

Local Methods for Measuring Tiled Display Alignment

Mark Hereld

Mathematics and Computer Science Division, Argonne National Laboratory
Computation Institute, University of Chicago and Argonne National Laboratory
hereld@mcs.anl.gov

Abstract

We describe a new method for accurately measuring the relationship between pixel coordinates of neighboring projectors in a tiled projection display. The method can be used to calibrate the alignment and distortions in the tiled display and to touch up existing calibrations with local corrections. Additionally, by simultaneously measuring the characteristic spacing of textural features on the screen surface, we show how an absolute measure of the screen coordinates might be derived. The techniques have application to information encoding in images, in situ measurement of projector distortions without interrupting display of image content, and commodity tile-aware smart projectors. We discuss the limits to precision of these measurements.

1. Introduction

Many computer vision application areas need to encode information in a visual field for capture and interpretation by camera systems: 3-D scanning, automated inspection, inventory tracking, pattern recognition, data transmission, and various closed-loop measurement systems. Solutions to these problems have included several kinds of so-called structured light, 2-D bar code schemes, and dynamically changing image content.

After a brief discussion of related work (Sec. 2), we describe (Sec. 3) crinkle, a scheme for encoding detailed data into an image, and discuss how it can be used to measure in detail the mapping between pixels of one projector and its neighbor in the vicinity of the measurements. We then describe (Sec. 4) a method for extracting a physical scale from textural cues on a display surface. We discuss the individual merits of crinkle and texture measurements. We then describe (Sec. 4) a technique for combining these measurements to extract a full calibration of the alignment of projectors in a tiled display and to extract transformations that can

be used to correctly place image data into the collection of frame buffers driving the array of projectors for optimal resolution.

2. Related work

Measuring and correcting misalignment and distortion in tiled projection displays by using cameras has been an active research topic for many years [6,8]. Several measurement techniques have been explored, as have several methods for applying corrections derived from the measurements.

Cameras have been used to measure projected marks spread over the tile either as grids of dots or sequentially as Gaussian features [12,2]. These methods enable the measurement not only of projector orientation and position but also of distortions inherent in projector lenses. In our previous work, we extended such methods to careful measurement of dense grids of dots, across the entire field of the display [7].

Others have added absolute marks either projected by lasers or permanently incorporated into the field of view of the measuring camera in order to ground relative measurements in camera coordinates to the geometry of the screen.

Still others have calibrated their cameras using test patterns on cards or in situ, to measure and remove the distortions of the camera optics, leaving purely projective coordinate effects. So-called calibrated cameras can then be used to directly measure positions of projected marks in a rectilinear screen-based coordinate system.

Many have focused on purely projective transformations (homographies) when correcting alignment issues, and therefore on measurement patterns that are relatively simple, typically sparse dot grids or simple rectangles. These methods enable the correction of orientation and positional errors of the projector with respect to the

screen and neighboring projectors. They do not account for distortion in the projector optics.

At least one method that has been studied relies on *local* measurements of the relative positions of pixels in one tile with respect to pixels in another [5]. In this approach, measurements are made at many places in the overlap regions of all the tiles. These are then used to generate a transformation from tile coordinates to a common coordinate system based on simulated annealing. In an extension of this technique [4], the tile-to-tile transformations were organized into a hierarchy of homographies, which sped the solution, led to better accuracy, and dodged convergence problems.

And others have taken a very casual approach to alignment, and even to geometry of projection surface. Alignment is taken as is, and all surfaces are fair game. Corrections are computed from a single point of view [11,9,12].

Correction methods, when armed with detailed information about the mapping between projector coordinates and screen coordinates, can be used to warp image content into the frame buffer in order to correct for a plethora of defects [10,12,2,7]. This strategy can be done in real time by using current generation commodity graphics cards in PCs, with little impact on frame rate [1].

A common practice is to use measurements to correct only the projective transformation of image content into the frame buffer of each projector driver [3,5]. This accounts for all orientation, position, and scale parameters and can therefore correct zoom, keystone, rotation, and registration errors.

In broad summary, these measurement techniques fall into two categories: *global* or *local*. In global techniques, the camera is used to capture a relatively large portion of the tiled display, or at least an entire tile of the display. Local techniques measure only small portions of the display at a time, typically tens or hundreds of pixels on a side. The method described in this paper is *local*. The measurements are then applied correctively by one of a few methods: projective transformation, optical model inversion, or general image warp on triangular mesh. Any of these correction methods can use the measurements described in this paper.

2. Crinkle

In this section we describe our method for encoding pixel positions in projected images, and we discuss how cameras can use these to make local measurements of pixel alignment between neighboring projector tiles. For this application, the chief advantage of our approach is that nothing needs to be known a priori about the alignment and registration of the tiles: any pixel in the image can be accurately identified by using only an image of its immediate surrounds. There is no need to place marks in one tile image at points predetermined to be proximate to marks placed in a neighboring or overlapping tile image. Neither is there a need to use closed-loop algorithms that actively move marks around in the field until alignment is achieved. Our method is *static* and *local*, leading to simple extraction methods.

The encoding scheme must be easily detected and recognized by software, provide subpixel accurate locations, fit well within a modest camera pixel format, be readable at any orientation or scale, and be separable from patterns from neighboring overlapping tiles. It is also advantageous for the brightness of the code to be independent of payload, for the coding unit to fit in a modes field of view, and for the scheme to support simultaneous measurements of overlapping codes.




Figure 1. A balanced 2x2 pixel representation of bit values and an example of a two-bit orientation key.

Crinkle is a scheme that we have devised that meets these needs well. It is built on encoded messages (for pixel position, tile number, etc.) using symmetric and antisymmetric pattern building blocks to represent 1s and 0s and registration cues to help recognition and unambiguous decoding (Fig. 1). The bit representations have the same number of *on* and *off* pixels, which aids in robust interpretation of the pattern in several ways. No streak of more than two *on* or *off* pixels can occur in the coding region. This scheme gives a large number of centroids from which to extract precise pixel positions and to infer by interpolation the positions of *off* pixels. It

can therefore generate a very accurate local map of pixel positions in camera coordinates.

To establish orientation, we considered “two-bit” patterns, 4 x 2 pixels, that would expose the orientation unambiguously, and could not be created by accidental sequences of the “1” and “0” patterns or by scale change. Figure 1 shows an orientation key that we use.

We combine these building blocks into a 2-D layout (Fig. 2). This template describes a crinkle with 54 bits of data plus 2 bits for an orientation key, 8 x 6 bits, or 16 x 12 pixels in the image. The crinkle is read left to right, top to bottom. The fields are described in the legend. The crinkle encodes the pixel coordinates of its upper left corner. Because the encoding is dense in *on* and *off* pixels both, approximately half *on*, every pixel position in the vicinity of the decoded crinkle can be deduced simply. These crinkles can be packed into the image in any convenient way, typically as a regular array with or without a gutter between adjacent crinkles. We typically guard each crinkle with a one-pixel-wide gutter of *off* pixels. This helps us to see the patterns and helps in automatic finding of the crinkle in the field.

O	T									
X										
Y										
U										

<i>Field</i>	<i>Value</i>	<i>Bits</i>
O	Orientation	2
T	Tile Number	16
X	X Coordinate	16
Y	Y Coordinate	16
U	Undefined	8

Figure 2. An example of a coding layout for a single crinkle.

Images of overlapping tiles projecting crinkle-encoded pixel coordinates in different colors can be decoded to generate coordinate transformations between these projectors. Figure 3 shows an example of such an image. The lighter crinkle pattern on the left is projected in red by one projector and actually covers the entire field of this image. The brighter (and fuzzier) crinkle on the right is at the edge of the tile projected in blue by a second projector.

In our example, we have separated the blue crinkles from on projector by simply using the blue color channel of the image. This method doesn't work generally because of differences in projector and camera filters. We expect that good results will be obtainable using a principal component analysis, but we are also

considering cluster analysis methods. Recognizing, measuring, and interpreting the crinkle is detailed but straightforward. We have already implemented a similar extraction in our work on DottyToto [7]; for details see the description therein.

In Figure 4 we sketch the steps used to extract these transformations. The signal-to-noise ratio is high because of the large number of pixels in the crinkle, the resulting large number of independent peak detections, and the statistical advantage of combining those measurements to extract the relatively small number of parameters describing the pixel positions in camera coordinates.




Figure 3. Example of crinkle separation by color (shown here in black and white). In the top image a very small portion near the edge of one tile is shown (approx. 40 x 30 pixels). The image on the bottom is the blue channel of the color image and shows easy-to-read crinkle patterns.

Our first experiments have yielded good results. Starting with the top image of Figure 3 we have correctly

measured the positions of the 99 *on* pixels in the red (left) crinkle to an accuracy of less than one camera pixel. For this measurement one camera pixel corresponds to 1/8 projector pixel. If the errors are systematically skewed in a single direction, a worst case, then we can place the entire crinkle in the camera field to that same 1/8 projector pixel. Combining errors quadratically we set the upper bound on our accuracy in relating positions of pixels from two crinkles (two corresponding projectors) at $\sqrt{2}/8$ in x and in y , less than 0.2 projector pixels. Note that if the pixel measurement errors are uncorrelated, then our accuracy could be better by a factor of as much as $\sqrt{99}$, roughly 10 times better, far below other errors.

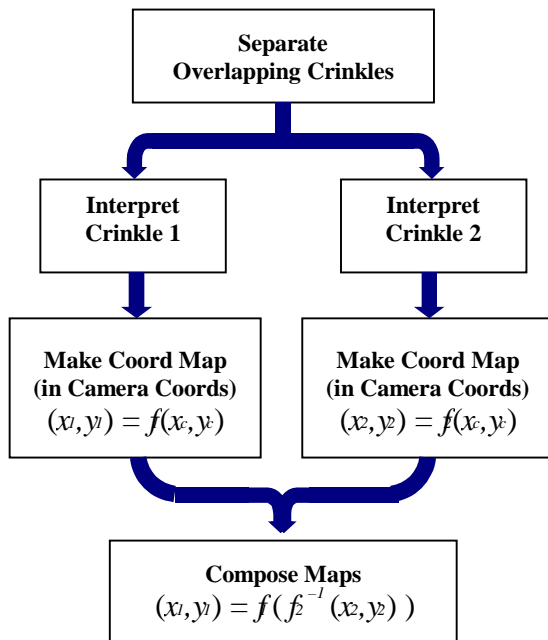


Figure 4. Extracting tile-to-tile coordinate mappings from a single image of overlapping crinkle patterns from two projectors.

3. Screen texture

Many, if not most, screen surfaces have texture that can be seen with a digital camera. If the characteristic length scales of this texture, whatever its source, are uniform across the display surface and can be measured with enough precision from images of it, we can exploit such measurements in our calibration of tiled displays.

In the case of many manufactured rear-projection screens, the back surface of the screen is covered with a film comprising very small beads. These are part of the technology for presenting a low albedo, and therefore black, surface to the viewer, a characteristic that significantly increases the dynamic range of the projection system. These beads are typically (for materials that we are aware of) several tens of microns in diameter (say, 70) and are selected to be very close to the same size.

Figure 5 shows a close up image of such a screen. The image is taken from the viewing side of the screen, from a distance of only a few centimeters. The field of view is approximately 5 pixels wide by 4 pixels tall (4 mm x 3 mm). The bright dots within each large pixel blob are the glass beads. The blobs comprising the pattern (the “X” shown by X-windows at startup) are individual pixels, a little under a millimeter “square.”




Figure 5. Detail of an image showing the texture from glass beads on the Jenmar Blackscreen material.

One way to exploit these two measurements together, crinkle and texture, has the potential to yield a full calibration of the alignment of a tiled display. In brief the steps we propose are as follows:

- Use crinkle to measure the pixel scale of the projector at several points.
- Use simultaneous texture measurements to attach a physical scale to each pixel scale.
- For each projector, compute the homography that best fits these data.
- Use techniques described in [4] to combine the measurements and arrive at a global coordinate system and first-pass transformation from each

projector pixel coordinates into global coordinates.

- Reuse the crinkle images to extract detailed tile-to-tile transformations in the overlapping or neighboring projector tiles.
- Use these detailed local maps to correct the first-pass maps derived earlier.

Finally, we derive an estimate of the physical scale of projector pixels on the screen and its uncertainty, in projector pixels per bead separation, in a few steps as follows. These calculations will help us design the measurement parameters required to realize the scheme just outlined. For the derivation we estimate errors as coming from counting statistics in both bead and crinkle cases. We believe that this measure is an upper bound to the real errors.

First, the camera image can be used to find the relationship between camera pixels and bead separation. If N beads are counted in an area, A , on the camera image in (camera pixels)², then the bead scale in camera pixels is given by:

$$S_b = \sqrt{\frac{A}{N}} \pm \frac{\sqrt{A}}{2N}. \quad (1)$$

On the other hand, the camera picture can also be used to measure the relationship between camera pixels and projector pixels. An m by n crinkle pattern contains $mn/2$ measurable projector pixels, each with an uncertainty of $\delta x = \delta y$ in its measured position. The projector pixel scale in camera pixels is given by

$$S_p = \sqrt{\frac{A}{mn}} \pm \frac{2}{\sqrt{m^2 + n^2}} \sqrt{\frac{2}{mn}} \delta x. \quad (2)$$

From this we can estimate the projector pixel scale factor, $S = S_b/S_p$:

$$S = \sqrt{\frac{mn}{N}} \pm \sqrt{\frac{mn}{N}} \sqrt{\left(\frac{1}{4N}\right) + \frac{8}{(m^2 + n^2)} \frac{\delta x^2}{A}}. \quad (3)$$

For example, if we assume the 16 x 12 projector pixel crinkles described above, an A about 300 x 200, N about 120 x 80, and δx about 0.5 camera pixels, we get $S = 0.05 \pm 0.005$ projector pixels per glass bead.

4. Discussion, future work, and conclusions

We have described a novel method for locally encoding projector pixel coordinates, in the form of 2-D *crinkles*, in an image so that they can be unambiguously deciphered with a camera-based measurement. The method is compatible with simultaneous decoding of data in overlapping projections and can be used to precisely measure the relation between pixel coordinates of these projectors. In this application, crinkles can be used to quickly measure and touch up minor misalignments.

We have also demonstrated the feasibility of a method for measuring the local scale of a projection on a bead-based rear projection screen in terms of projector pixels per screen scale length. If local scale length is constant across the screen, then the measurement can be coupled with the crinkle measurements to map absolute pixel positions for each projector of a tiled display in terms of the physical position on the screen. Failing that, such measurements made at several points on a tile can be used to compute a rough correction to keystoning and projector-to-projector zoom mismatch.

In particular with regard to the texture measurements described, the practicality of the methods discussed will hinge on measurement uncertainty. To aid later calculations and experiment designs, we have estimated the errors of these measurement methods.

The methods described in this paper can be extended to apply to other materials, circumstances, and applications.

- Sources of uniform screen texture (naturally occurring or engineered for the purpose) other than the glass beads could form the base.
- Projectors equipped with through the lens imaging might be designed to automatically read crinkles in neighboring projectors to create a commodity-based tiling capability.
- Crinkles might be embossed onto ordinary image content to enable live alignment measurement and correction without disrupting production use of the display system.

To further the work on this project, we plan the following:

- Integrate the separate steps into a single software bundle, and package it for routine use.
- Test methods for quickly and robustly extracting a global coordinate system for applications where these techniques are used for full calibration.

- Implement a demonstration of the end-to-end correction using a hand-held camera, perhaps a networked iPaq.
- Explore the limits to measurement precision of embossing applications of crinkle, that is, where crinkle data is impressed over application image data.

We have described two measurements that might be applied fruitfully to the problem of accurate alignment of image content in tiled display systems. The anticipated errors of measurements based on these techniques are promising.

Acknowledgments

We acknowledge helpful discussions of the ideas leading to this paper with Ivan Judson, Rick Stevens, Mike Papka, Ed Frank, the Futures Lab en masse, and Aditi Majumder.

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38.

References

- [1] J. Binns, G. Gill, M. Hereld, D. Jones, I. Judson, T. Leggett, A. Majumder, M. McCrory, M. E. Papka, and R. Stevens, Applying geometry and color correction to tiled display walls, poster presented for *IEEE Visualization 2002*, Boston, MA, October 2002.
- [2] M. S. Brown and W. Brent Seales, 2002, A practical and flexible tiled display system, *IEEE Proceedings of the 10th Pacific Conference on Computer Graphics and Applications (PG'02)*, pp. 194 - 203.
- [3] C. J. Chen and M. Johnson. Fundamentals of scalable high resolution seamlessly tiled projection systems, *Proceedings of SPIE Projection Displays VII*, 4294:67-74, 2001.
- [4] H. Chen, R. Sukthankar, G. Wallace, and K. Li. Scalable alignment of large-format multi-projector displays using camera homography trees, *IEEE Visualization 2002 (Vis 2002)*, Boston, MA, pp. 339 - 346, 2002.
- [5] Y. Chen, H. Chen, D. W. Clark, Z. Liu, G. Wallace, and K. Li, Automatic alignment of high-resolution multi-projector displays using an un-calibrated camera, *IEEE Visualization 2000*, Salt Lake City, UT, October 8 - 13, 2000.
- [6] M. Hereld, I. Judson, and R. Stevens. Introduction to building projection-based tiled displays. *IEEE Computer Graphics and Applications*, 20(4):22-28, 2000.
- [7] M. Hereld, I. Judson, and R. Stevens. DottyToto: A measurement engine for aligning multi-projector display systems. In *Projection Displays IX, Proceedings of SPIE Vol. 5002*, ed. Ming Hsien Wu, Santa Clara, CA 2003, pp. 73 - 86. (Also Argonne National Laboratory preprint ANL/MCS-P958-052, 2002.)
- [8] K. Li, H. Chen, Y. Chen, D. W. Clark, P. Cook, S. Damianakis, G. Essl, A. Finkelstein, T. Funkhouser, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J. P. Singh, G. Tzanetakis, and J. Zheng. Early experiences and challenges in building and using a scalable display wall system, *IEEE Computer Graphics and Applications*, 20(4):671-680, 2000.
- [9] R. Raskar, Immersive planar displays using roughly aligned projectors, *Proceedings of IEEE Virtual Reality 2000*, pp. 109 - 106.
- [10] R. Raskar, M. S. Brown, R. Yang, W. Chen, G. Welch, H. Towles, B. Seales, and H. Fuchs. Multi-projector displays using camera-based registration, *IEEE Visualization '99*, pp. 161 - 168, San Francisco, October 1999.
- [11] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays, *Computer Graphics (Proc. SIGGRAPH)*, pp. 179 - 188, 1998.
- [12] R. Yang, D. Gotz, H. Towles, and M. S. Brown, Pixelflex: A dynamically configurable display system, *IEEE Visualization '01*, 2001.

The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory ("Argonne") under Contract W-31-109-ENG-38 with the U.S. Dept. of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.