7. International Immersive Projection Technologies Workshop
9. Eurographics Workshop on Virtual Environments (2003)
J. Deisinger, A. Kunz (Editors)

Color Consistency for Digital Multi-Projector Stereo Display Systems : The HEyeWall and The Digital CAVE

W. Kresse, D. Reiners and C. Knöpfle

Fraunhofer-IGD, Fraunhoferstrasse 5, 64283 Darmstadt / Germany

Abstract

Digital projectors have a significant advantage over CRTs for IPT setups: brightness. But they also have a number of disadvantages, one of which is color consistency. This problem is exacerbated when using the Infitec method for stereo separation, which in itself has some strong advantages for CAVE and tiled wall setups. In this paper we will describe a method for color and brightness correction of multi-projector display systems. The method itself is used in two new projection systems, which are currently under construction at Fraunhofer-IGD: The HEyewall and the Digital CAVE. The HEyeWall is the first stereo capable tiled display worldwide. The Digital CAVE is the first CAVE with digital projectors and stereo separation based on Infitec(tm). In this paper we present these new IPTs in more detail and also present our experience with digital projectors.
To calibrate all the involved projectors photometric measurements of the different projectors are used to calculate a common gamut in a linear colorspace. Input colors are mapped into this gamut and from there mapped into the individual projector's colorspace. This method allows to adjust the rendering output of two or more projectors with different color gamuts in such a way that the projected images are photometrically calibrated. Since the correction has to be done for each pixel, a straightforward implementation would be very slow and far away from realtime. Consequently we will outline a method how to improve performance and overcome this limitation.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Display Algorithms

1. Introduction

Multi-Projector IPTs have been around for more than 10 years now. CAVEs, Powerwalls and other systems have become commonplace in large companies for design review, visualization and group discussion of different problems. Some of these have employed stereoscopic methods to allow displaying 3-dimensional models. Different technologies allow the separation of images for the left and right eye to create the illusion of a world behind the wall.

By far most of these systems have used analog CRT projectors. In the beginning there was just no other technology available, only over the last couple of years digital projectors have become available and competitive.

And these analog systems have some significant advantages. They are extremely flexible. Geometric distortion and convergence can be adjusted for a large number of zones of the display independently, colors and brightness can be adjusted in a wide range, sometimes also variably for different parts of the screen. The image forms only very close to the screen, thus irregularities in the light path such as not quite perfectly flat mirrors don't really matter. As long as the adjustment range of the projectors is not exceeded it is almost always possible to get a high-quality image and to get a very good match between different projectors for stereo displays.

Compared to the nimble elegance of the CRT dancers digital projectors are chunks of rock. As all the pixels are positioned in a fixed regular grid relative to each other on a DMD chip or LCD panel all geometric adjustments influence them all in the same way. Many projectors don't have any kind of lens shift to move the image on the screen, if they do they only have a vertical lens shift, only the more expensive ones have a horizontal lens shift too. None of them have a convergence adjustment; some even have different, non-adjustable focal planes for each color channel. They have a clearly visible pixel structure, all pixels have a black border, on DLP systems they also have a black dot in the middle due to the mechanical structure of the DMD device. This can be a significant problem in situations were the user is close to the screen, e.g. a CAVE. In terms of brightness some of the newer ones allow different adjustments for different parts of the screen. One of the major problems of the digital projectors is the black level, they can not display a sufficiently dark black. Typical digital projectors have contrast ratios (full white to full black) of 500:1 or less. Typical CRTs have contrast ratios that are ten times better or more. This is an important problem, especially for styling applications.

Another problem is the color consistency. There are significant differences in the color spaces of projectors, even the ones produced by the same manufacturer, even produced on the same day. Furthermore, color and brightness can change depending on the working temperature and age of the lamp. For the typical application, which uses a single projector, this is not a big problem, for multi-screen immersive projection systems this can make the difference between immersion and looking at a patchwork of pictures, especially for tiled displays. There is, to our knowledge, only a single extremely high-end system that allows adjustment of its primary colors, all others suffer from this problem. This problem becomes imperative to solve when using Infitec for stereo separation (see section 3).

So why would one want to use digital projectors at all? The simple answer: brightness. CRTs have a brightness in the 250-300 ANSI Lumen range, which in combination with the losses incurred by the stereo separation system and the screen size can be equivalent to a moonlit night. Digital projectors can go up to more than 10000 ANSI Lumen, typical ranges for higher-end devices are 4000 to 7000. These brightnesses allow photometrically correct display of non-night situations, allowing the evaluation of styling or simulation at a new quality level, and generally make for a more comfortable working environment.

Thus the question opens up how to overcome the shortcomings of the new technology. Most aspects like optical and geometric adjustability, pixel structure and black level are intrinsic properties of the projector and cannot be influenced. There is no alternative to testing and measuring the different available systems and selecting the best compromise available. The color problem, however, can be approached in software.

We start by describing previous works in the IPT and color correction areas in section 2. Different stereo separation technologies and their applicability to digital projectors are described in section 3. We are building two instances of multi-screen digital projector stereo displays, a 5-sided CAVE system and a tiled stereo wall. The specifics of these are described in section 4. Our color correction algorithm is described in section 5. It is a purely software-based and given a remote-controllable Photometer can be fully automated. Section 6 deals with the utilization of current graphics hardware to speed up the color correction of displayed images. Results are presented in section 7, followed by open problems and suggestion for future work in section 8.

2. Previous Work

The CAVE [3] was one of the initial multi-screen IPT systems and a large number of similar installations have been created worldwide. A CAVE still gives one of the best feelings of immersion and thus is widely used in situations where spaces utilizing all 3 dimensions are displayed. CAVEs suffer from numerous problems, one of which is the color uniformity across screens. These variations can come from intrinsic projector problems as described above, as well as differences in screen materials used for the top, sides and floor.

Tiling a planar display surface to produce a higher resolution image has first been used in the University of Minnesota's Powerwall [5] and University of Illinois, Chicago's ImmersaDesk [4] systems. The rising power of common PC graphics cards makes them an increasingly attractive setup to harness the power of a cluster of PCs. Numerous monoscopic systems using a variety of projectors and sizes have been designed and built in the last couple years[6, 8, 1]. To our knowledge none of the current tiled walls are stereo-capable.

Brightness uniformity for displays has been investigated by different groups. In [10], a method to automatically calibrate an immersive display towards uniform luminace and color across the screen has been presented. The problem of disparate color gamuts, however, remained unsolved. The paper also investigated and compared several new and developing immersive projection technologies, especially concerning their photometric and colorimetric properties. Another method for generation of an alpha-map to achieve photometric uniformity is discussed in [13].

An interesting method of matching several projectors to a reference projector concerning luminance and color reproduction is presented in [12]. The approach computes look-up tables of a sub-sampled input domain. This effectively only works for projectors that differ in the brightness of their primaries, not in their actual color, as the look-up tables can not transport information between primaries (e.g. they can not add a fraction of the red intensity of a color to the green channel to change the color). There are some further shortcomings in this approach. The luminance response is linearized, resulting in an unnatural gamma value of 1.0, compared to about 2.4 of typical displays. Furthermore, the black level of the involved projectors is not accounted for in the chromaticity matching step, which yields varying chromaticity levels depending on pixel input values. By applying the black of the worst projector to all others, the contrast ratio is reduced.

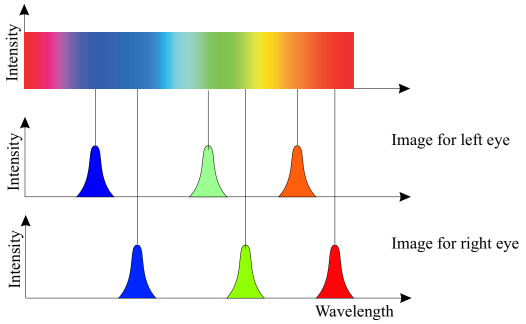

Figure 1: *Basic principle of Infitec*

3. Stereo Separation

Stereoscopic viewing is a key aspect in modern VR applications. The idea is to calculate different images for the left and right eyes and present them to the user so that he sees the correct image for every given eye. Doing that needs special hardware to block the other image from the corresponding eye, which can be done in an active and several different passive ways. The alternatives are:

- Active stereo (shutter)
- Image separation based on polarized light (linear, circular)
- Image separation based on single color bands (anaglyph)
- Image separation based on multiple color bands (Infitec)

Shutter, polarized light and anaglyph are fairly well known technologies within the VR industry, thus will not be explained in more details.

Infitec[9] is similar to polarization concerning the projector setup. Two projectors are required, both equipped with filters. Furthermore the user has to wear special glasses with matching filters. The difference: the *in*terference *fi*lter *tec*hnology uses spectral interference filters to select three narrow wavelength bands out of the visible spectrum. The three bands are different for each eye (see fig. 1).

The main advantage of this technology is that the image separation is significantly better than everything else. Furthermore there are no special requirements for the mirrors and screens, even very diffuse screens can be used. Because of the limited filter size, Infitec is currently only available for digital projectors.

The main disadvantage of Infitec is the visible color difference between left and right eye (see fig. 6). For one filter the illumination is very reddish, for the other one greenish. One reason is that standard data projectors have a color spectrum with varying intensities. Another reason are the sensitivity characteristics of the color receptors of the eyes. For example the red color band for the left eye stimulates the red and the green receptors, while the other red color band stimulates the red receptors only.

All serious stereo technologies for projection systems have in common that light has to pass through special filters and glasses. Thus the brightness is diminished in some way. We compared shutter, polarized light and Infitec with each other. Infitec consumes more brightness than the other methods, but with current projectors that is not as much of a problem as it used to be. In fact, many setups can't run the projectors at full brightness as they would glare the user.

To summarize, one can say that anaglyph is not an option for professional applications. Shutter has specific demands concerning the graphics hardware and projector technology, restricting its use to expensive - or dark - systems. Polarization and Infitec are similar and differ in three main issues: Crosstalk, color shift and screen requirements. Infitec is not a clear winner here, but it has other charactericties which made it the primary choice for our application cases.

4. Application Cases

We are building two digital multi-projector stereo display systems, the HEyeWall and the Digital CAVE.

4.1. HEyeWall

The goal of the HEyeWall project is to setup a high-quality, high-resolution, stereo display built using common off-the-shelf components. While a number of tiled walls have been built using digital projectors in the recent past, none of them are stereo-capable.

The central problem of tiled walls is to hide the fact that they are. The critical areas are the borders between tiles, which have to be hidden as good as possible. Overlapping the projection areas of the different projectors, which is often used to simplify the adjustment[16, 15, 18, 2], would have resulted in disturbing bright seams in dark areas due to accumulated black levels, which was not an option for us. Instead we decided to use a high-precision mechanical setup that would allow us to adjust the projectors finely enough to have them abut without overlaps.

Another problem is that the projector image does not stop at the edge of the area that is filled by pixels. Around that active area is a border which, depending on the projector, can be brighter than a black image. These borders have to be suppressed by mechanical blinds, in our case close to the screen. These blinds are also used to hide a small (2 pixel) region at the border of the screen, which is used as a sacrifice to prevent having to readjust the projectors for every single pixel of drift. The software can easily be adjusted to shift the image to correct this kind of problem.

To achieve a uniform brightness on the screen we use the correction texture of [11]. One major problem area that is left is the screen material. When viewed orthogonally the border between tiles will look good. But when viewed at an angle, for most screen materials the border will be very visible. The goal was to find a screen material that takes all light coming in from the back and distributes it as uniformly as possible to the front. After looking at a large number of screens we found the Gerriets Opera [7] to be the best. It is nearly perfectly diffuse and hides the seams between tiles very well. The disadvantage is that it destroys every polarization very effectively. Thus we had to choose Infitec as the stereo separation method.

The HEyeWall consists of a 6 x 4 array of tiles filling a 5 x 2.5 m area, giving an effective resolution of 6k x 3k pixel, with a single pixel area of $0.64mm^2$. Each tile is driven by 2 projectors, one for each eye, and each projector has an associated dedicated PC. We are using Christie Vivid LX-41 XGA LCD projectors, as they provide the necessary brightness to get a bright image using Infitec at a reasonable price.

4.2. Digital CAVE

In 2002 we started planning our new CAVE. Based on our previous experience we put together a number of requirements for a next generation CAVE. The most important one was more brightness compared to the old CAVE, where CRTs were used. So we decided to use digital projectors. Since we wanted to use standard PCs for the rendering, shuttering for stereo viewing was not an option. So we had to decide between polarization filters and Infitec. At the end we took Infitec because we had to reuse the floor of our old CAVE. Replacement was and is nearly impossible. Unfortunately it does not preserve the polarization.

After we had determined the projector and stereo technology, we had to think about the screens. Beside the various factors which influence visual quality and photometric consistency, interreflection between projection walls is an important issues for CAVEs[17]. Thus we were looking for screens with very low reflectivity, ideally black ones. As of today only a single company manufactures a real black screen, with low reflectivity and very good contrast. Since this screen has a well visible texture, its visual quality was too low for our applications. After testing several screens, we concluded that none of them is ideally suited for the CAVE and our requirements for photometric consistency (low reflectivity, low hotspot behavior, no color shift). We choose the Stewart Filmscreen 150 as the best of the worst.

For the past months we struggled with all the problems mentioned in the introduction. Often the only possibility to solve one of these problems was by replacing the whole LCD panel or even the whole projector. The very limited possibilities for manual adjustment of modern LCD projectors are the main reason. But thanks to the projector people we were able to solve most of them.

We plan the digital CAVE to go operational in March 2003. Its extends are 2.4 m cubed, with 5 sides used. For the rendering we use 10 standard PCs with nVidia Geforce4 graphics cards and a GBit Ethernet network. As software we use the Open-Source scene graph OpenSG and the VR-System Avalon, a joint development of IGD and ZGDV e.V. Darmstadt. The projectors are EIKI LC-UXT1.

5. Color Correction Algorithm

It is fairly easy to adjust two or more displays towards a common white point by mixing the three primary colors. But most displays already show severe color differences when simply displaying pure red, green or blue. Varying color primaries denote non-identical color gamuts, which means in essence that each display can utilize several colors that cannot be reproduced on the others, no matter how red, green, and blue are mixed. A drastic example of differing color gamuts is the Infitec stereo technology. As a result of the spectral selection, displaying any RGB value for both eyes will result in very distinctive colors for each eye (see the left column in figure 6 for examples).

To guarantee that any input-RGB color looks exactly the same on *all* displays, the common color gamut of all displays has to be determined, the gamma (or more general: non-linear) behavior before and after the correction has to be taken into account for each color component, and all input-RGB values have to be adjusted accordingly.

The drawback of any color and luminance matching approach is that the common color gamut will always be smaller than the individual gamuts, resulting in less saturated color reproduction, and the darkest display defines the maximum representable luminance for *all* displays. The major advantage is full color uniformity between all displays which have been calibrated in such a way (e.g., the right and left eye of the Infitec stereo technology).

All color computations are performed in a device independent color space, such as CIEXYZ.

5.1. Required measurements

For each projector and each color component, luminance (cd/m^2) and chromaticity (CIE-xy) have to be measured at several input levels. Since our measuring device was not connected to a PC, measurements were taken manually in 10% steps.

Luminance *and* chromaticity also need to be measured for black. Since this black is added by the projector to all displayed colors, *it needs to be subtracted from all other measured data before proceeding*. Failing to do so will lead to incorrect color information, especially for low intensity colors (see figures 4, 5). Furthermore, input values that are calibrated will change their properties after display.

5.2. Finding the common color gamut

The first step is to find the range of colors that can be reproduced by all displays. Their individual luminance capabilities are not relevant for this step.

After arbitrarily choosing one of the projectors as a starting point, we map the RGB primaries of the next projector into its color space. If any negative values occur, that particular primary can not be displayed by the first projector:

If only one color component of a primary is negative, it has to be clipped in CIE-xy space towards the corresponding primary of the second display. There exists a special case where the connection of the affected and the corresponding primary does not intersect the original gamut. In this case the primary of the original gamut lies within the second gamut, and is therefore the common primary.

For the case of two negative components, the process has to be repeated with reversed roles by mapping the RGB primaries of the original gamut into the color space of the new display. Now the maximum possible number of negative values is one.

The process can then be repeated with the next projector.

Another, probably more straightforward approach would be to directly perform 2D triangle clipping in CIE-xy space for all RGB primary sets, and then use the largest fitting triangle of the intersection as common gamut.

The resulting gamut (cf. figure 3) is guaranteed to be displayable by all involved projectors.

5.3. Finding the maximum intensities

After establishing the common gamut, for each projector the primary intensities have to be computed which reproduce the new, common primaries.

We start with arbitrarily set of luminances for each primary. For each projector, the primaries are mapped into display color space. The sum of all three red components (green and blue respectively) must not exceed 1, therefore we need to solve for the three scaling factors.

The resulting intensities are the maximum luminances with which this particular display can reproduce them. By setting the luminances to the lowest value of each projector for each color component, we determine the final primary luminances.

However, the newly common white, defined by adding all three primaries, may not exhibit suitable chromaticities. But in fact the white point of the new gamut can be chosen arbitrarily: by defining a new white point, such as the native white of the involved projectors, or a standard white such as D6500, and scaling the new primaries accordingly, we can calibrate the new gamut - and thus all displays - towards that white point. Care has to be taken that no resulting color values are larger than 1.

From the final primaries and the original primaries of each projector, an RGB-to-RGB color conversion matrix can be computed which maps incoming input values to the new gamut for any of the projectors.

As a result, an input value such as RGB=(1,0,0) will result in different replacement-RGB values for each display, where the "pure red" that is common for all involved displays is a combination of each display's individual primaries. The common gamut and thus the inverse primary mapping functions for each display can be computed directly from the display's color matrices, once the non-linear behavior has been taken care of.

5.4. Accounting for non-linear behavior

Every display device exhibits a unique response when mapping pixel values of each color channel to luminances. This response depends on all kinds of factors: display technology, projector type, and especially image settings such as black level (a.k.a. "brightness") and intensity (a.k.a. "contrast")[14].

Due to the non-linear behavior of the human visual system, which is more sensitive to luminance changes in dark areas than in bright ones, usually an exponential function, the display system gamma, is used in most displaying and imaging devices. For example, compared to full luminance at pixel level 255, the very common gamma exponent of 2.4 (native to most CRT monitors and TFT displays) exhibits 10% luminance at a pixel value of about 100, instead of a value of 25.

Color computations and black level adjustment need to be performed in a linear space. Therefore the response curves of the display (one for each color channel) have to be measured and any linearly computed color information has to be mapped by the inverse curves before sending the result to the projector. For reliable color calibration, this has to be done for each color channel individually, since there are usually (slight) variances between the channels.

However, stopping here would lead to a distorted representation of displayed images: they would be mapped with a gamma value of 1.0, a linear curve, which will display low intensity image areas too bright. Therefore, we need to apply a gamma curve to our input values to transform *them* into linear space. We could simply use the original display response functions here, to simulate the original behavior: located between input values and projector, the color calibration first applies the projector response curve to transform to linear space, performs the calibration, and reverses the adjustment using the inverse curves.

However, if the response curves differ for the color channels, or there are wash-outs in the response curves themselves due to wrongly adjusted black level or overexcitement in intensity, this would result in unnecessary color shifts for grey levels, or other unwanted behavior.

5.5. Replacing the gamma curve

Actually, there is no real reason to use the display response curve in this step. All that is required is an arbitrary function to map the input values to a linear space - or to be more precise, to something with a similar response as a "well-known" display device, and using the result for computations in linear space. Therefore we can actually choose an arbitrary, ideal function for this step. This could be the averaged response curve of all projectors and color channels, to maintain a rough similarity with the behavior of an uncalibrated projector. The most important part is to use unified curves for all color channels to assure that chromaticity values for grey ramps (or any other color ramp) do not change.

But it is even more convenient to use a smooth function such as the common gamma exponent 2.4. Also a response curve of a digital camera is conceivable, if images taken by that camera are to be displayed.

After transforming the input values into a linear space, color calibration can commence.

5.6. Accounting for the color of black

As stated in 1, digital projectors suffer from a bad black level. It is not possible to improve a bad black, but there are several possibilities how to treat it within a color consistency algorithm.

In any case, the most important factor is that treating "black" as a *color* is a prerequisite for color calibration. Luminance and color can not be treated separately. "Black" may be dark - but it still contains color information and will influence low intensity colors 4.

Since the black color is a constant term, it must not be incorporated in the gamut computations, since then it would scale with the involved color. Instead, it has to be treated separately. The straightforward approach of using the worst black of all projectors and adjust all others to that value is a fairly bad idea. (e.g. reduced contrast).

Our approach is to incorporate the measured CIEXYZ color of each projector's black into the calibration algorithm. Knowing that the projector will add the black color afterwards, we convert the CIEXYZ value into display color space and subtract it from the calibrated output value while still in linear space.

In case of negative values in the result, the affected color channels need to be below the black level for that channel, which is not possible. These values have to be clamped to zero, and the color cannot be represented photometrically correct.

If the resulting values are all positive, we are fine and the resulting color can be displayed *without any loss of contrast and any influence of bad black*. In fact we incorporate the black value as valid part of our regular display color by sending the target color minus black to the projector.

The drawback of this approach is the loss of pixel information that is below the black level - at the benefit of being able to utilize the full common gamut range for all color intensities, and increased contrast.

The visible effect for strong black levels is that with the first approach the image will look all foggy, since the black level is present at every pixel. For the second approach, the "fog" will only be visible in areas that would be darker than the fog level - the rest of the image is as brilliant as with perfect black.

5.7. Transforming input RGB values

The full pipeline to calibrate a given input color value for a projector to allow photometric consistent and colorimetrically calibrated display is as follows (see fig. 2).

5.8. Treatment of calibrated input data

For photometrically consistent display of already calibrated input data, such as images in a calibrated color space, high dynamic range images, etc., the pipeline changes. The initial gamma transformation, as well as the RGB transformation, are no longer required; the color black can be subtracted directly.

However, now it is possible that the resulting color can no longer be displayed with the calibrated projector, because the color lies outside the common color gamut, or it is too bright. For luminances exceeding the range of the common primaries, tone mapping has to be applied to transform all colors into a luminance range the projectors are capable of[19].

Colors outside the gamut need to be clipped to the gamut, using either a straightforward technique such as clipping towards the white point, or a perceptual approach using gamut mapping.

In both cases it might also be useful to mark pixels that cannot be displayed photometrically consistent with pseudo colors, as information to the user.

The pipeline is shortened accordingly, and a new step "Gamut Mapping" is added after the black level subtraction.

6. The Algorithm in Hardware

The described algorithm will create a (4x3) color matrix for every display, which needs to be applied to every pixel's value to get the corrected colors.

As this has to be the last operation before actual display, i.e. after lighting and texturing, it cannot be done as a preprocess on the model data, it actually has to be done at the very end of the graphics pipeline.

The currently available programmable graphics hardware on the PC platform has enough capabilities to support this kind of pixel-pipeline programmability, and we have realized a proof-of-concept implementation using nVidia's RegisterCombiners. It uses three general combiners and the final

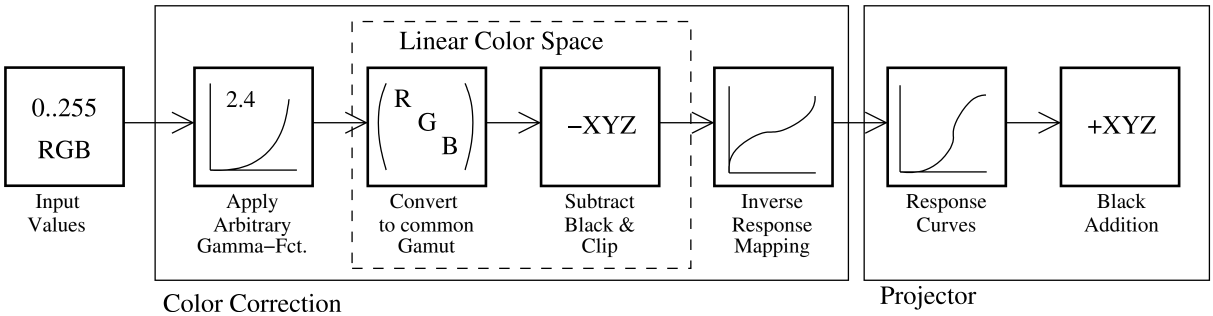

Figure 2: *Steps of the algorithm for uncalibrated inputs*

combiner and thus still leaves some capacity for other effects. It does use the `NV_register_combiners2` extensions though due to the large number of necessary constants and thus only runs on GeForce3 and later nVidia chips.

The same effect can easily be achieved using ATI-specific extensions or the generic Open GL 1.4 Fragment Programs.

Instead of applying the correction to every calculated pixel the same effect can be achieved as a post-process by rendering the image into a texture in the normal way and drawing a screen-filling rectangle using the abovementioned fragment program/register combiner setup for correction. Which version is better depends on the performance of direct-to-texture rendering, which is not very well supported in current drivers.

7. Results

For our test setup we used one pair of the CAVE projectors LC-UXT1, fitted with Infitec filters. To measure chromaticity and luminance, a Minolta Chromameter CS-100 was used.

Due to manual transfer of the color values, we measured each primary only in 10 intensity steps, resulting in 3x10 plus black = 31 CIE-xyL colors per projector.

Table 2 contains measured samples of the two Infitec projectors, the resulting common gamut and preset target white (which is the native white of the projectors, reduced in brightness), as well as the comparison values after color correction. In addition, for better comparison and error estimation the average XYZ distances of each sampled color are given (see table 1).

The photographs of the results before and after calibration were taken with an Olympus E-20p digital camera. The visual and quantitative color deviations of the calibrated images were much smaller in reality, leading us to the conclusion that the spectral response of the E-20p does not match the response curves of the standard observer very closely, especially for the green component. Note that in the calibrated images the mouse cursor is still visible in its original, uncalibrated "white".

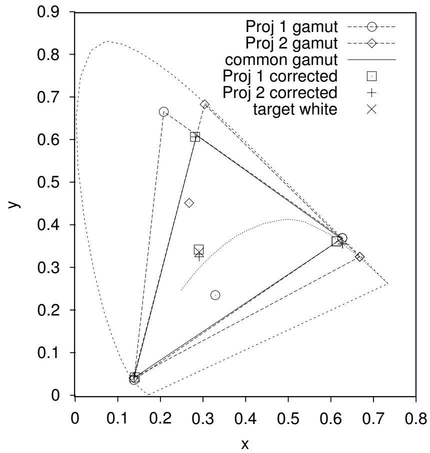

Figure 3: *Common gamut and verification for two test projectors.*

	Projector 1		Projector 2	
Color	Uncalibrated	Calibrated	Uncalibrated	Calibrated
R	91.802	0.320	1.018	0.692
G	6.341	0.641	131.297	1.366
B	113.168	1.143	54.627	1.073
W	175.645	1.480	179.462	1.906
50%	34.715	1.070	41.242	0.422
20%	1.650	0.094	1.538	0.114

Table 1: *Absolute XYZ Differences between projectors and common gamut.*

8. Conclusions and Future Work

In conclusion one can say that digital projectors are not easy to handle and need special care when used for IPTs. When the user stands very close to the screen (e.g. CAVE instal-

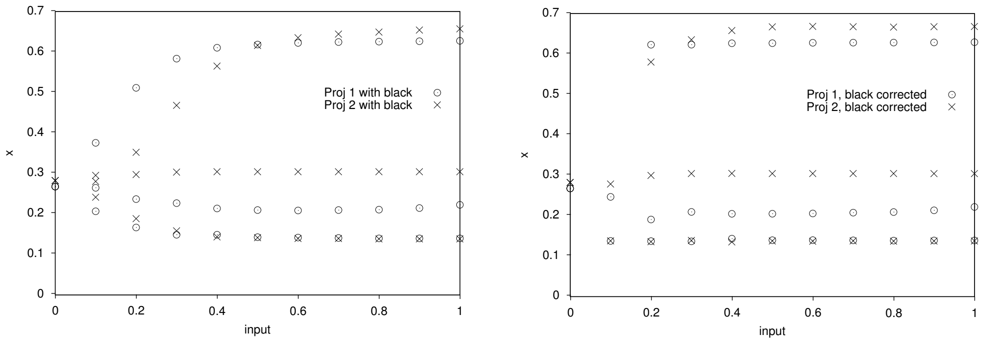

Figure 4: *x color coordinate ramps for red, green and blue, uncorrected (left) and corrected (right). Note that the non-linearity of projector 2 in the upper left (x color coordinate of the green primary) and projector 1 in the lower left (x color coordinate of the red primary) are due to very low luminances of these primaries (cf. figure 2)*

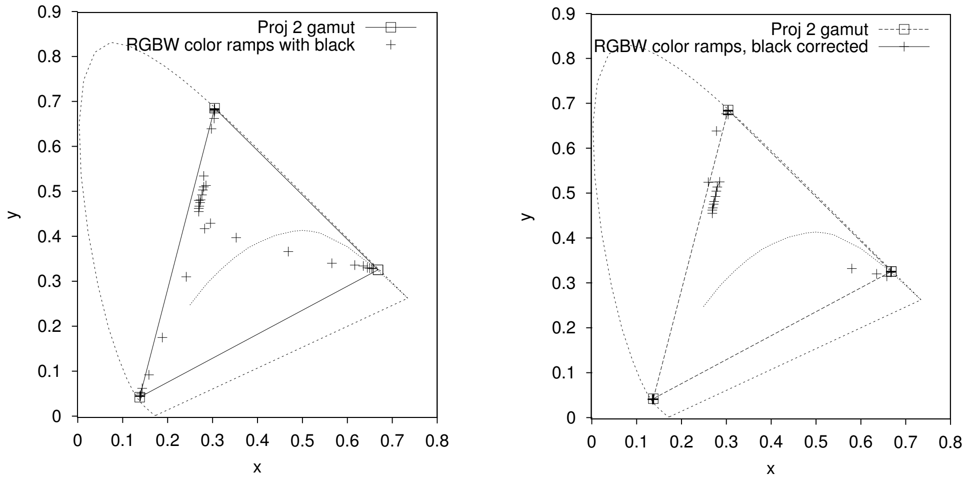

Figure 5: *R, G, B and grey ramp chromaticities without (left) and with black correction (right).*

lations) pixel structure, convergence and focus are major issues.

Although the black level at the first glance seems rather unimportant, the inclusion of its value within the whole calibration process, and its correct treatment as a color is an essential step for consistent and calibrated color representation down to very low luminance pixels. Even for extreme gamut differences, as introduced by the spectral filters, a quite satisfactory color consistency can be achieved. Standard projectors, with only minor differences in gamut, should be able to be matched even better.

The quality of the results can be improved by more accurate measurements, especially for the display system gamma map. As our current photometer does not have a digital output, all values have to be transferred into the system by hand, which is both tedious and error-prone. We have ordered a new photometer, which will allow us to automate the process. It will allow us to measure all interesting values, giving us more data to work with, especially in the low-intensity regions.

For the application of gamma look-up tables, the use of a floating point framebuffer can increase the accuracy and might prevent color banding due to rounding errors.

The developed algorithm is able to adjust different projectors to each other, but in the current form it cannot deal with color and brightness inconsistency within a single projector. Some projectors, especially cheap ones, exhibit a quite noticeable color shift already across a single screen. By putting the correction matrix coefficients into textures it can be made more adaptive; in the extreme case it would be possible to have a different matrix per pixel. We will investigate the localized characteristics of our setups further to see where further refinement stops making sense.

By combining the presented approach with techniques for

	Projector 1			Projector 2			Common Gamut			Measured Result 1			Measured Result 2		
Color	L	x	y	L	x	y	L	x	y	L	x	y	L	x	y
R	118.72	.627	.371	13.41	.668	.328	16.26	.625	.370	16.20	.615	.365	16.50	.629	.359
G	67.82	.209	.668	327.51	.305	.685	57.45	.286	.613	58.30	.283	.610	54.80	.284	.610
B	17.02	.139	.039	11.31	.138	.044	4.14	.140	.045	4.05	.142	.046	4.64	.140	.049
W	202.72	.357	.268	351.51	.269	.458	77.86	.293	.340	77.70	.292	.346	76.00	.293	.329
50%	39.22	.357	.265	82.81	.276	.495	14.75	.293	.340	13.50	.292	.336	13.90	.293	.330
20%	2.18	.307	.223	4.95	.286	.527	1.63	.293	.340	1.63	.285	.360	1.47	.283	.320
0%	0.28	.268	.184	0.49	.283	.420									

Table 2: *Measured primary colors and a number of grey levels for the uncalibrated projectors, the common gamut and target grey levels, and the verification after calibration.*

uniform brightness [11] and by automating the whole process with techniques presented in [11], we are planning to provide a simple-to-use tool for calibration and color matching of multi-projector display systems, such as tiled displays or CAVEs.

References

1. Han Chen, Grant Wallace, Anoop Gupta, Kai Li, Tom Funkhouser, and Perry Cook. Experiences with scalability of display walls. In *IPT2002*, 2002.

2. Y. Chen, D. Clark, A. Finkelstein, T. House, and K. Li. Automatic alignment of high-resolution multi-projector displays using an un-calibrated camera. In *Proceedings of IEEE Visualisation 2000*, 10 2000.

3. Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the cave. *Proceedings of SIGGRAPH 93*, pages 135–142, August 1993. ISBN 0-201-58889-7. Held in Anaheim, California.

4. M. Czernuszenko, D. Pape, D. Sandlin, T. DeFanti, G. Dave, and M. Brown. Immersadesk and infinity wall projection-based virtual reality displays. *Computer Graphics*, 31(2):46–49, 1997.

5. Paul Woodward et al. Powerwall. http://www.lcse.umn.edu/research/powerwall/powerwall.html, 1994.

6. Tom Funkhouser and Kai Li. Large-format displays. *IEEE Computer Graphics and Applications*, 20:20–21, 2000.

7. Gerriets GmbH. Opera milchig matt specifications, 2003.

8. M. Hereld, I.R. Judson, and R.L. Stevens. Introduction to building projection-based tiled display systems. *IEEE Computer Graphics and Applications*, 20, 2000.

9. H. Jorke. Infitec - Wellenlängenmultiplex Visualisierungssysteme, 2002.

10. Wolfram Kresse and Dirk Reiners. Can we trust that image? Photometric Attributes of Current Projection Systems. In *IPT2002 Proceedings*, Orlando, March 2002.

11. Wolfram Kresse, Frank Schoeffel, and Stefan Mueller. An autocalibraion tool for the photometric and colorimetric consistency of ipts. In *IPT2000*, Ames, 2000.

12. Aditi Majumder, Zhu He, Herman Towles, and Greg Welch. Achieving color uniformity across multi-projector displays. In T. Ertl, B. Hamann, and A. Varshney, editors, *Proceedings Visualization 2000*, pages 117–124, 2000.

13. Aditi Majumder and R. Stevens. Lam: Luminance attenuation map for photometric uniformity in projection based displays. In *Virtual Reality and Software Technology 2002*, 2002.

14. Charles Poynton. Color technology, 1997.

15. R. Raskar, M. S. Brown, R. Yang, W.-C. Chen, G. Welch, and H. Towles. Multi-projector display using camera-based registration. In *Proceedings of IEEE Visualisation 1999*, 10 1999.

16. Ramesh Raskar, Greg Welch, Matt Cutts, Adam lake, Lev Stesin, and Henry Fuchs. The office of the future: A unified approach to image-based mpdeling and sptially immersive displays. In *Proceedings of ACM SIGGRAPH 1998*, 1998.

17. Frank Schoeffel, Wolfram Kresse, Matthias Unbescheiden, and Stefan Mueller. Do IPT systems fulfill application requirements? a study on luminance on large-scale immersive projection devices. In *IPT1999*, Stuttgart, 1999.

18. R. Surati. *Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays*. PhD thesis, M.I.T., 1999.

19. Greg Ward-Larson, Holly Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high-dynamci range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, 1997.

Figure 6: *Comparisons between uncorrected (left) and corrected (right) image pairs. Note that the screen the pictures were taken from has a pretty strong highlight (especially visible in the top row).*