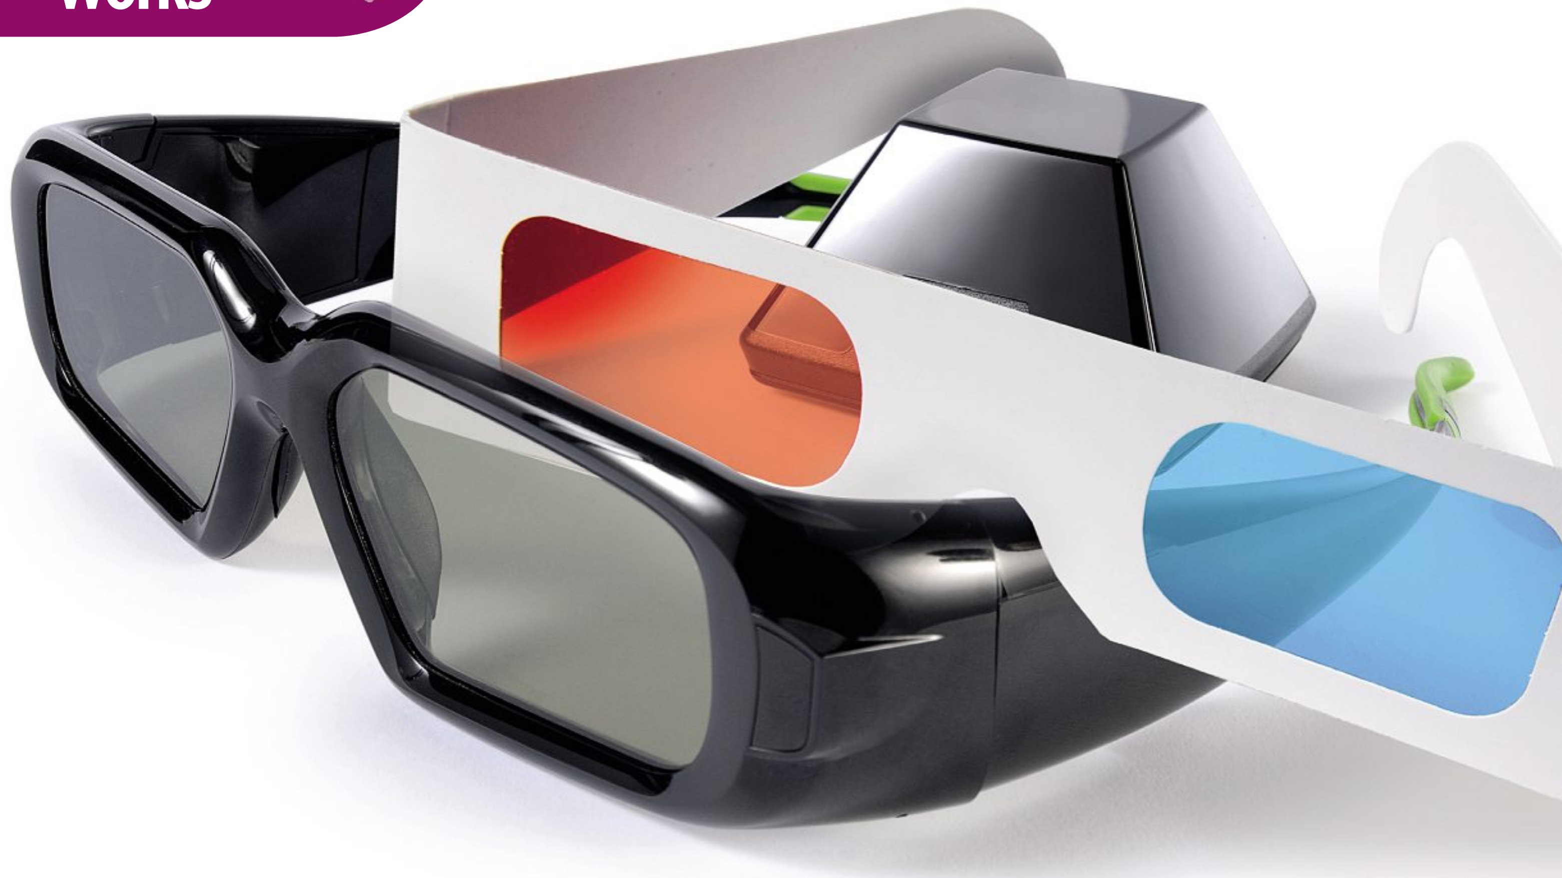

How 3D TV works

3D TV has become a huge trend. **Julian Bucknall** explains how it fools your eyes

Over the past 12 months or so, a new technology has evolved. Like most new developments, it has its detractors and its proponents. Some say it's harmless fun and the naysayers should lighten up, while others warn against using it for extended periods of time. Yet others say it's just a gimmick that manufacturers are using to part the techno-gullible from more of their money. This technology is 3D TV.

Let's start with some history. In the 1950s, there was a craze for 3D movies. Not very well-made ones by any means – most were sci-fi or horror B-movies – but they certainly made a splash. Film producers were trying to get their audiences back into the cinemas and away from their new-fangled TVs, and movies like Creature From the Black Lagoon and It Came From Outer Space served that role well.

These 3D movies (the first was a 1922 flick called The Power of Love) used a red/green anaglyph dual-strip system. Before discussing what this means, let's take a look at how we see in three dimensions and perceive depth.

Stereoscopic vision

Our eyes are roughly 2.5 to 3in apart. This separation means that the image each eye receives is slightly different. The light from distant objects reaches each eye roughly in parallel, whereas the light from nearby objects travels at different angles (the nearer the object, the more different the angles). This is known as convergence. The other process that's going on is focusing. When looking at distant objects, the lens in the eye is relaxed (or rather, the muscles that squeeze the lens are relaxed). The closer the object, the more work the lens has to do to keep it in focus. The brain uses all this effort, plus the image recorded by the light-sensitive cells in the eye (the rods and cones) to produce depth perception.

When we're out and about, walking around, we're unaware of the amount of work that's going on to stop us accidentally walking into doorframes or walls. The eyes are continually feeding information to the brain, which it interprets as 'this object is close, that one is further away'. In essence, the convergence and focus points are equal for scenes viewed in the real world. When we look at a normal TV screen or a monitor, there is no depth perception – our eyes are just focused on the screen, and it's as if we're simply looking at a flat object (which we are, of course). There's no convergence needed for the 2D image on the screen either – it's just flat. So how do we turn it into something with depth?

The early 3D movies made use of convergence (and ignored focus). If the camera recorded the same scene via two lenses positioned 3in or so apart onto two separate film stocks, then the two films could be played back in sync – one film for the viewer's left eye and the other for the right eye. But how do we ensure that each eye only sees what it's supposed to?

Early techniques

Back in the '50s, the answer was to play back the black-and-white film in two different colours on the same screen. The film for the left eye was blue (or cyan to be more precise) and the film for the right eye was red. If you looked at the screen, you'd see the scene blurred between red and cyan, but if you looked at the screen wearing glasses where the left lens was red and the right one cyan, you'd see something completely different.

The red lens would absorb all the red light hitting it and would only let through the cyan light. The cyan lens would let through the red light and absorb the cyan. Each eye would therefore only see the scene in the colour meant for it, so the left eye would see the left film and the right eye the right film. This system is known as the anaglyph technique, and is a passive system. It works well for black-and- ▸

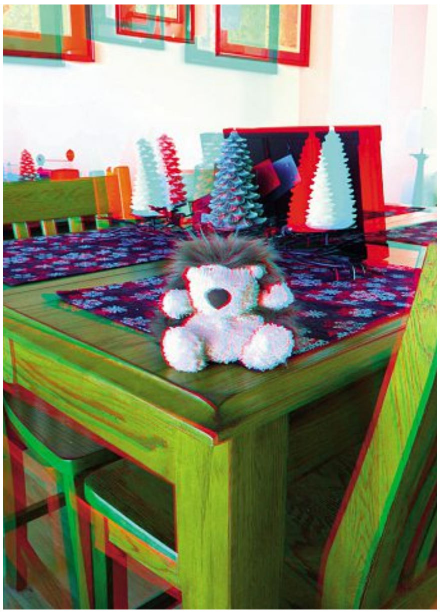

▲ **Figure 1: A red/cyan anaglyph of a 3D scene. To see it properly, you'll need red/cyan glasses.**

▲ **Figure 2: A cross-eyed stereogram. To view it, cross your eyes until you see four images, then uncross them a little until the two in the middle overlap and the scene should pop out in 3D.**

white movies, since there's no colour in the scene to be incorrectly absorbed and confuse the viewer. You soon forget about the colour cast. For an example of an anaglyph image, have a look at Figure 1 while wearing a pair of red/cyan glasses (available cheaply on eBay).

Because the light reaching the eyes obeys the 'distant objects send light in parallel, near objects at an angle' rule, the brain can perceive an illusion of depth though convergence. However, the eye is only able to focus on the screen – there is nothing else there to focus on. A 3D movie will show things 'closer' and 'further away', but we can't focus on whatever we want to – we can only see in focus what the director wants us to concentrate on. For shock value, this generally means objects that seem to come close to the viewer's face.

This difference between the convergence and focus points in 3D movies means that you're likely to experience eye strain and headaches if you watch something in 3D for too long, because your eyes are trying to do a lot of work that isn't necessary.

Polarised light

Moving back to 3D movies, the next big invention was the use of polarised light. Polarised light vibrates in a single plane, whereas the light waves in normal sunlight, for example, oscillate about many planes – some horizontally, some vertically, most in between. The lenses in polarised glasses only let through light in a single plane, which is a handy way of reducing the amount of light that reaches your eyes in bright sunlight.

This time, the projectors display the left and right image streams using polarised light (the projectors essentially have big polarised screens in front of them), with the left images shown with horizontally polarised light, and the right with vertically polarised light. The viewer wears glasses with the left lens geared to horizontally polarised light and the right to vertically polarised light. Each lens only lets through the light with the correct polarisation for that eye. Providing the viewers keep their heads vertical, they'll see a 3D effect because each of their eyes sees a different set of images. Again, it's all about convergence rather than focus, so the same drawbacks (eye strain and headaches) can appear. However, this time there's no colour cast to the movie.

This polarised light system first appeared in the early to mid 1950s, and quickly supplanted the old-fashioned anaglyph (two-colour) system, which has since been relegated to static images rather than films.

On the small screen

Although this system promises a great deal, it's extraordinarily difficult to convert for use with TVs in the home. Since there's no projector (which can easily be modified to produce polarised light), you would have to coat the screen with some kind of polarising film first, which would cut down the light output. So what can we do instead? Cut to a couple of years ago, when plasma and LCD screens finally became fast enough to work with a new method of creating a 3D effect.

With this new method, you're still going to be wearing a special type of 3D glasses to watch something on the screen, but this time the glasses are active – not passive like in the anaglyph and polarised systems. What this means is that the glasses work in concert with the TV in order to control what each eye sees.

The lenses in the glasses use LCDs (liquid crystal displays) to block and then allow light though to the eye. The glasses alternate quickly between the left lens being opaque (right lens clear) and the right lens being opaque (left

Spotlight on… How to make 3D

There's only one way to make 3D movies from scratch: film the action with two cameras connected at a fixed distance from each other (generally a little further apart than human eyes). The film is then edited – usually digitally – so that each image stream is in sync. If you think about it, a movie shot this way has a smooth, continuous depth of field. Figures 1 and 2 were made this way: I shot two images, the second slightly offset from the first (they're both shown in Figure 2), then used an app on my iPhone to create the anaglyph in Figure 1.

Animations are especially easy: you render each frame as two images, each from two slightly different viewpoints. In fact, 3D video games work on this premise too.

The other way to make 3D movies is to shoot an ordinary 2D film and then convert it. The conversion is totally artificial, and is done using specialist software on fast computers. First, the editor assigns several depth of field planes to the image (typically three to eight). Each object in the frame is then assigned, cut out and placed in one of these planes. The result is a 3D space that consists of 2D images pasted on well-defined planes. The planes can be manipulated independently of each other (tilting them may improve depth of field, for example). Once the software has a complete definition of the 3D space (and note that the editor may have to fill in gaps from the cutting and pasting by cloning other areas), it can render the frame as two images, left and right. This technique produces a result without the smooth depth of field of true 3D and can tend to look awkward or wrong.

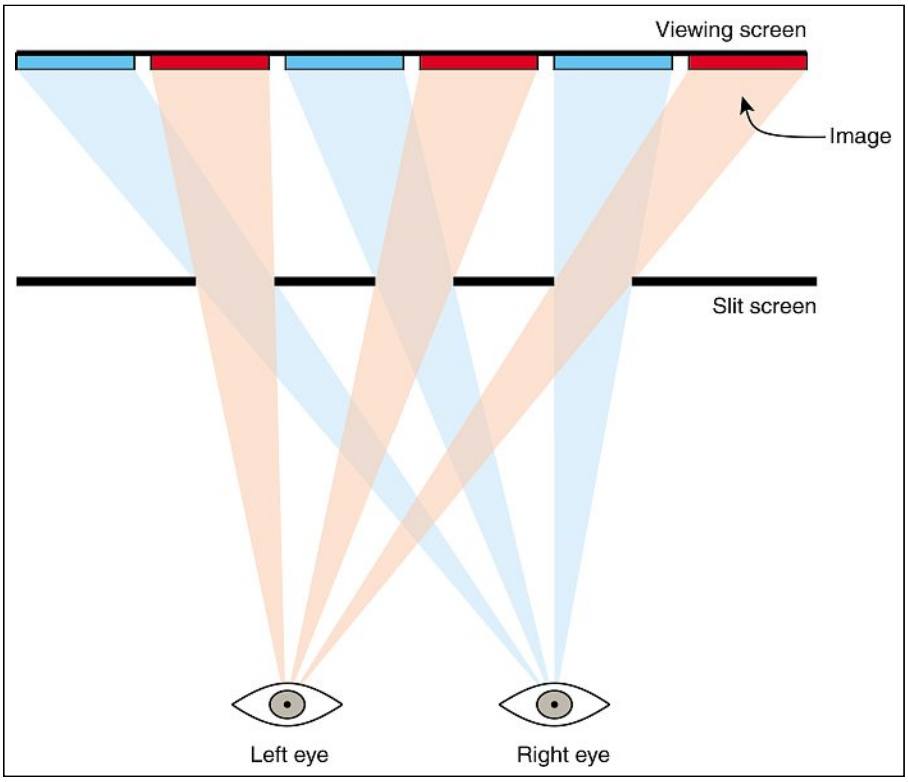

▲ **Figure 3: How the parallel barrier method works for viewing footage in 3D.**

lens clear), back and forth. The glasses communicate with the TV using infrared, so that this on/off cycle is synchronised with the TV image. The TV alternates between a left-eye frame and a right-eye frame very quickly at the same rate as the glasses are shuttering opaque/clear. This TV/glasses synchronisation means that your left eye only sees the left-side frames and your right eye the right-side frames. If you weren't wearing the glasses, you'd see a slightly blurry video.

Just like with movies on film, this image-switching and TV/glasses handshaking happens quickly enough that each eye sees continuous motion and the view just for that particular eye. For many years, plasma and LCD screens couldn't refresh fast enough for this trick. It's only recently, with the advent of 3D-ready TVs and monitors, that screens have become able to refresh quickly enough for the process to work effectively. In essence, the panel TV must refresh at 120Hz – double the refresh rate of a standard LCD panel.

This means that in order to watch a 3D movie on a TV, you must have a 3D-ready TV, an IR transmitter (although some now use radio or Bluetooth instead) and a pair of shutter glasses (as they're called). Although there's movement towards standardisation of the synchronisation chain, it's still advisable to buy all the parts from the same manufacturer.

Stereograms

Another technique for viewing 3D images (but not films) is the so-called cross-eyed stereogram. Here you're presented with two images side-by-side. The left-hand image is supposed to be viewed by the right eye, and the right-hand image by the left eye. In order to achieve this, you look at the images cross-eyed, say from a distance of about two feet away. You'll see four images initially. You then slowly uncross your eyes until a middle image is formed, which will suddenly seem to pop out into 3D. Try this with Figure 2.

An alternative is the parallel stereogram where again there are two images, but this time the left eye views the left image and the right eye the right one. Some people can see a 3D effect with these more easily than with the cross-eyed version. ■

Glasses-free HD

Despite the impressive progress made with passive and active glasses, there's still the nirvana of watching 3D without the need for any glasses at all. The best solution so far is lenticular displays. These displays have myriad tiny lenses (called lenticules) on a corrugated film on the screen. The screen uses these to display two images simultaneously in such a way that the lenticules aim the left frame in a different direction from the right frame. If your head is in the so-called sweet spot, your left eye will only be able to see the left frame and your right eye the right frame.

Unfortunately, there are numerous problems with this system that make it difficult to implement. First of all, the left and right frames must be interlaced into a single image in such a way that the lenticules direct the light in the correct course. The screen must have the lenticular film correctly and precisely aligned. The viewing angle is somewhat small – you must be perpendicular to the screen and the sweet spot is fairly small (at least when compared to a 3D-ready TV, for example). Nevertheless, lenticular displays are looking promising, although there's quite a long way to go before they're a practical alternative.

Making 3D images

There are many ways to create 3D images, like the ones I made for this article (see page 88). All of them involve taking two pictures offset from each other by roughly 3in. A cross-eyed stereogram is the easiest to make: after all it's just the two pictures side by side, with the right one on the left and the left one on the right.

For an anaglyph image, you can do it the hard way by using Adobe Photoshop (just search Google for "make anaglyph with photoshop" to find a few tutorials explaining how to proceed), or you can use a program that will do it automatically. For the images in this article, I used an app called 3D Camera on my iPhone. It prompts you to take the left picture and the right picture, then lets you position the two images over each other (I used the hedgehog's nose as a common target). After that, you can create the anaglyph in moments by clicking the appropriate icon.

Recording 3D movies is a more difficult and involved process. See page 88 for details. ■

A more popular autostereoscopic system is called parallax barrier. This is the system that's used on the new Nintendo 3DS gaming console. The effect is produced by two screens, one on top of the other (the 3D screen is produced by Sharp). The top LCD screen produces a set of very fine slits, each a pixel wide, through which the bottom screen can be viewed. Since the device is close enough to your face, your eyes will see different views of the display on the second bottom screen. This is all that's needed for the brain to convert the two different images into a single 3D rendition. Figure 3 shows a stylised example with three slits, where the left frame is coloured red and the right frame cyan.

The parallax barrier suffers from the same problem the lenticular screen does, in that there's only a small zone that produces the 3D effect, but for a small device like the Nintendo 3DS, that sweet spot is where you'd normally view the screen anyway – about two feet away. Again, just as with the lenticular screen, the left and right frames have to be precisely interlaced so that the 3D effect is produced. For games, this is relatively simple: instead of producing a single frame for the game's 'video', you have to produce two, each from a slightly different viewpoint (requiring twice the processing). 3D movies have to be converted from their default form to an interlaced form.

So there you have it. We've gone from the very earliest to the most up-to-date methods of showing 3D content. Given the state of flux, when are you going to put your money? Personally I think the Nintendo 3DS shows the way, but it'll be a while before large screens will be able to do the same. **PCP**

Julian M Bucknall has worked for companies ranging from TurboPower to Microsoft, and is now CTO for Developer Express.
feedback@pcplus.co.uk