

Edge Blending

White paper

Edge blending is a technique used to describe the process of visually combining several projected images to make a single seamless image.

Typically, this is done to increase the size of a projected image to make a very wide image, or by combining a number of lower resolution devices together to increase the total resolution of a display. In both cases, a number of projectors are needed which are firstly overlapped and then visually joined together using an edge blending technique.



Figure 1: A wide image that we want to project using three projectors

At first glance it might appear that the easiest way to produce this image across three projectors would be to line up the projectors next to each other. This method is technically referred to as image butting because the projected images are butted up next to one another. To try and accomplish this, the projectors are carefully arranged so that the right-hand edge of the left-hand projector sits perfectly next to the left-hand edge of the right-hand projector.

Unfortunately, achieving a seamless image using this technique is near impossible. In principle it sounds easy, but in practice it is very hard to do. There are a couple of reasons for this. Firstly, different projectors have different characteristics, even if you are using the same brand and model. For example, an image may drift as they warm up and the brightness and colors will also differ slightly even if the same projector brand and model are used. Secondly, it is a difficult task to manually adjust a projector to exactly the right position so that they match up at the edges. Some areas will overlap whilst others will have a gap. This results in a tiled appearance, rather than a seamless one as shown in Figure 2.




Figure 2: Edge Butting

The answer lies with the use of image blending. When image blending, we overlap projectors to create what is referred to as an overlap region. For a detailed description about overlapping regions, please read our explanation here. In this overlapping region the same proportion of the image is projected from both projectors to provide a duplicate image region. This means we can no longer simply cut the projected image into exact parts to match the exact resolution of the total number of pixels used in all projectors, but we now have to compensate for overlapping regions.




Figure 3: Wide image with overlapping regions

Figure 3 shows an image with overlapping regions which are overlapped in the final projection. A larger overlapping region provides better control of the overall image, but this has to be balanced against the reduction in image size and the loss of overall number of pixel used in the display.

Now that we have an overlap region, it is a little easier to match up the projectors without having obvious discontinuities in the image. But now we have a new problem: the overlap region is twice as bright as the rest of the image, as this area now has the brightness of two projectors as shown in Figure 4.



Figure 4: Double brightness in the overlap region

To overcome the double brightness in the overlapping regions, we can add a brightness (luminance) control function in each of these areas. On first appearances, we could halve the luminance of the overlap region, so that when the two projectors combine, it comes back to the full brightness on a single projector. Unfortunately that makes alignment almost as hard to line up as image butting. The solution is to blend the images together. This is where the blending part of the image blending name comes in: we gradually drop off the luminance from one side as we increase the luminance from the other as shown in Figure 5.




Figure 5: Blending in the overlapping regions gradually produces a visually correct image

As one projector gradually takes over from the other, this overcomes the problems we identified above: Firstly, because one projector slowly gives way to the other, changes in color and brightness are less apparent. Secondly, it is also quite error tolerant. If the projectors don't perfectly line up, the gradual blending makes it hard to notice.

The amount of edge blending performed in each region is controlled by a gamma function. A gamma ramp (as often used in computer graphics to deal with luminance), can be used where the brightness drops off in a non-linear curve.

A gamma ramp has a single parameter, called the gamma value. If the gamma value is 1, then the gamma ramp is linear. If the gamma value is less than 1, then the blend region quickly becomes bright and slowly approaches full brightness. If the gamma value is greater than 1, then the ramp stays darker and rapidly approaches full bright.




Figure 6: Examples of gamma ramps: red is $\gamma = 0.3$, blue is $\gamma = 1$, and green is $\gamma = 3$.

The mathematical relationship is $o = i^\gamma$, where o and i are the output and input values in the range $[0..1]$ and γ is the gamma value. Typically, a larger overlap region, blended with a lower gamma value will produce the best blends. This is because it will give a slow change over from one projector to the other.



Figure 7: Image once blended