swave

Display Technology Requirements for AR Smartglasses

Dr. Edward Buckley
Vice President, Augmented Reality Solutions

Swave Photonics | Whitepaper

May 2024

swave

hey@swave.io

Table of Contents

1. Introduction

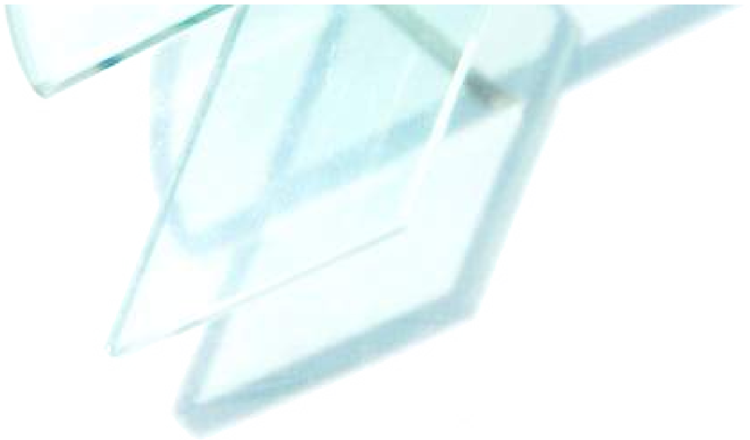

There is no doubt that Augmented Reality (AR) is the mobile computing platform of the future. However, the lack of a compelling mass-market AR product is due to the inability of existing display architectures to meet requirements for all-day wearable AR that pose extremely challenging - and, in many cases, competing - requirements.

For example, while it is generally accepted that truly immersive experiences require a wide field of view (FOV), these requirements are in direct opposition to the large eyebox and small form factor necessary for a usable, socially acceptable product. The same is true of brightness; while a high brightness display is needed to compete with ambient conditions, this requirement is in direct opposition to the need to enable a low enough power consumption for all-day use given a reasonable capacity battery. Last, and as is the case for all consumer electronic products, mass-market adoption places strict constraints upon the bill of materials – constraints which are impossible to meet when system complexity is high, when yield is low and when component technologies are so difficult to make that they represent a cost and supply risk.

The reason that mass-market AR does not yet exist is because, until now, none of these requirements could be simultaneously met by any display technology. Whilst conventional panel-based (i.e. LCOS or DLP) systems are mature and relatively low cost, they are too large to be integrated into fashionable eyewear while maintaining a reasonable resolution. MicroLED-based systems, on the other hand, are currently unable to meet the yield requirements consistent with a consumer product, and their reliance on extremely expensive, inefficient waveguides to provide a useful FOV/eyebox combination only serves to exacerbate such issues.

* * *

Swave's technology takes a **fundamentally different** approach by exchanging optical for computational complexity, avoiding the physics-based limitations that constrain current approaches and instead leveraging the proven paradigm of Moore's law.

As we will show in this white paper, this approach is the **only one** which has the potential to deliver on the promise of all-day wearable, mass-market AR.

2. Limitations of Current AR Systems

2.1 Eyebox & FOV

The issues faced by light engines for AR applications, and the proposed solutions, can be understood by considering the nature of etendue in an optical system. In an AR system, this is defined as the product of the panel area **A** and its emission solid angle **Ω** and, as in any other optical system, the best we can hope for is to conserve etendue. In a near-eye display application, the panel area defines the eyebox and the solid angle sets the FOV so that:

$$\mathbf{A}\ x\ \mathbf{\Omega} = \mathbf{eyebox}\ x\ \mathbf{FOV} = \mathbf{constant} \tag{1}$$

Form factor requirements dictate that the panel size should be no more than a few millimetres on one side, and since we typically would like an eyebox of around 10mm, the panel must be optically magnified. A direct consequence of ***Equation (1)*** is that any magnification of the eyebox is associated with a concomitant reduction in FOV. A popular way of overcoming the etendue limitation of ***Equation (1)***, while maintaining a reasonable form-factor, is to use a waveguide but, unfortunately, many of the fundamental issues associated with current AR display system designs trace directly back to their use.

2.2 Issues with Waveguide-based Systems

Although waveguides seem an attractive proposition for solving the etendue limitation of ***Equation (1)***, there are many practical disadvantages; we summarise these in ***Table 1*** below and provide details in the corresponding sections.

Issue	Reason
Cost, Yield, Manufacturability	• Diffractive waveguide structures are very difficult to manufacture repeatably; Section 2.2.1 • High index materials are expensive and difficult to use; Section 2.2.1
Efficiency and image quality	• Diffractive waveguides have a very low efficiency; Section 2.2.1 • Uniformity issues in diffractive waveguides place an additional heavy burden on the light source; Section 2.2.1 • Vergence-accomodation conflict cannot be addressed; Section 2.2.2
Mass-market applicability	• Prescription correction requires additional optics; Section 2.2.3

Table 1 - Issues with current waveguide technologies.

2.2.1 Efficiency and Image Quality

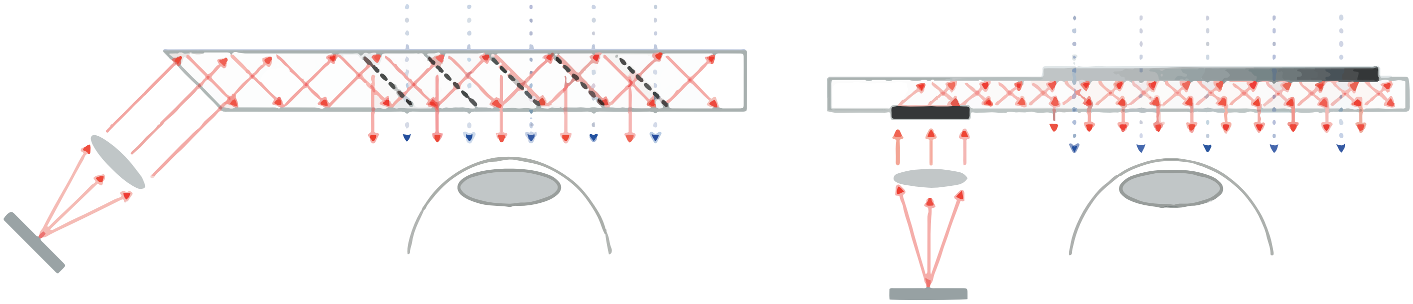

Figure 1 - Illustration of eyebox expansion in a waveguide by pupil replication. Incoming light rays undergo multiple bounces in the waveguide, replicating the input eyebox as they exit. After **[1]**.

Waveguides are often used to perform pupil replication in a similar manner to ***Figure 1*** so that a wide FOV and large eyebox are simultaneously achieved. Although a simple concept, it is exceptionally difficult to achieve in practice under the efficiency and image quality constraints imposed by AR product requirements.

It is inevitable that waveguides contribute to size, weight and cost of an AR product, regardless of the technology employed. However, designs suitable for wide FOV typically necessitate high index materials which are expensive and difficult to manufacture. Diffractive waveguides are generally preferred for lightweight designs, but are exceptionally difficult to realize in practice because imperfections in grating manufacture typically result in poor efficiency, uniformity issues across the FOV and eyebox, and degradation of display acuity. Manufacture is expensive and yield is low, and ultimately waveguide design goals are fundamentally at odds with the laws of physics.

Worse still, the poor uniformity of diffractive waveguides places a heavy burden on the light source. To make a uniform display, it is necessary to increase the brightness of the light source according to the lowest transmission in the FOV. For a LCOS light engine, this is disastrous; if the uniformity is ***N*:1** across the FOV, then the light source brightness needs to be increased by a factor of ***N***. The use of zonal illumination for LCOS, or even perhaps microLED, helps this situation somewhat but does not provide a complete situation because in general it is not possible to predict with any certainty where the lowest transmission will be.

So, any light emitting element must be capable of being driven ***N*** times harder, with drive circuits that can supply ***N*** times as much current. While this may be theoretically possible, it is certainly not possible to maintain peak light source efficiency over ***N*** times the range of display brightness. Waveguide uniformity issues therefore ultimately translate to reduced efficiency and increased cost.

2.2.2 Vergence-Accommodation Conflict

Waveguide-based systems are stereoscopic; that is, one display per eye is used to synthesise a stereo effect. However, such systems mandate a fixed-focus image plane, with the plane often placed at a depth of approximately two metres. This means that the display imagery can only ever be naturally augmented with the outside world at one specific plane, which conflicts with many of the proposed AR use cases which specifically require accurate placement of images in depth. A good example of a real-world AR scenario which would require accurate depth control is shown in ***Figure 2***; there are multiple planes which need to be used to provide indications to the user at various points in time.

If a display system were to present these notifications at a fixed depth plane, the experience will inevitably be unnatural and uncomfortable. The latter is a serious consideration and is the well-known problem of vergence-accommodation conflict (VAC), arising when a viewer expects to see a real-world image at some distance ***(Figure 3, left)*** but instead accommodates on a display that is showing the image ***(Figure 3, right)***.

Although VAC is known to cause nausea, and will certainly provide comfort issues that are inconsistent with the requirement for all-day wear, there are no known viable solutions to this problem for waveguide-based AR displays.

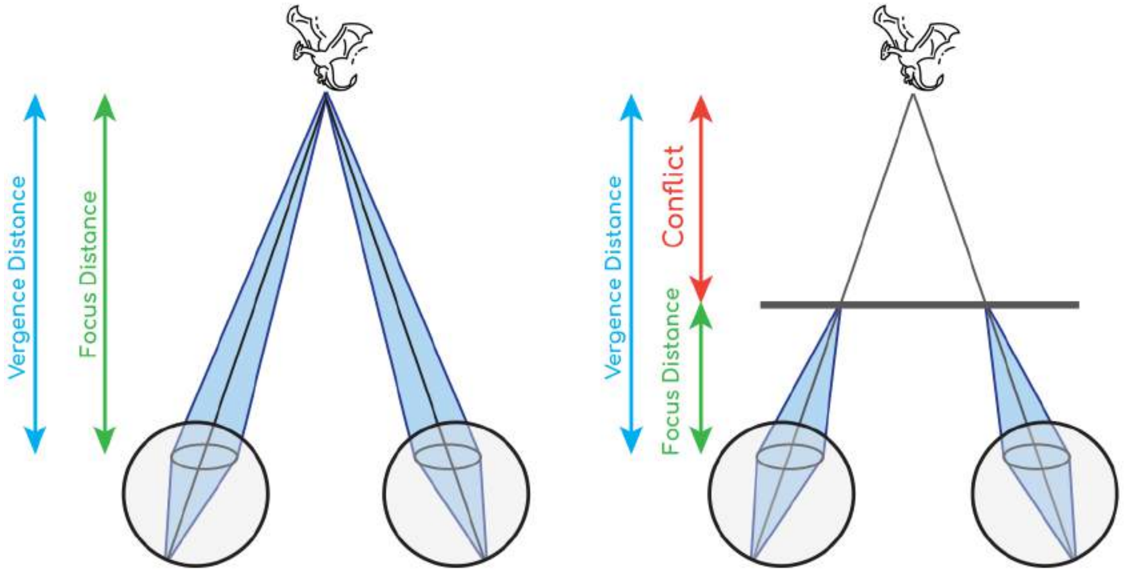

Real World

Stereoscopic 3D Display

Figure 3 -
An illustration of vergence-accommodation conflict (VAC) in an AR display. The user expects to see an image at some given distance (left), but instead accommodates on the display which is rendering it. The difference between the two is the conflict.

2.2.3 Prescription Correction

In general, AR display systems must be corrected for wearers that use prescription eyewear. In a waveguide-based system, this can only be addressed by adding corrective optics, matched to the user's prescription. Not only does this make size and weight goals very difficult to achieve, as well as significantly increasing cost, it complicates a brand's go-to-market strategy and associated logistics. Given that more than 60% of adult Americans require some kind of corrective prescription **[2]**, this is certainly a barrier to mass-market adoption.

Figure 2 - An example AR scenario which requires the display of images at various depth planes.

3. Overcoming Waveguide Limitations

Since the requirements for waveguides in AR display systems are pushing the laws of physics, the ultimate result is sub-optimal tradeoff against size, weight, cost, yield, throughput, social acceptability and mass-market adoption. For mass-market, all-day wearable AR to be successful, it is clear that removing the waveguide would be very desirable. There are only two known techniques for achieving a waveguide-free display with the required characteristics; either by directly displaying the desired image into the user's eye using some type of pupil-tracked display, or by constructing a holographic (i.e. diffraction-based) display with such a large cone angle that the eyebox/FOV tradeoff defined in ***equation (1)*** still produces an acceptable number for each.

3.1 Retinal Projection Displays

Direct retinal projection displays **[3]**, **[4]** are perhaps the most obvious way to remove a waveguide, and have the added attraction of a very low optical power requirement. However, such displays require extremely accurate pupil tracking, which must be performed not only at a reliably high accuracy and precision, but also over ~6 orders of magnitude illuminance, for a huge variation of pupil sizes and all while accounting for ocular reflections and eyelash occlusions.

Furthermore, if a retinal display is realized by a scanned-beam laser projector the above issues are compounded by well-known concerns regarding eye safety and image quality due to scan patterns that become visible under eye motion.

Hence, although direct retinal displays are theoretically attractive, they in fact present technical challenges that are, at the very least, as order-of-magnitude as difficult as fabricating a waveguide without addressing the issues of prescription correction and VAC.

3.2 Holographic Displays

A second approach is to exploit holography via a diffraction-based display so that optical complexity is exchanged for computation. This is not a new approach, and has been successfully demonstrated in a number of products **[5]**, **[6]**, but to use this technology to develop an AR display requires some fundamental breakthroughs. We begin by describing the background to such a system, the requirements imposed upon the microdisplay - which is conventionally referred to as a spatial light modulator (SLM) when used in holographic mode - and Swave's breakthrough enabling HXR technology.

4. Holographic Displays for AR

A holographic display is a type of computational display in which we choose the image I that we want to form, and compute a hologram pattern h which, when illuminated with coherent light, gives rise to an approximation to I by the process of diffraction. The huge advantage of employing this approach is that diffraction can be used to explicitly control the wavefront of light. This makes holography the only display technology that can provide “true” 3D images - but, for reasons which shall become clear later, our interest lies in forming 2D images at a depth which can be arbitrarily adjusted.

4.1 Basic Principles

A basic 2D display geometry is shown in ***Figure 4*** below; when the hologram h is illuminated by coherent collimated light of wavelength λ, the complex field formed in the back focal plane of the lens of focal length f is the 2D spatial Fourier transform $F(h)$ of the hologram pattern. Practically, h is displayed on a SLM, which has N pixels of size Δ and a finite number of levels, and humans perceive intensity $I = |F(h)|^2$ with the phase of the reconstructed image I being irrelevant. Computational aspects are outside of the scope of this paper, but we note that real-time computation of holograms in custom silicon is a solved problem and was first demonstrated some years ago **[7]**.

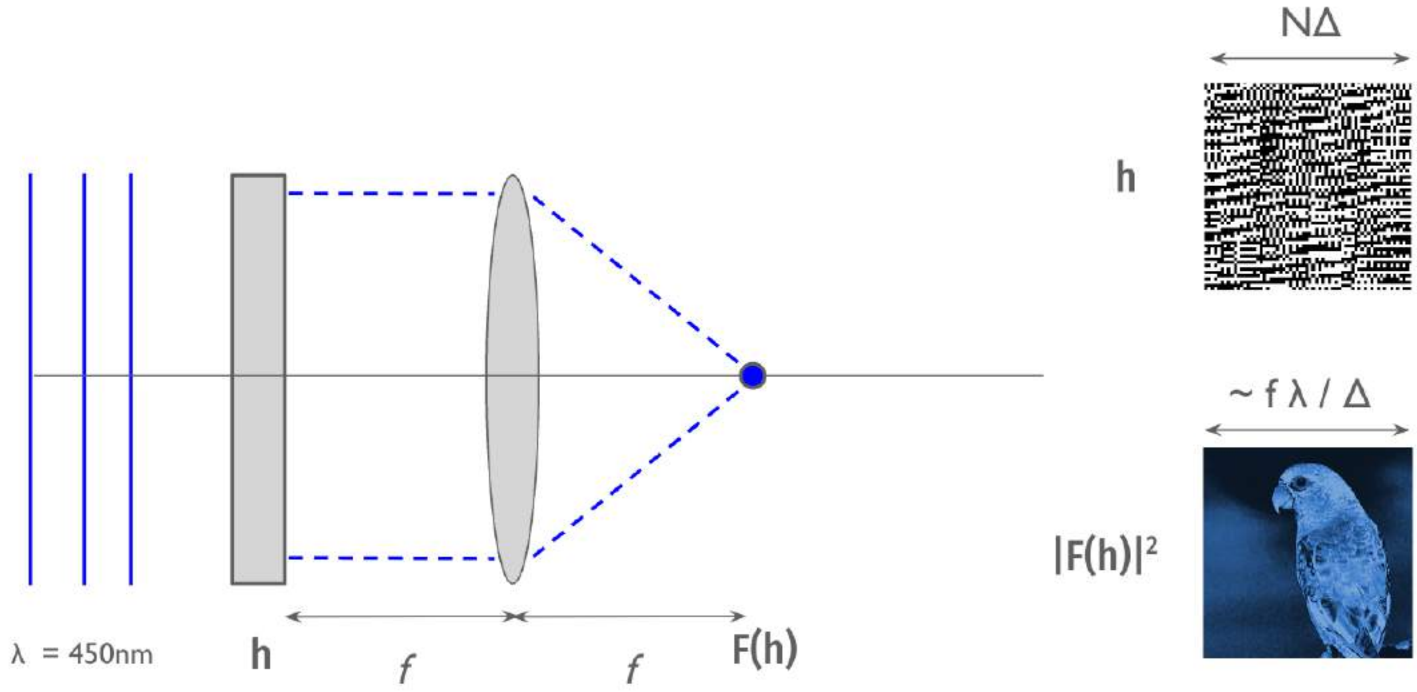

Figure 4 - An example of 2D image display using a hologram. The hologram h is illuminated by coherent light at λ = 450nm and the desired intensity image, I, is formed by an optical Fourier transform $|F(h)|^2$ at the back focal plane of the lens. In a display application, the hologram h is shown on a spatial light modulator with N pixels of size Δ.

4.2 SLM Requirements

Due to the very different operating regime of a holographic display, the key SLM requirements for AR applications are less immediately apparent than for conventional direct-view display panels. Here, we briefly outline the requirements for SLM pixel size, pixel count and speed.

4.2.1 Pixel Size

From the grating equation, we know that the diffraction angle from pixels of size **Δ** when illuminated by light of wavelength **λ** is:

$$\theta = 2 \sin^{-1}(\lambda / 2\Delta) \tag{2}$$

and hence, the smaller the pixel size, the larger the diffraction angle. Due to this nonlinear relationship, and as shown in ***Figure 5*** below, the addressable FOV grows very quickly as pixel size is reduced. For high performance all-day wearable AR, a FOV of approximately 45 degrees is required. If we assume that we will need to optically magnify the SLM by a factor of at least two to achieve the desired eyebox size, then from ***equation (2)*** we will need a diffraction angle of at least 90 degrees. It follows directly from ***Figure 5*** that a pixel size significantly less than **λ** is required, and we refer to these as "nanopixels".

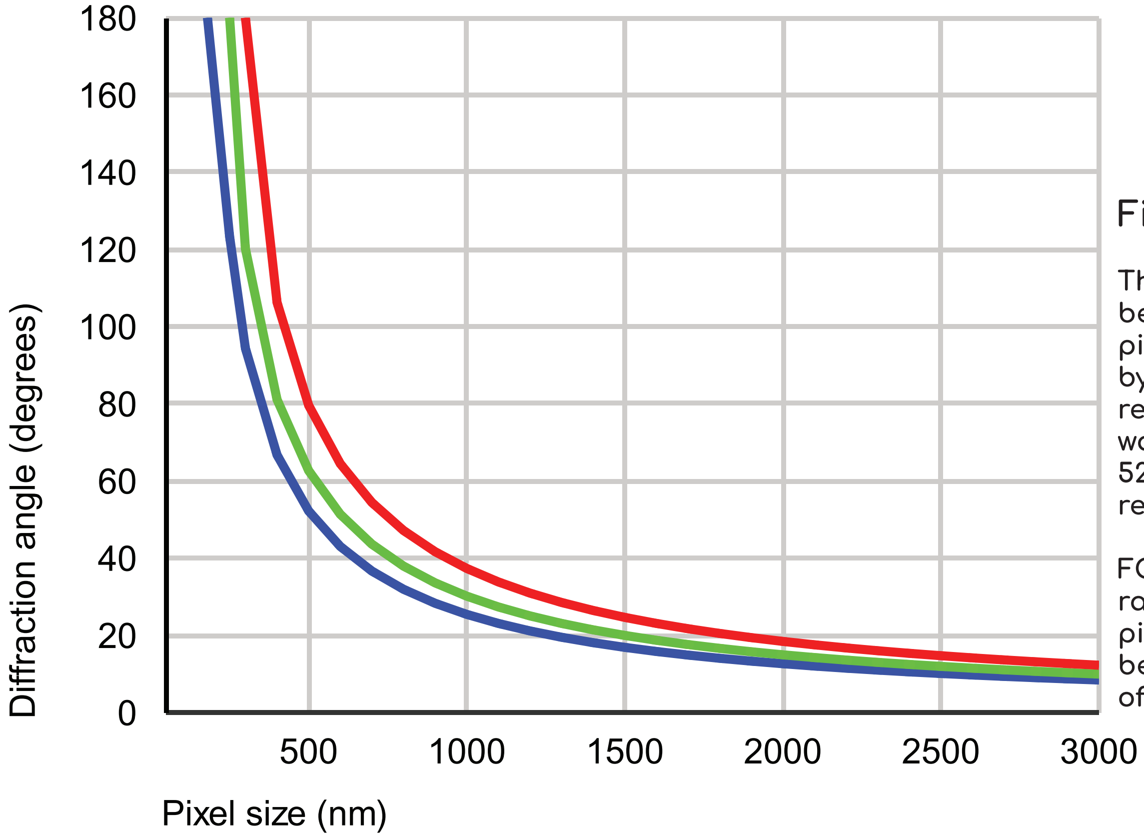

Figure 5 -

The relationship between FOV and pixel size described by *equation (2)* for red, green and blue wavelengths of 640nm, 520nm and 440nm respectively;

FOV increases rapidly as the pixel size decreases below the wavelength of light.

4.2.2 Number of Pixels

In a holographic display, the number of SLM pixels is not the same as the number of pixels in the image because the number of levels available from a SLM pixel is typically small. Consider an ***M* × *M*** input image at 8 bits per pixel and an SLM that is capable of showing ***N* × *N*** nanopixels at ***n*** bits per nanopixel. If the number of bits ***n*** is small, which is necessary to make the pixel logic size compatible with nanopixels **Δ < λ**, then it follows that we need ***N* >> *M*** to maintain the same amount of information.

On the other hand, if we choose not to maintain the same amount of information, we change the contrast ratio in the reconstructed image. The quantized nature of the modulator determines the noise floor, and each addressable image in the reconstructed image comes with some associated quantization noise. In a holographic display, then, contrast is proportional to the ratio of the number of image pixels to the available nanopixels on the modulator i.e. **~ (*N/M*)**^{2}. Since contrast and resolution are both important in AR, ideally the SLM should provide on the order of ~100 million nanopixels.

4.2.3 SLM Speed

Previous approaches to holographic image display have used a fast modulator, so that many independent holograms per video frame are displayed for the purposes of quantization noise averaging **[8]**. However, although this helps to reduce the variance of the noise, thereby improving uniformity, it cannot reduce the mean value of the noise. Hence, a fast modulator cannot improve contrast. So, where the SLM has so many nanopixels compared to the image that the contrast is likely to be large in the first place, high frame rate operation is unnecessary and the SLM can operate at video rate. As we will see in ***Section 4.3.1***, if the SLM is bistable, the average frame rate can be reduced further still.

4.3 Swave's HXR Microdisplay

To make holography feasible in an AR form factor requires a SLM technology that can deliver ~100 million nanopixels in ~5mm × 5mm, while operating at video rate. There is no off-the-shelf display that can achieve a subwavelength pixel size; even the most advanced LCOS displays are limited to ~3μm pixel pitch due to fringe-field switching effects.

To address this, Swave has developed a proprietary microdisplay that enables a sub-300 nm pixel by depositing non-volatile phase-change materials on a CMOS backplane using a standard 300 mm process. Phase change materials are well known and understood due to their use in many high volume products (such as rewritable DVDs, for example) and operate as a modulator in this application by providing a refractive index change as the material is thermally actuated between its amorphous and crystalline material phases.

The use of CMOS enables scaling to ~tens of Gigapixels, if required, and HXR is the only known microdisplay technology that can satisfy the key performance criteria outlined in ***Section 4.2***.

4.3.1 Bistability

Bistability results from the ability to change the optical state of a material by altering its material properties. A well known example of a bistable display is E Ink's, in which the optical state of a pixel is set by rotation of one or more microcapsules. Swave's pixel, on the other hand, is bistable because its optical properties are determined by the refractive index difference between amorphous and crystalline material phases.

Bistability generally confers a huge power consumption advantage because it eliminates the necessity to continuously refresh the display. Most displays need to be refreshed at some rate to display images - even if the image is static - because a charge relating to the pixel value is held on a capacitor, which will leak away unless refreshed. Typically, a refresh rate of 90 Hz is required in high ambient light conditions to prevent flicker; for field-sequential displays, this needs to be higher still to avoid color breakup artifacts. In a bistable display, however, there is no requirement to refresh the image at all unless it has changed; once the pixel value is set, the state is optically maintained until it is changed again. Hence, for still images, the only power required is that necessary to illuminate the nanopixels.

This fact is particularly relevant to AR, in which content changes much more slowly than the typical refresh rate required by a conventional display; during a day of use, a typical AR display will show messages and symbols for the majority of the time, with only relatively short bursts of video. Under these circumstances, a bistable SLM allows the average frame rate of the AR display system to be dramatically reduced. Swave's technology is the only AR display technology which can exploit this fact, significantly reducing power without introducing unacceptable flicker artifacts.

4.3.2 DyanamicDepth™

Holography is a wavefront control technology. Although we can use this technique to deliver “true” 3D, in the sense that all 3D motion cues such as motion parallax are present, this is unnecessary for AR. For an AR display that can deliver content which is seamlessly and naturally augmented with the outside world, it is sufficient to generate a 2D plane in which the depth of the plane is controllable. With a holographic display, this is trivial to achieve computationally i.e. without any additional optics, and in fact using the same SLM, and the same optical system, we can position content anywhere from ~0.5m to optical infinity. We call this feature DynamicDepth™ and, not only does it enable compelling AR scenarios such as visual search, it also eliminates VAC.

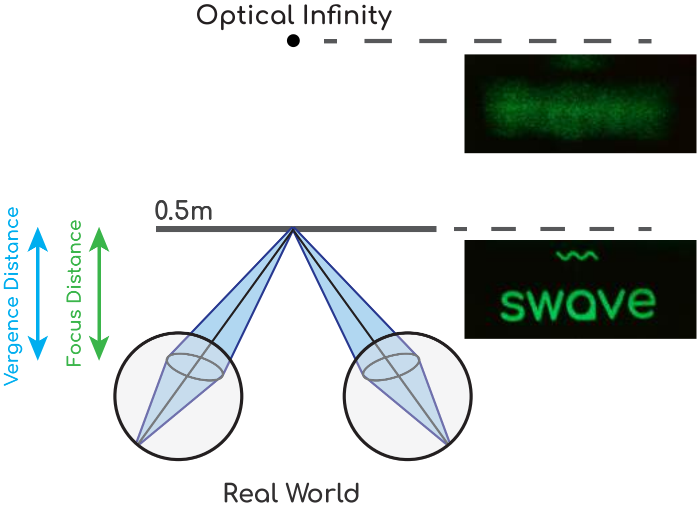

Figure 7 - Rendering of an image at a distance of 0.5m by Swave’s holographic display, correctly matching the focus distance of the observer. The corresponding image at optical infinity is now out of focus, which is exactly the behavior that the human visual system expects. The focus change is achieved entirely computationally.

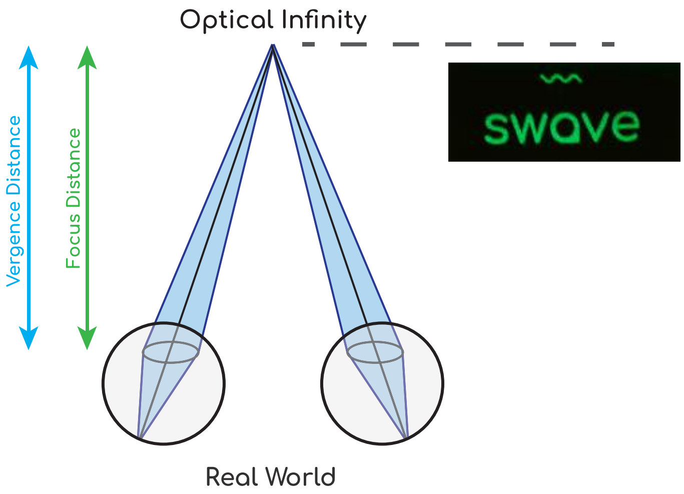

Figure 6 - Rendering of an image at optical infinity by Swave’s holographic display, correctly matching the focus distance of the observer.

An experimental proof of DynamicDepth™ is shown in ***Figure 6 (above)***, in which Swave’s holographic display is configured to provide an image at optical infinity; if the viewer’s eye is accommodated to any other distance other than infinity, VAC will result. This is a general consequence of any fixed focal-plane AR display.Since holography allows us control of the optical wavefront, it is straightforward to computationally modify the focal plane of the image, which is illustrated in ***Figure 7 (left)***. The image is rendered at the same plane to which the user is accommodated, and the image at optical infinity is out of focus. This is exactly the behaviour expected by the human visual system, and the behaviour which addresses VAC. What we have demonstrated above is a programmable defocus aberration, but practically there is no restriction on the complexity of the aberration that can be encoded. It is this property that also allows the display system to be corrected for the precise prescription of the user.

5. Summary

By using the requirements derived in previous sections, it is straightforward to put together an exemplary set of specifications for AR glasses that could be realized using Swave's HXR technology and known, low-cost optical components. In the example shown in ***Table 2*** below, we have chosen to employ a 16k x 16k SLM to realize a wide FOV AR display, but we note that an 8k x 8k SLM may be more appropriate for smaller FOV designs.

Uniquely, Swave's HXR microdisplay has such an ultra-small pixel pitch, and consequently such a large diffraction angle, that the eyebox/FOV tradeoff of ***equation (1)*** still produces an acceptable number for each. It is therefore unnecessary to use a pupil-replicating waveguide, and a much simpler and more efficient holographic combiner can be used instead.

Specification	Value	Notes
SLM Pixel Size	280 nm	Set to approximately half wavelength
SLM Pixel Count	16k x 16k	
SLM Active Area	~5mm x ~5mm	
Eyebox Size	~10mm x ~10mm	Optically magnified by ~2x per *equation (1)*
FOV	45° Diagonal	Optically demagnified by ~2x per *equation (1)*
Addressable Image Size	2400 x 1400	Chosen to maximise contrast
Maximum Acuity	~60 Pixels per Degree	
Image Depth	0.5m to ∞	Enabled by wavefront control in hologram computation

Table 2 - Exemplary specifications for an AR system employing a holographic display, which can be realized using Swave's HXR microdisplay and known, low-cost, optical components. Since pupil replication is unnecessary, a low-cost holographic combiner can be used instead of a waveguide.

6. Conclusions

Swave has developed a radically new SLM technology that provides a huge number of ultra-small pixels, which are the fundamentals required for next-generation holographic displays. Such a display will uniquely enable the mass adoption of all-day wearable consumer AR because it is the only display technology capable of simultaneously delivering on size, weight, cost and functionality requirements.

The only display technology capable of simultaneously delivering on size, weight, cost and functionality requirements.

By exchanging optical complexity for computation - riding the trend of Moore’s law rather than fighting against the laws of physics - Swave is able to deliver a wide FOV, large eyebox display that does not require complex and expensive waveguides, can account for prescription eyewear without additional optics, can eliminate VAC, provide content that is truly augmented to the virtual world, and enable untethered all-day operation. This is the future of AR that we have been promised.

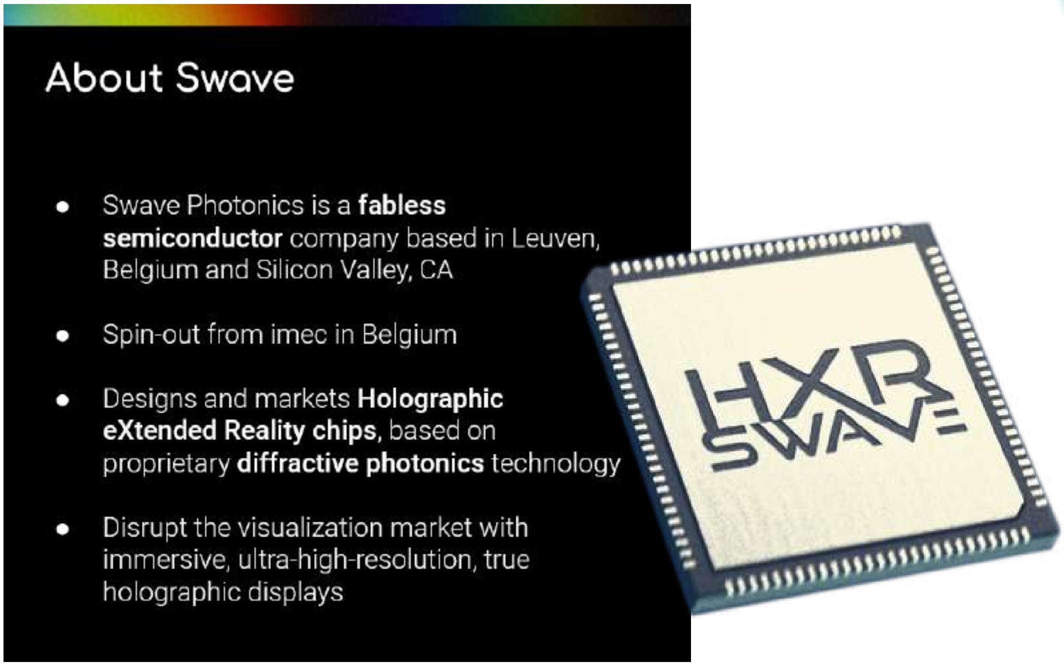

About Swave

- Swave Photonics is a **fabless semiconductor** company based in Leuven, Belgium and Silicon Valley, CA
- Spin-out from imec in Belgium
- Designs and markets **Holographic eXtended Reality chips**, based on proprietary **diffractive photonics** technology
- Disrupt the visualization market with immersive, ultra-high-resolution, true holographic displays

7. References

[1] S. J. Park, C. Keum, H. Zhou, T.-W. Lee and W. Choe, “Progress and Prospects of Nanoscale Emitter Technology for AR/VR Displays,” Advanced Materials Technologies , vol. 8, no. 20, p. 2201070, 2022.

[2] “The Vision Council,” 2021. [Online].

[3] Y. Wu, C. P. Chen, L. Mi, W. Zhang, J. Zhao, Y. Lu, W. Guo, B. Yu, Y. Li and N. Maitlo, “Design of retinal projection displays enabling vision correction,” Optics Express, vol. 25, no. 23, pp. 28223-28235, 2017.

[4] L. Mi, C. P. Chen, Y. Lu, W. Zhang, J. Chen and N. Maitlo, “Design of lensless retinal scanning display with diffractive optical element,” Optics Express, vol. 27, no. 15, pp. 20493-20507, 2019.

[5] “Light Blue Optics launch Light Touch™ interactive projector,” 5 January 2010. [Online]. Available: https://www.businesswire.com/portal/site/LightBlueOptics/.

[6] “LAND ROVER INTRODUCES WORLD FIRST LASER HEAD-UP DISPLAY FOR RANGE ROVER EVOQUE,” 19 September 2014. [Online]. Available: https://media.landrover.com/en-gb/news/2014/09/land-rover-introduces-world-first-laser-head-display-range-rover-evoque.

[7] [Online]. Available: https://www.youtube.com/watch?v=hrMfuFXFdf8.

[8] E. Buckley, “Holographic Laser Projection,” Journal of Display Technology, vol. 7, no. 3, pp. 135-140, 2011.

swave

hey@swave.io

swave