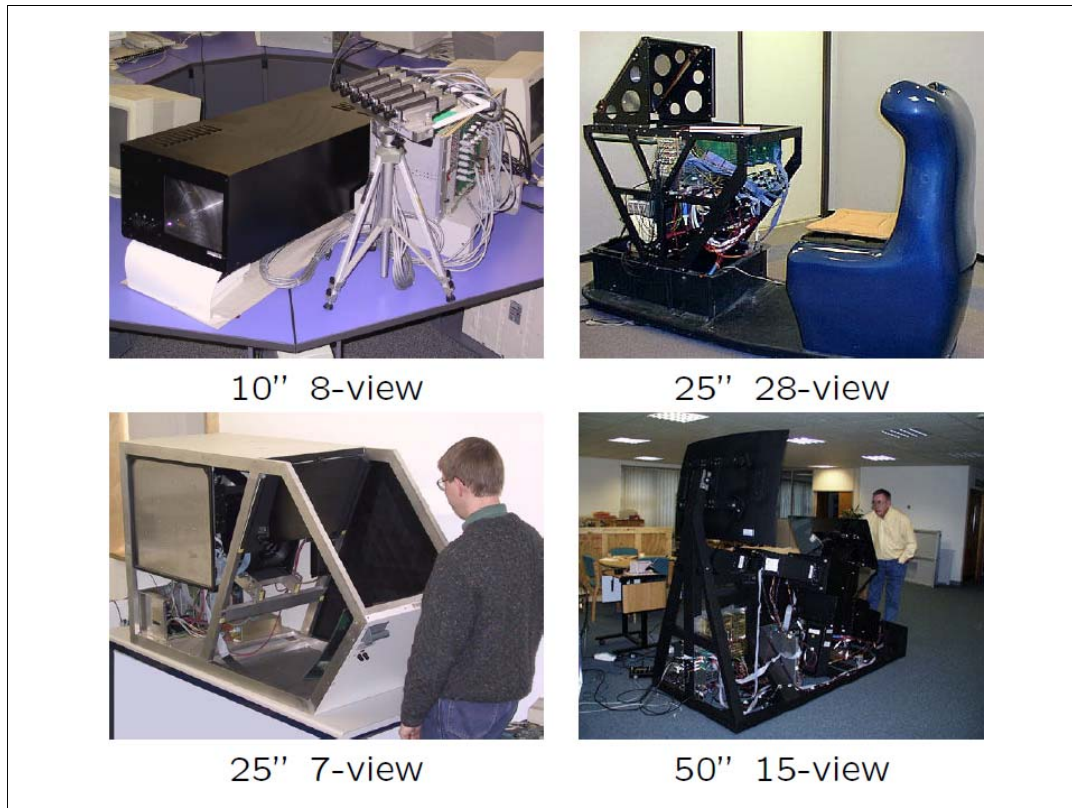
The overall analysis can be found in these two paper:

2002 “Analysis of the viewing zone of multi-view autostereoscopic displays”, N. A. Dodgson, in “Stereoscopic Displays and Virtual Reality Systems IX”, *Proc SPIE* **4660**:254–265, ISSN 0277–786X

1996 “Analysis of the viewing zone of the Cambridge autostereoscopic display”, N. A. Dodgson, *Applied Optics: Optical Technology & Biomedical Optics* **35**(10):1705–1710, 1 April, ISSN 0003–6935

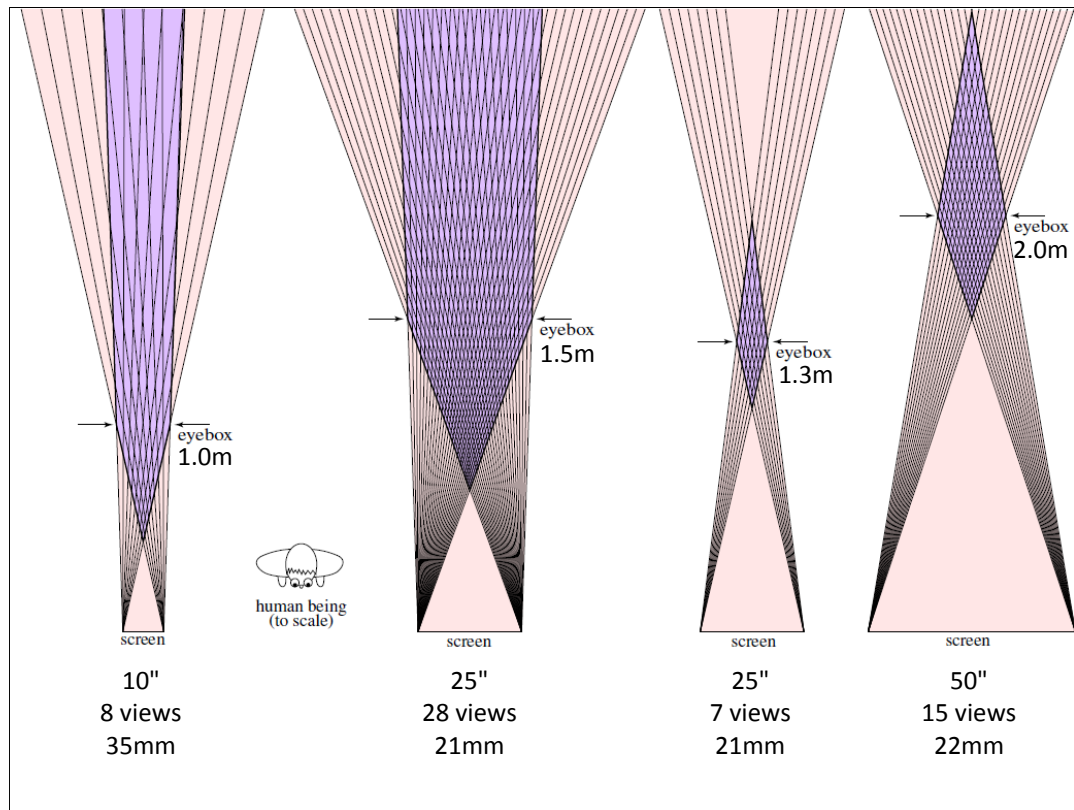


It transpires that, at the optimal distance (the “eyebox”), an eye sees the whole screen showing exactly one view. Forward or back from the optimal distance the image is made up of parts of various views, with increasing numbers of views contributing as we get farther from the optimal distance.



Let's now look at four practical examples, where I was able to compare the theoretical results to the actual viewed imagery.

These are the four Cambridge displays: they have different sizes (10", 25", 25", 50"), and different numbers of views (7, 8, 15, 28).



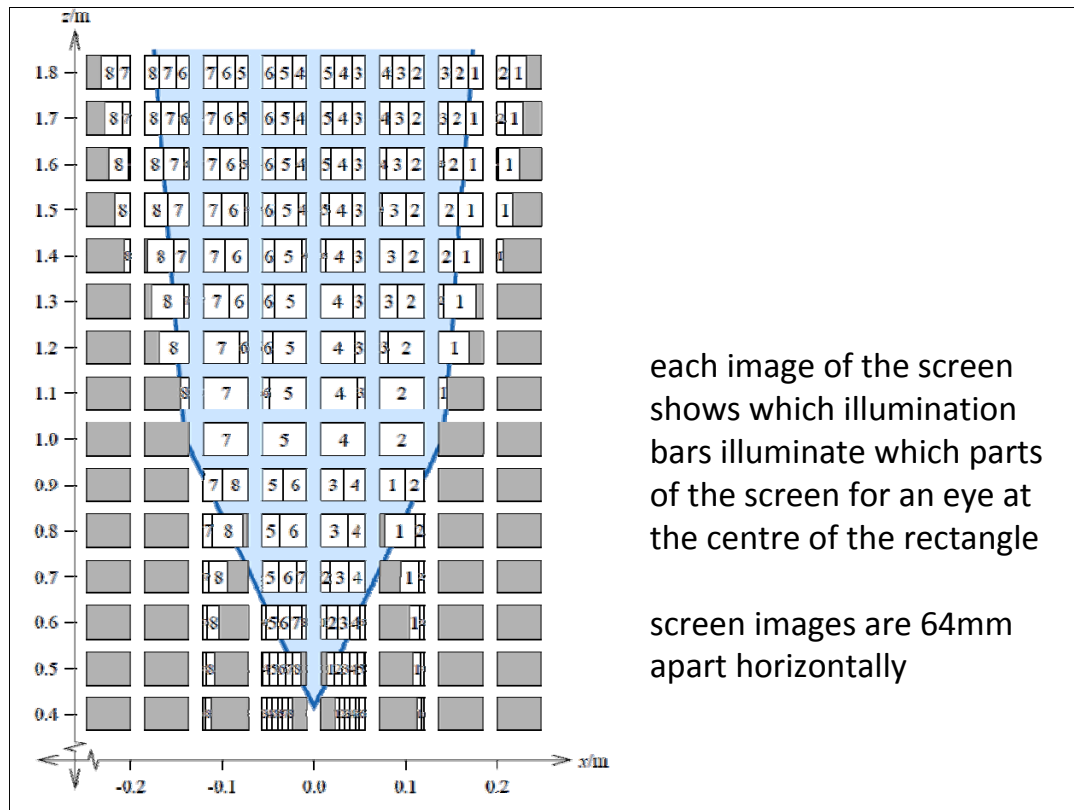
For a small screen (10", 8 views), we get a viewing zone that allows multiple people to stand one behind the other to see stereoscopically. We managed to get up to six people simultaneously viewing this display in 3D. Normal operation had one person, with a second occasionally looking over the shoulder of the first. This is typical professional use of a display: single operator with occasional collaborative work.

If we enlarge the screen but not the number of views (25", 7 views) then the viewing zone becomes a diamond. This is essentially a single-viewer system as it would be hard for a second person to have their head in a position where they can see stereoscopically without having the first viewer's head in their way.

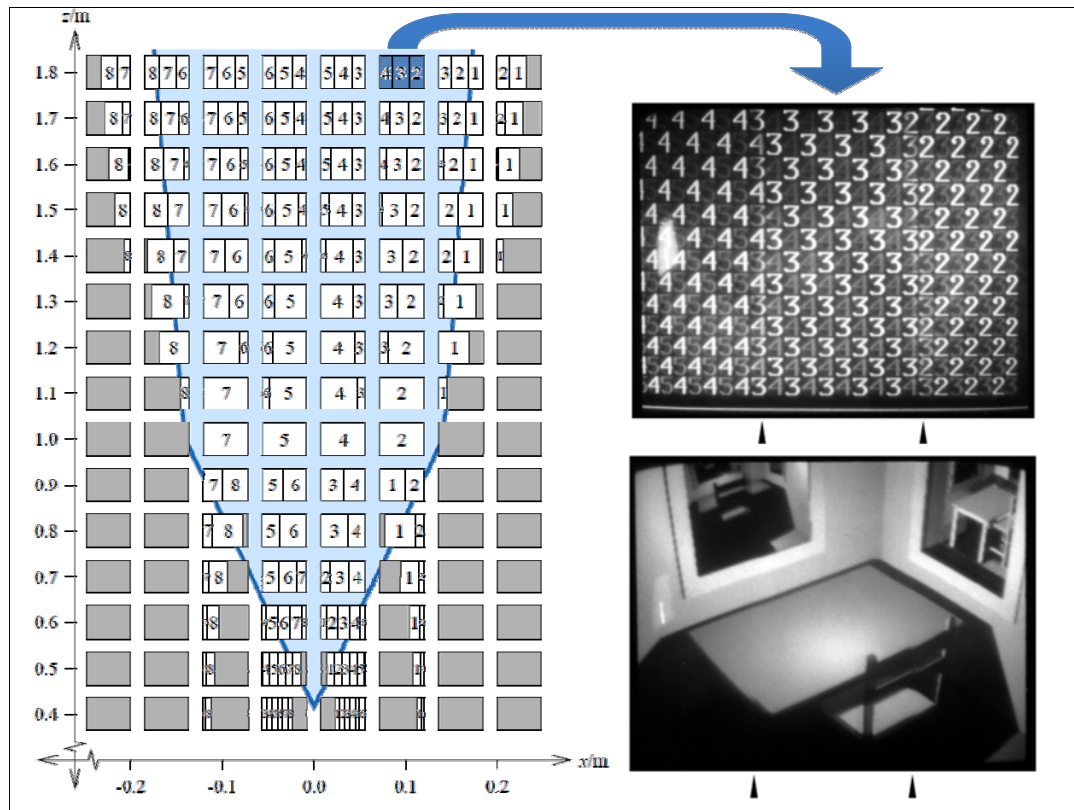
If we enlarge the number of views (25", 28 views) then things improve enormously. We can now have several people all seeing stereoscopically.

It is obvious that increasing the number of views increases the size of the viewing zone.

This analysis shows that increasing the screen size can quite dramatically reduce the size of the viewing zone.



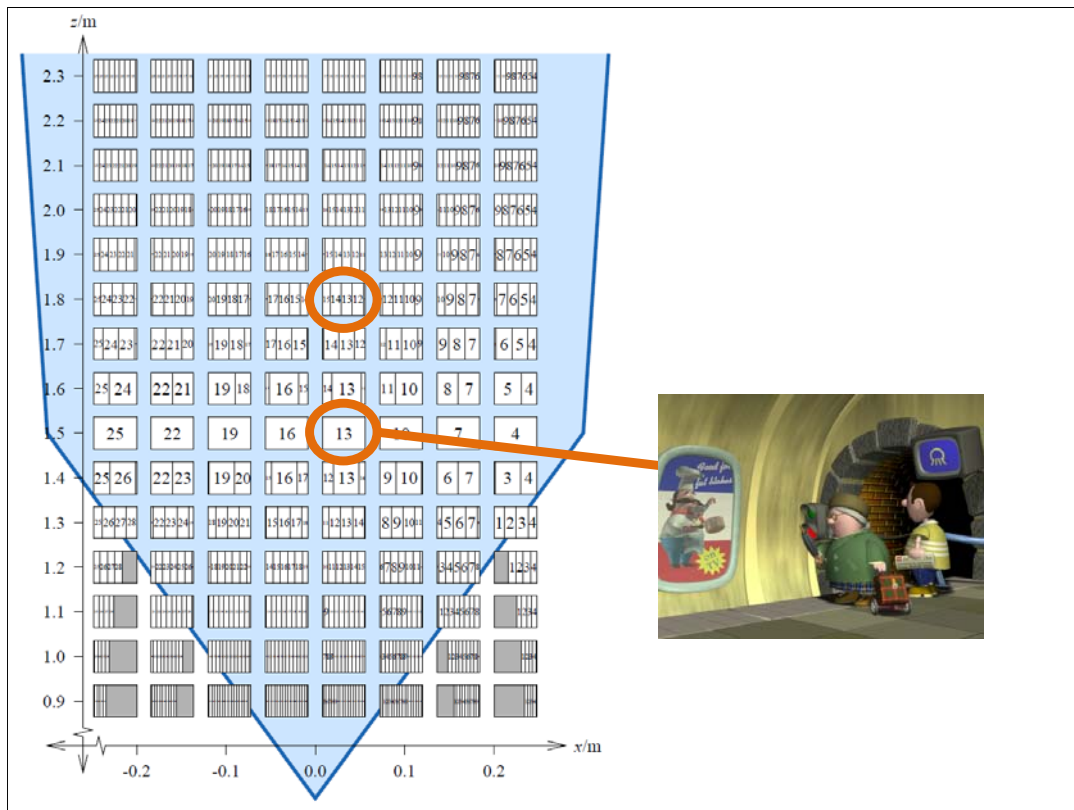
These charts proved useful in understanding what the viewers saw on the screen. At the optimal distance, an eye sees exactly one view. As we move forward or back from that distance, you can see that the image on the screen is made up of stripes taken from the correct portions of different views.



These photographs off the 10" display's screen show what is really seen.

As predicted, the image (upper image) is made up of parts of views 4, 3, and 2 (you can also see some ghosting coming through from adjacent views).

If we look at a real image (lower image) rather than the test card (upper image), then we find that we can barely see the joins between the pieces of different views, because adjacent views are so similar in content, being taken from viewpoints close together.



This is the 25" 28-view display. You see that you do not need to move far from the optimum distance for there to be multiple views contributing to the perceived image.



This is the view from the optimum distance,

(Granny sequence, View 12, Frame 21)




This is the view if we step back from the optimum distance. The orange lines show you where the joins are between the different views. At first glance, you cannot tell that this image is drawn from multiple different views.

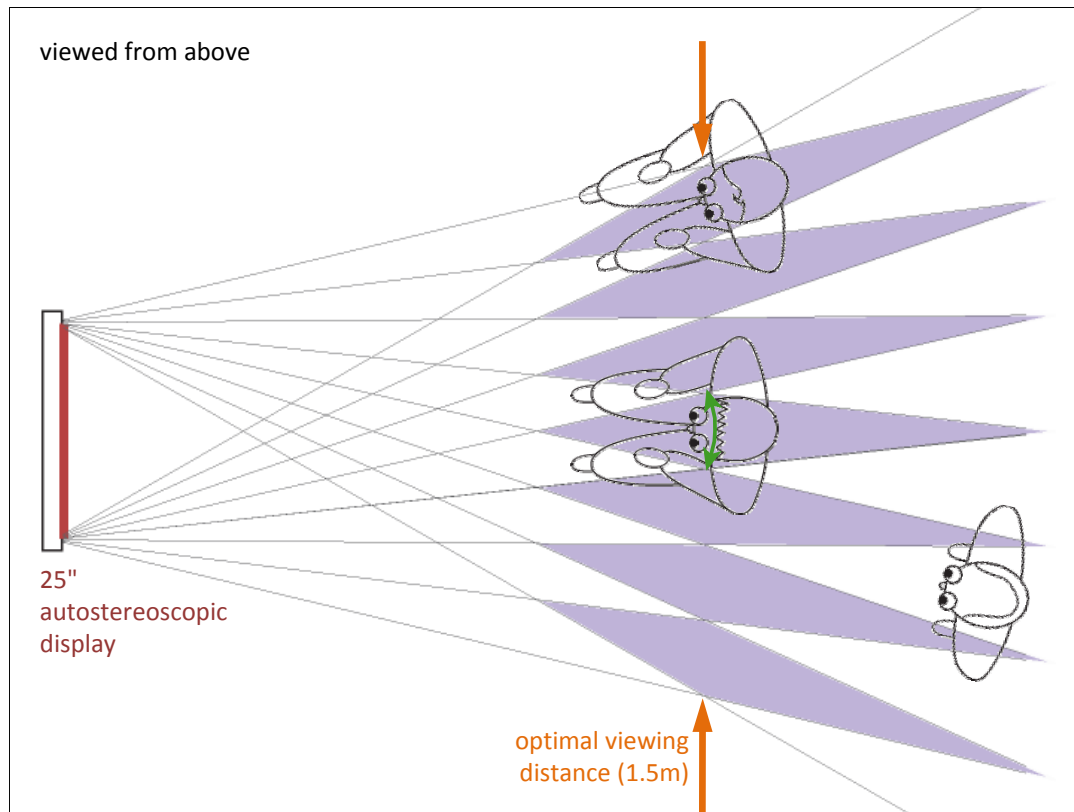


In practice, the display had a little blurring as you moved from one view to the next, this makes the joins even harder to see.

Thus, although the image is made up of parts of multiple views, it has acceptable quality. This is because the viewing zones are close enough together that one view is very similar to the next. In this case the camera has moved only 21 mm to the side.



Consider now, the 25", seven-view display. This is similar to the commercial displays produced by Philips and Stereographics. However, our display had only one viewing lobe, whereas theirs have multiple lobes. Ours is essentially a single-viewer display.



Any autostereoscopic display produces a viewing zone in space. Your head must be inside this zone in order to see 3D. You can move forward, or backward, so long as your two eyes both stay in the zone. You can move your head side to side, so long as both eyes stay in the zone. This is quite restrictive and only allows one viewer.

Most commercially available displays produce multiple zones. This allows multiple viewers (two seated viewers are shown). But there is still an optimal viewing distance, which means that you need to set up your living room so that the sofa is at the correct distance from the screen.

It also introduces the strong possibility that someone will stick their head in the wrong place (viewer standing at right) getting an inverted 3D effect and a head ache.

People need to understand the technology to know where they should put the chairs in their living room and people will need to learn how to spot when the 3D is wrong and how to move themselves to get into one of the zones. Because of this, glasses-free systems are going to be even more difficult to get into homes than with-glasses 3DTV.

Lessons learnt over 20 years

- 3D is **needed** only in technical niches
 - visualisation of complex 3D structures (sci/med)
 - remote manipulation (robots)
- 3D **might** be useful elsewhere
 - advertising
 - games: video arcades, computer games
 - entertainment: theme parks, TV, gambling



We spent five years commercialising our technology.

We discovered that 3D is only needed in two particular places: visualisation of complex scientific and medical structures; and remote manipulation. These are niche markets.

We investigated other markets, video arcades and theme parks in particular. We found it difficult to make a case that the 3D experience is of sufficient benefit to be worth the extra expense.

Sources and Acknowledgements



DigiTimes (Newsight 70" display)
copyright owner unknown
www.digitimes.com/photogallery/ShowPhoto.asp?ID=4464



Fashionably Geek
© Nicole Lindner
fashionablygeek.com/dresses-and-skirts/3d-glasses-dress/



Flickr (White House photostream)
US Government
www.flickr.com/photos/whitehouse/4291192697/



3d-forums.com
copyright owner unknown
www.3d-forums.com/autostereoscopic-displays-t1.html



3d-forums.com
copyright owner unknown
www.3d-forums.com/autostereoscopic-displays-t1.html



Unless noted in the left hand column, all images and diagrams in this presentation are © Neil A. Dodgson



This presentation and its Powerpoint file are copyright ©2011 Neil A. Dodgson, all rights reserved

This presentation and its Powerpoint file are copyright ©2011 Neil A. Dodgson all rights reserved.

This PDF handout is copyright ©2011 Neil A. Dodgson. You may make a single copy for your own use but you may not make derived works. In particular, you may not use the photographs, diagrams, or text without prior permission.

Contact information: www.cl.cam.ac.uk/~nad
Neil.Dodgson@cl.cam.ac.uk