Super-Resolution Composition in Multi-Projector Displays

Christopher Jaynes and Divya Ramakrishnan

Metaverse Lab, Dept. of Computer Science
University of Kentucky
jaynes@metaverselab.org

Abstract

Recent advances in multi-projector research have made large-format tiled display walls feasible. Resolution of the tiled display is defined by the resolution of the component projectors and the tiling. In domains where display surface size is constrained or the number of pixels per unit area on the display must be increased, traditional tiling may not be feasible. We propose a multi-projector method that exploits overlapping regions to achieve increased resolution and contrast on the display surface.

Given a set of overlapping projectors and a target image, component images are derived and then rendered by each projector resulting in an image on the display surface exceeding the resolution limits of the individual projectors. The relative viewing geometry between the projectors is taken into consideration while computing the component images that should be projected. Sub-pixel shifts between the displays are viewed as integer disparities in the high-resolution target image. Results demonstrate that for a variety of target images, there is perceptible resolution improvement on the display surface. Super-resolution image composites have a significantly higher sampling rate on a display surface without a corresponding increase in display area.

1. Introduction

This work introduces the Super-Resolution Composition problem for multi-projector display and presents an initial approach to achieving Traditional multi-projector high-resolution is achieved by tiling component projectors together, calibrating their relative geometry [4,8,19,21], potentially their radiometric differences [15,24], and rendering a coherent image across all projectors. Due to significant advances in automatic calibration, tiled display is feasible, very large-format displays are fairly straightforward to deploy in unconstrained environments, and novel tiled configurations are being explored [25].

This work focuses on the super-resolution composition (SRC) of multiple overlapping and calibrated projector frustums. In traditional multi-projector displays, accidental overlap occurs quite frequently and SRC can be used to increase resolution and display contrast in these regions. For displays that require high-resolution but have insufficient surface area to support tiling, projectors can be fully overlapped to exploit both the increased resolution and contrast that SRC affords.

Recent research advances in multi-projector displays both from the computer vision and graphics communities inspire this work [5,9,10,11,12,14,26]. Efforts have focused on building large-scale displays from clusters of projectors that automatically calibrate [4,12,20,22], and correctly render a single, coherent tiled image [5,10,12,14,19]. The work here is an alternate approach to achieving higher resolution display from multi-projectors.

Multi-projector displays are loosely configured and may have significant regions of overlap. Several researchers have addressed these overlapping regions to provide intensity blending and color uniformity [15,19,24]. Recently, these regions have been exploited to actively remove shadows and provide increased contrast ratio [3,16].

Indeed, regions of the display surface illuminated by more than one projector can be viewed as an Optical Framebuffer in which operations may take place by appropriately adjusting images in the different projectors. This Framebuffer Optique, and its potential use for intensity and contrast enhancement, as well as cooperative rendering of linear computer graphics operations (e.g. light source compositing) was first noted by [16]. This paper demonstrates that the Framebuffer Optique can be exploited to produce high-resolution images by compositing low-resolution components.

Solution to the SRC problem has direct utility in multi-projector display environments as is demonstrated here. Images that exceed the resolution limits of a single projector can be more accurately visualized using multiple projectors and the SRC method without increasing display size. This dramatically increases the sampling rate of the rendered images *on the display surface* to better approximate the fidelity of the human eye. High-resolution insets in a projected display are also made feasible through this technique.

The technique requires at least two overlapping projectors whose view frustums illuminate a planar display surface.

During an initial calibration phase, a camera is used to recover a homography between each projector and a base projector frame. A base projector frame is aligned with the high-resolution space and only differs in its resolution. Sub-pixel overlap defines pixels in the high-resolution target frame. Figure 1 depicts this situation.

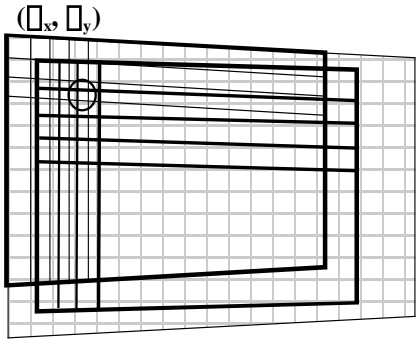

Figure 1: Multiple projectors illuminate a planar surface. A single projector defines the Framebuffer Optique, whose high-resolution pixels (shown in circle) are defined by the overlapping sub-pixel shifted components.

Because each projector is related to the target frame via a general homography, the relative two-dimensional shift and sampling rate will change across the display. The component homographies are approximated by a set of two-dimensional sub-pixel shifts (referred to as the linear shift matrix) that represents the sub-pixel disparity of one projector with respect to the target image reference frame, for a given image sub-region.

Component images are then estimated in the frequency domain where the target image is sub-sampled and phase shifted according to the sampling rate and shift matrix for each component. The resulting amplitudes and phase values are then optimized according to a cost function that measures the difference between the target image and the image that results from adding the multiple sub-sampled, shifted components together.

2. Previous and Related Work

Super-resolution composition is directly related to a significant research history on super-resolution reconstruction from image sequences. Many of the techniques described here are related to the super-resolution reconstruction problem and are closely related to the frequency domain approach [2,7,18,23]. For a concise overview of the more common approaches to super-resolution reconstruction, the reader is referred to [2,6].

Although multi-projector display systems are the focus of an active research community, there has been relatively little work related to exploiting projector overlap. Initial research focused on computing overlap regions automatically from camera-based calibration techniques [4,12,19,22] and monitoring active displays for reactive calibration [21]. Other methods seek to address overlap regions through intensity blending [9,19], geometric correction [4,12,14,19,22], and color balancing [15,24]. These techniques are important in that loosely configured, automatically calibrated tiled displays are almost certain to contain regions of projector overlap that must be addressed. In contrast to these efforts, we exploit the utility of overlapping regions rather than seek to remove or attenuate their effects.

Only recently, methods for exploiting projector overlap have been introduced. Detection of radiometric changes on the display surface can be used to intensify or attenuate different projectors to actively remove shadows on an interactive display [3,13]. Welch. et. al. demonstrated that the rendering problem can be decomposed into several components and added together on the display surface to support distributed computation of lighting and rendering effects [16]. The Framebuffer Optique is described in the same work as a virtual Framebuffer that can be accessed via manipulation of the component images and added together through image superposition. We demonstrate how the Framebuffer Optique can be exploited to produce high-resolution image overlay by solving the super-resolution composition problem.

3. Image Decomposition and Compositing

There are three steps in deriving an appropriate set of component images. First, each projector must be calibrated to the target reference frame, in order to determine the relative sub-pixel shift between each projector and that frame. Once calibrated, an initial set of component images is derived via the known shifts and sub-sampling rates for each component. This phase of the algorithm can be performed in parallel for each projector. Given these initial estimates, a global optimization step minimizes the difference between the sum of the components and the target image.

3.1 Calibration

The goal of calibration is to derive an accurate mapping from each projector's framebuffer coordinates to the high-resolution target frame. This mapping must be accurate to less than a pixel and presents a significant challenge in practice. In order to compute the sub-pixel disparity between projectors and a target frame, pixel correspondences between any two projectors must be computed to far less than a pixel accuracy and nonlinear effects due to both camera and projector optics must be known and corrected.

The epipolar relationship between pairs of projectors and a known surface model is sufficient to compute pixel correspondences for every pixel in the Framebuffer. In

the results shown here, however, the display surface is constrained to be planar and the full projective relationship between any two devices, *i* and *j* can be modeled as a 3x3 homography matrix that maps pixels in projector *j* directly to pixels in projector *i,* through the display plane. The homography can be automatically computed given a sufficient set of matchpoints between the two projectors in question.

Prior to estimating the calibration matrix, the radial distortion coefficients, k_1 and k_2, of the camera and all projectors are independently estimated in an offline process. Prior to estimating the linear homography between each camera and projector, points in each device are unwarped using the recovered distortion coefficients to compute point positions in an undistorted frame of reference. Recovered homographies, then, map points from one undistorted frame to another.

Matchpoints are then computed between a camera that observes the display surface, and each projector in the display. For each projector, a set of Gaussian target fiducials centered at randomly selected framebuffer pixels are iteratively displayed. The target is captured in the camera and a match-point pair, using the undistorted coordinates, is stored. The sub-pixel location of the target in the camera is computed through an optimization routine that estimates seven parameters of a distorting homography that is specific to each matchpoint pair. We have shown that this technique can recover match-points at close to $1/5^{th}$ pixel accuracy and the reader is encouraged to read [20] for more details and analysis of the subpixel matchpoint finding algorithm.

Given a sufficient number of correspondences (25 for the results shown here), a homography between each projector *p* and the display camera *c*, H_p^c, is then computed using linear least squares estimation. These homographies represent the mapping between each projector and the camera used to calibrate the display.

Each projector to camera matrix is then converted into a homography between each projector and the target image reference frame. An arbitrary projector Framebuffer is selected as the base projector reference frame. Although the target frame and base projector Framebuffer are of different resolutions (the target is significantly higher), this projector defines the target space up to an unknown scale by assuming that the base Framebuffer is axis-aligned with, and shares the origin of the target image space. Therefore, the relationship between any projector and this target image frame, H^b_i, can be written as a composite homography from the projector (*i*) to the camera (*c)*, and then to the base projector frame (*b*):

$$\mathbf{H}_i^b = \left(\mathbf{H}_b^c\right)^{-1}\mathbf{H}_i^c \qquad (1)$$

This situation is depicted in Figure 1. Projectors overlap on the display surface and are each shifted with respect to the base frame. Shifted overlapping pixels give rise to the higher resolution space where sub-pixel shifts define integer pixels in the high-resolution target frame. For results shown here we assume the sub-pixel calibration accuracy up to $1/4^{th}$ of a pixel resulting in resolutions in the target frame that are 4-times those of component projectors. We assume a constant sampling rate from one projector to the target frame. Although general configuration of projectors will lead to non-uniform sampling rates (particularly in cases of extreme off-axis projection and skew), a mean sampling rate is a good approximation for common projector setups.

In order to demonstrate the calibration accuracy that can be achieved under controlled conditions using accurate, undistorted camera matchpoints, a homography was derived between two overlapping frustums. A high-resolution image was then captured very close to the display surface using a digital still camera. Figure 2 shows two 4x4 pixel grids drawn on the high-resolution image corresponding to the two overlapping projectors. The white grid, corresponding to projector 1 pixel boundaries, was drawn by hand for demonstration purposes while the black grid was defined by the homography from projector two's frame to the frame of projector 1. The white circle on the image depicts the center pixel *p* in the framebuffer of projector 2. This user estimated center can then be compared to the sub-pixel disparity predicted by the calibrated homography (centers of the black squares). In this case, calibration predicted a subpixel shift of *p* with respect to *q* of (0.09,-.12) while inspection reveals a shift of projector 2 with respect to projector 1 of (0.13,.11).

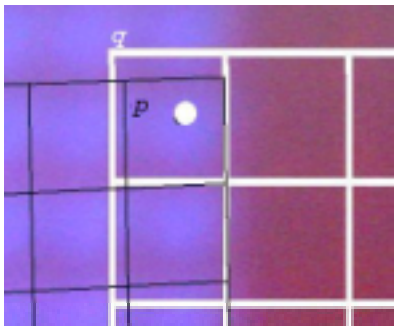

Figure 2: Calibration accuracy for overlapping projector frustums. White lines show pixel boundaries for base projector. Circles correspond to pixel centers estimated by human subject while center of actual grid predicted by homography (shown as black lines) is approximately off by $1/5^{th}$ of a pixel.

According to our formulation, the low-resolution component images are modeled as sub-sampled versions of the image target with a *uniform* two-dimensional shift and sampling rate with respect to the target frame of reference. Clearly, the projective warp between two projectors describes a more general displacement than uniform shift. In order to derive an appropriate component image, without the undue computation that a per-pixel solution would require, the homography can be approximated by a set of two-dimensional shift vectors between a projector and the target frame.

The homography between projector *i* and the base frame is decomposed into a *linear shift matrix* that represents two-dimensional offsets between the projector and the base reference frame. Each entry in the shift matrix corresponds to a region in the target reference frame for which the two-dimensional offset is assumed to be constant. Once computed, the shift matrix replaces the more general homography and regions in the component frame are related to the target frame through a constant, two-dimensional offset. We present a method for constructing a linear shift matrix S_i^B from a given homography H_i^B given target error tolerance ε.

The two-dimensional disparity δ, between a component projector reference frame *r* and target frame *b* is written as the difference between the locations of a pixel in frame *r* the same pixel in frame *b* (given by the known homography).

$$\delta = p - \mathbf{H}_r^b p \tag{2}$$

The disparity in *x* and *y* directions (δ_x, δ_y) is independently given by:

$$\delta_x = p_x - \frac{H_1 p_x}{H_3 p_x}$$
$$\delta_y = p_y - \frac{H_2 p_y}{H_3 p_y} \tag{3}$$

where $\mathbf{H}_k$ is the vector formed by the k^{th} row of the matrix $\mathbf{H}_r^b$. If we assume that δ_x and δ_y are independent, the disparity value is a linear function of the pixel position in the component frame.

As *x* ranges from zero in the component projector to x_{max}, the resolution of projector in the *x* direction, the disparity values will vary in accordance with the line equation given by Equation 3. This line is divided into *k* equal regions such that the disparity values in the region are all within ε_x of one another. Conceptually, these *k* regions are columns in the component image that will use the same approximate *x*-shift values, $\sim\delta_x$, for the purposes of deriving the component image corresponding to pixels contained in that column.

Given the line equation for independent disparities in the *y* direction (Equation 3), a similar process divides the component frame into rows of uniform *y*-disparity with error tolerance ε_y. These regions are combined to produce regions in the component image containing approximate values for two-dimensional shifts that are within $\varepsilon = \sqrt{\varepsilon_x^2 + \varepsilon_y^2}$ of the values represented in the actual homography. Therefore, for a given error tolerance ε (fixed to be .2 pixels for the results shown here), the homography can be decomposed into areas of uniform disparity. These region-based two-dimensional approximate shifts, $(\sim\delta_x, \sim\delta_y)$, and the corresponding offset of the region itself, (O_x, O_y), are used to derive the a component image for a single projector. It should be noted, that, as error tolerance approaches zero, the region-based two-dimensional shifts simple become the two dimensional shifts given on a per-pixel basis by the homography itself and no accuracy is lost. We select an error tolerance that reduces computation while still retaining a perceptually good result.

Figure 3 shows ten regions corresponding to a 5x2 shift matrix, computed from the homography between a projector to the base frame for a two-projector setup. Given a fixed error tolerance, the number of regions computed by this process is related to the amount of projective (off-axis) distortion induced by the viewing geometry.

3.2 Component Image Estimation

Component image estimation is performed in the frequency domain where initial images are first estimated and then optimized. We do not construct component images in the spatial domain because overlapping pixel intensities and the constraints they represent are difficult to characterize. We are strongly motivated to formalize the problem in the frequency domain because, a *perceptually good* image must take into account models of the human visual system (HVS). Successful perceptual models, in particular those that assign importance to particular aspects of an image, are defined as weighted frequency filters. Because the problem is ill-posed (i.e. we can only approximate a target image with a given projector setup), these perceptual models must be incorporated into an objective function during the optimization phase of reconstruction. Although, we do not yet incorporate an HVS model, we expect that improved results will require an appropriate model of perception.

A given target image is first converted to the resolution of the Framebuffer Optique, defined by the sub-pixel shift pattern recovered in the calibration phase. This target image $I(x,y)$, is converted to a corresponding discrete Fourier transform (DFT) $F_I(u,v)$.

A component image for a particular projector is estimated in two phases. First, sub-components for the *n* regions of uniform disparity are derived. Following that, sub-components are combined to form a single component image for the projector.

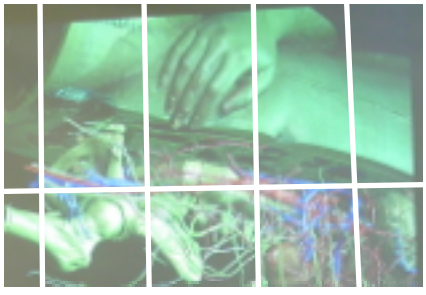

Figure 3: Automatic decomposition of a homography into regions of uniform two-dimensional disparity. White lines denote boundaries of uniform 2D displacement.

The target image DFT is sub-sampled at a rate of R_x and R_y based on mean sampling-rate derived from the calibration phase for a particular component projector with respect to the target frame. The DFT of the target image $F_T(u,v)$, and the sampled component image, $F_s(u,v)$, are related via aliasing:

$$F_S[u,v] = \sum_{p=0}^{N-1} \sum_{q=0}^{M-1} F_T\left[\frac{k}{MR_x} + pf_{sx}, \frac{l}{NR_y} + qf_{sy}\right] \quad (4)$$

where $f_{sx}=1/R_x$ and $f_{sy}=1/R_y$ are the two-dimensional sampling rates in x and y directions. Again, non-uniform sampling rates can be supported by this framework and can be estimated on a per-region basis in a manner similar to that used to estimate two-dimensional disparity. Future work will address this issue.

The sub-sampled DFT signal is shifted, for each region r, by the corresponding entry in the Linear Shift Matrix, $(\sim\Delta_x^r, \sim\Delta_y^r)$, plus the offset (O_x^r, O_y^r) for that region. The shifting property of the Fourier transform relates spatial domain translations to a phase shift in the frequency domain as:

$$F_{\Delta r}[u,v] = e^{j2\pi((\sim\Delta_x^r + O_x^r)u + (\sim\Delta_y^r + O_y^r)v)} F_S(u,v) \quad (5)$$

Equation 5 holds only for a stationary signal and shifting the signal for a finite region of the target image may result in invalid frequency coefficients due to boundary problems. In practice, a region of size w x h is extended by the magnitude of the phase shift to include neighboring pixels. For regions on the image edge, the edge values of the image are copied into a border whose size is defined by the magnitude of the phase shift. The resulting Fourier series $F_{\Delta r}[u,v]$ is the frequency space representation of the sub-sampled, shifted image for a specific region in the component image. Each of the n regions is then composed into a single Fourier Series using the distributive property of the Fourier transform over addition. This results in a single Fourier series for a component projector, p:

$$F_\Delta^p[u,v] = \sum_{r=0}^{N} F_{\Delta r}[u,v] \quad (6)$$

3.3 Global Optimization

Once an initial set of component images have been estimated, they must be adjusted to more closely approximate the contents of the target image. There are a number of ways to formulate the problem of determining the component image contents given their relative position to the target reference frame and a target image. Figure 4 depicts the theoretical problem to be solved. Pixel values for projector 1, shown in black, must be determined with respect to the constraint that the total energy supplied by all projectors sums to the expected value for the high resolution pixel under consideration, shown as k_1 in the image. We are exploring several different approaches including linear programming and direct optimization of the free parameters (see Conclusions). Here we present a heuristic-based iterative relaxation method used for the results here.

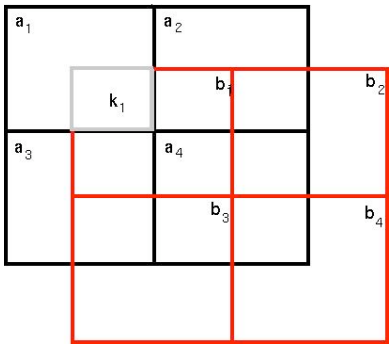

Figure 4: Ideal two-projector composite case. Pixel values a_1 and b_1 must be determined so that $k_1=a_1+b_1$. Because the value of b_1 also influences the value of k_2, k_3, and k_4, the pixel values for a single component are not independent.

Component images, computed for each overlapping projector, are derived independently. These initial components do not take the additive nature of the superposition process into account and must be adjusted in a global optimization phase to achieve increased contrast. In order to do so, the intensity values contained in the target image at each pixel are multiplied by the number of projectors that overlap that pixel. This defines a new target of increased contrast that is approximated by the projector components during the optimization phase.

Each derived component image for each of the k projectors, F_Δ^p, is treated as the initial estimate for an iterative algorithm that seeks to minimize the difference between the sum of the different components and the image target. We model the superposition process on the display surface as additive and minimize:

$$\left\| F_T[u,v] - \sum_{i=0}^{k} F_\Delta^i[u,v] \right\|^2 \quad (7)$$

In future work, Equation 7 will be biased by a perceptual model that assigns importance to particular frequency events in the reconstructed image. The unbiased error metric, must be augmented by additional constraints if we are to converge to component images at are reasonable. Several constraints are available and we are still exploring

how the choice of constraints at this phase of the algorithm influences composite image quality. Currently, the outermost projector boundary pixels are fixed to be exactly *x/k* of the intensity of the high-resolution pixel it overlaps (For example, pixel a_1 in Figure 4) where *x* is the target value and *k* the number of projectors that overlap the super-resolution pixel. This defines a set of pixel values that can no longer be adjusted.

The remaining component image pixels, *p*, are then visited in random order in each image, the largest difference between the high-resolution target values contained within component pixel *p* and their current values (defined by the sum of the component pixel and its overlaps) is corrected.

Correction can either involve directly adjusting the low-resolution component pixel to match the target pixel value of the largest error exactly or may involve a weighted adjustment of the component pixel intensity. The pixel value is adjusted to exactly match the desired value for the target pixel if, in doing so, the maximum error of the other high-resolution pixels contained within *p* do not exceed the previously measured maximum error.

If direct correct if pixel *p* yields an error that is greater than the maximum error contained at *p* prior to adjustment, the intensity of the component pixel in question is adjusted according to, $\Delta I = \lambda D_t - (1-\lambda) D_{t+1}$. Where, lambda is scalar weight that can dampen local changes in component images and will ultimately control the amount of variation in the component images from the initial estimates. D_t and D_{t+1} are the maximum pixel differences between the intensity of the component pixel and any high-resolution pixel it overlaps both before the value is adjusted and after.

Starting from the periphery, pixels are iteratively adjusted according to the algorithm in all component images that can influence the value of a high-resolution target. This algorithm is applied iteratively to all images until either until it converges and no pixel values are adjusted in a single pass over each image component or the global error function is smaller than a predefined threshold. Determining this threshold for a given calibration accuracy and target image is a matter of current work and in practice, the global error measure is observed by a human operator who determines when to terminate the algorithm when it appears to be no longer improving.

4. Results

The technique was evaluated for two and three projector scenarios where the projectors were configured to have significant overlap and each projector's optic axis was oriented to be within approximately 15-degrees perpendicular to the display plane. Distance from the projectors varied for different test cases but ranged from 6 feet to 12 feet from the display. The technique was tested for 1024x768, 800x600, and 640x480 component projector resolutions.

Projectors were calibrated using 25 matchpoints per projector and the projector to target image homography was recovered in each test case. Shift matrices, and the corresponding image regions were derived from the component homographies to within 0.25 pixel error (estimated using a technique similar to the one discussed in Section 3.1).

For each projector, an initial component image was derived using the techniques described in Section 3.2. Each of these images (still in the frequency domain) was then provided to a single machine for global optimization (see Section 3.3). The resulting Fourier series for each projector was then converted to a spatial signal by each of the *k* different machines to produce a component image that was loaded into the Framebuffer for display.

The time required to derive initial component images is negligible compared to the time spent during the global optimization phase. Optimization time is directly related to the number of parameters being optimized (defined by the number and resolution of the components) as well as the complexity of the search space defined by the target image. Current times preclude using the SRC technique for video-rate imagery on a multi-projector display. We expect that a closed-form approach to determining the optimal components given the setup is feasible and is the subject of future work.

Since the resulting image is only present on the display surface and is not represented digitally, quantitative comparison of the SRC result against the target image is not possible. Instead we present several examples of reconstructed images by capturing a high-resolution image of the display surface (sometimes at very close range) to demonstrate the perceptual result of the technique.

Figure 5 depicts a close-up image of the "pinky finger" taken from the Human Body Dataset. A 2048x2048 resolution still image of a human torso (shown fully in Figure 3). The SRC image shown in Figure5b contains greater contrast, less aliasing at boundaries, while enhancing high-frequency features such as the point tip on the fingernail.

(a) (b)

Figure 5: Close-up of pinky finger from 2048x2048 resolution image (see entire image in Figure 3). (a) Single projector at 1024x768 resolution. (b) Two composite images computed using SRC method. Note that the surface pattern is more visible in image (b) due to increased contrast and resolution. The sharp tip structure on the fingernail is present in (b) and not in (a).

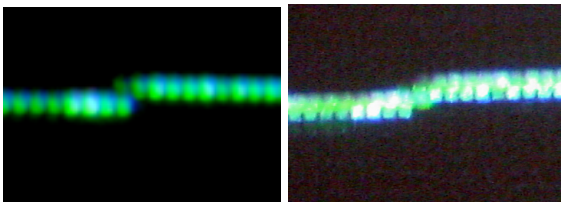

(a) (b)

Figure 6: Close-up view of line (veins) segments from a larger image from the human-body dataset (see Figure 3). Target resolution is 2148x2148. (a) Single projector view of image. Aliasing of line-segment is a result of down sampling in the projector. (b) Super-resolution reconstruction.

Composite results are perhaps best observed at the pixel level resulting images on the display surface. Figure 6 shows an extreme close-up of a curve from the Visible Human dataset. The image corresponds to a vein from an image that is 2148x968 pixels in resolution.

A computer desktop containing text and images at a resolution of 1280x1024, was rendered using a single projector at 1024x768 and reconstructed using the SRC method. Results are shown in Figure 7.

A close-up view of a section taken from a 2048x2048 resolution aerial image dataset shows the improved image quality that SRC affords. Figure 8 shows a close-up of a curving sidewalk region and several pedestrians taken from the "Aerial" dataset experiment. Previously unseen structure (such as the white shirt) is visible in the SRC image.

5. Conclusion

We have introduced the Super-Resolution Composition reconstruction problem and demonstrated how an accurate decomposition of a high-resolution target image into its overlapping, sub-pixel, shifted components can be used to improve the resolution of a projected display. Results show perceptible improvement for two and three-projector overlapping displays given a variety of target images.

In the near future, we expect to implement the SRC solution in hardware (particularly the Fourier transforms), and devise a more efficient optimization method using better initial estimates. We are also exploring direct methods to recover composite images. Under certain constrained configurations, the constraints on component images can lead to a large and sparse system of linear equations that can be solved faster than the iterative method described here. In this way, video and interactive applications can be supported on a super-resolution display. In addition, we are working on faster sub-pixel calibration techniques that allow users to "steer" projectors to regions of imagery where high-resolution insets are desired.

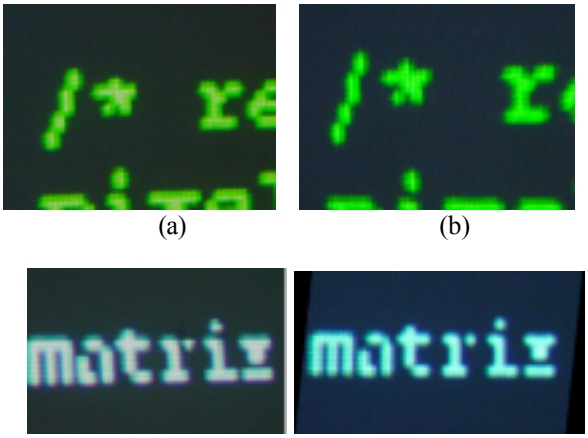

(a) (b)

(c) (d)

Figure 7: Close-up views of text projected on display surface. (a) Single projector at 1024x768. Sampling rate is too low and details (top of asterisk, for example) are lost. (b) Super-resolution composite. (c) Single projector at 1024x768 close-up of word "matrix". Compare circle smoothness and completeness in letters "r" and "i".

In longer-term research, we are exploring other operations that may be able to decompose the target into a set of component images that are ultimately overlaid in the Framebuffer Optique. Examples include distributed rendering and blending of color and intensity.

6. Bibliography

[1] S. Baker and T. Kanade, "Limits on super-resolution and how to break them", *Int. Conf. Computer Vision & Pattern Recognition*, volume 2, pp. 372--379, 2000.

[2] S. Borman and R. L. Stevenson, "Super Resolution from Image Sequences - A Review", *Proceedings of the 1998 Midwest Symposium on Circuits and Systems,* Notre Dame, IN, 1998.

[3] H. Chen, R. Sukthankar, G. Wallace and K. Li. "Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees", *IEEE Visualization (Vis2002),* Oct. 2002.

[4] C. Cruz-Neira, D. Sandin, and T. Defanti, "Surround-screen projection based virtual reality: The design and implementation of the CAVE", *Proc. Of ACM Siggrgaph*, 1993.

[5] G. Demoment, "Image reconstruction and restoration: overview of common estimation structures and problems", *IEEE Transactions on Signal Processing*, vol. 37, no. 12, pp. 2024-2036, December 1989.

[6] M. Elad and A. Feuer, "Restoration of Single Super-Resolution Image From Several Blurred, Noisy and Under sampled Measured Images", *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1646-58, December 1997.

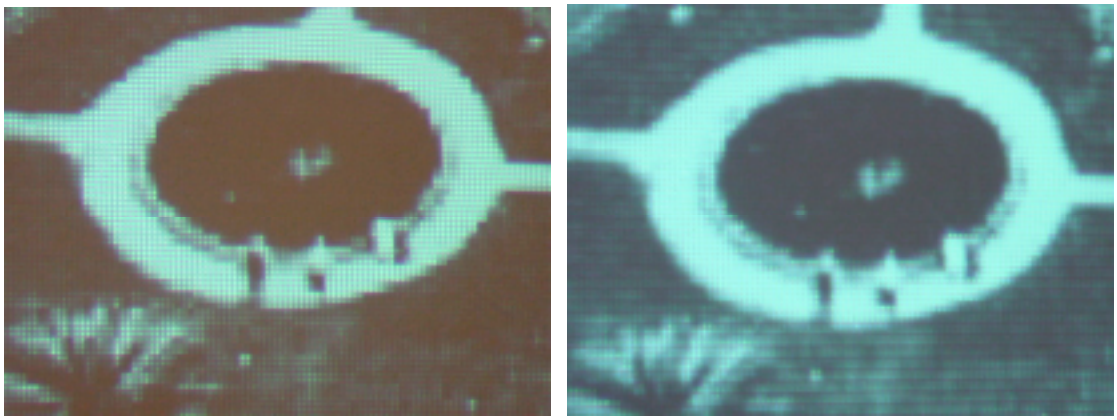

(a) (b)

Figure 8: Result of super-resolution composition on the aerial image dataset. Close-up view of pedestrians on a sidewalk. (a) Single projector at 800x600 resolution. Note aliasing problems on sidewalk in (a) and pedestrian's torso at left is not fully resolved in the imagery due to insufficient sampling. (b) Two-projector superposition result. Compare with results in (a).

[7] H. Foroosh (Shekarforoush), J. Zerubia and M. Berthod, ``Extension of Phase Correlation to Sub-pixel Registration'', *IEEE Trans. Image Processing*, vol. 11, Issue 3, pp. 188-200, 2002.

[8] M. Hereld, I. R. Judson, R. L. Stevens, "Introduction to Building Projector-Based Tiled Display Systems", IEEE Computer Graphics and Applications, pp. 22-28, 2000.

[9] G. Humphreys, M. Houston, Ren Ng, R. Frank, S. Ahern, P. D. Kirchner and J. T. Klosowski, "Chromium: A Stream Processing Framework for Interactive Rendering on Clusters", *ACM Transactions on Graphics*, vol 21(3), pp. 693-702, Proceedings of ACM Siggraph, July 2002, San Antonio, TX.

[10] G. Humphreys and P. Hanrahan, "A Distributed Graphics System for Large Tiled Displays," *IEEE Visualization*, 1999.

[11] C. Jaynes, S. Webb, and R. M. Steele, "A Scalable Framework for High-Resolution Immersive Displays", *International Journal of the IETE, Special Issue on Visual Media Processing,* Vol. 48, No. 3&4, pp. 273-285.

[12] C. Jaynes, S. Webb, R.M. Steele, M. Brown, and W.B. Seales, "Dynamic Shadow Removal from Front-Projection Displays", *IEEE Visualization*, San Diego, 2001.

[13] Z. Li and A. Varshney, "A Real-Time Seamless Tiled Display for 3D Graphics", Proceedings, Seventh Annual Symposium on Immersive Projection Technology, March 24 - 5, 2002, Orlando, FL

[14] A. Majumder, Z. He, H. Towles, and G. Welch, "Achieving Color Uniformity Across Multi-Projector Displays", *11th Ann. IEEE Visualization Conference,* 2000.

[15] A. Majumder and G. Welch, "Computer Graphics Optique: Optical Superposition of Projected Computer Graphics", *Fifth Immersive Projection Technology Workshop*, *Seventh Eurographics Workshop on Virtual Environments,* Stuttgart, Germany, Springer-Verlag, May 16-18, 2001.

[16] S. McCanne and M. Vetterli and V. Jacobson", "Low-complexity Video Coding for Receiver-driven Layered Multicast", *IEEE Journal on Selected Areas in Communication*, 1999.

[17] Nguyen, N., P. Milanfar, G.H. Golub, "A Computationally Efficient Image Super-resolution Algorithm", *IEEE Transactions on Image Processing*, vol. 10, no. 4, pp. 573-583, April 2001.

[18] R. Raskar, G. Welch, and H. Fuchs (1998), "Seamless Projection Overlaps Using Warping and Intensity Blending", *4th Intr. Conf. on Virtual Systems and Multimedia (VSMM)*, November 18-20, Gifu, Japan.

[19] M. Steele and C. Jaynes, "Parametric Subpixel Matchpoint Recovery with Uncertainty Estimation: A Statistical Approach", IEEE Workshop on Statistical Analysis in Computer Vision in Conjunction with CVPR 2003

[20] R.M. Steele, S. Webb, and C. Jaynes, "Monitoring and Correction of Geometric Distortion in Projected Displays", *International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, Plzen, Czech Republic, 2002.

[21] R. Surati, "Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays", *MIT*, 1999.

[22] B. C. Tom and A. K. Katsaggelos, "Reconstruction of a High-Resolution Image by Simultaneous Registration, Restoration, and Interpolation of Low-Resolution Images", *Proc. 1995 IEEE International Conf. on Image Processing*, pp. II-539-542, Washington, DC, Oct. 1995.

[23] G. Wallace, H. Chen, and K. Li, "Color Gamut Matching for Tiled Display Walls", In Proc. *Immersive Projection Technology and Virtual Environments*, pp. 293-302, 2003.

[24] S. Webb and C. Jaynes, "Design and Implementation of the Oracle: A High-resolution Object Visualization Device", Technical Report, Dept. of Computer Science, University of Kentucky, 2003.

[25] G. Welch, H. Fuchs, R. Raskar, M. Brown, and H. Towles (2000), "Projected Imagery In Your Office in the Future", *IEEE Computer Graphics and Applications*,