

Adaptive Instant Displays: Continuously Calibrated Projections Using Per-Pixel Light Control

Daniel Cotting
ETH Zurich
dcotting@inf.ethz.ch

Remo Ziegler
ETH Zurich
zieglerr@inf.ethz.ch

Markus Gross
ETH Zurich
grossm@inf.ethz.ch

Henry Fuchs
University of North Carolina at Chapel Hill
fuchs@cs.unc.edu

Abstract

We present a framework for achieving user-defined on-demand displays in setups containing bricks of movable cameras and DLP-projectors. A dynamic calibration procedure is introduced, which handles cameras and projectors in a unified way and allows continuous flexible setup changes, while seamless projection alignment and blending is performed simultaneously. For interaction, an intuitive laser pointer based technique is developed, which can be combined with real-time 3D information acquired from the scene. All these tasks can be performed concurrently with the display of a user-chosen application in a non-interfering way. This is achieved by using an imperceptible structured light approach enabling pixel-based surface light control suited for a wide range of computer graphics and vision algorithms. To ensure scalability of light control in the same working space, multiple projectors are multiplexed.

Categories and Subject Descriptors: I.3.2 [Graphics Systems], I.3.1 [Hardware Architecture], I.4.1 [Digitization and Image Capture], I.4.8 [Scene Analysis], I.3.6 [Methodology and Techniques]

1. Introduction

Nowadays, computer and home entertainment screens generally are fixed in size and position. Due to these limitations, users are required to adapt to given setups, instead of the display systems continuously accommodating the users' needs. Designing display systems, which can be activated on demand by the users anytime and anywhere, introduces many interesting challenges. Techniques to project screens at arbitrary locations need to be developed, and the systems need to recalibrate and recover the display geometry continuously in case an undistorted view is desired [RCWS98]. Furthermore, novel forms of multimodal interaction beyond mouse and keyboard need to be devised in order to use the new display opportunities to full capacity. Although steady progress in research has been made to counter the restrictions of traditional computer displays, most attempts have not focused on flexibility and elegance of the systems.

Relying on a novel technique of per-pixel light control, this paper presents a complete, intuitive framework for realizing instant displays with the help of large tiled projections. Simple bricks containing cameras and projectors are combined in a flexible, scalable way to achieve increasingly immersive, interactive environments. Dynamic changes are continuously accounted for by a recalibration procedure based on imperceptible per-pixel light control.

2. Related work

To build our instant displays, we rely on projection technology, since currently no other technology provides a way nearly as competitive and effective to build large, high-resolution screens. Therefore, our work is tightly related to large

tiled displays, where multiple projections are aligned and blended to create wide-area screens [Sur99, Che01, RGM*03, GSSB03]. Once these systems are set up, they all focus on achieving a static projection, without considering the projection of dynamic on-demand screens in flexible user-defined areas. Even though the PixelFlex system [RGM*03] is able to automatically rearrange its projection layout, it was not designed to provide on-demand screens on arbitrary surfaces. The IBM Research steerable projections [Pin01] and subsequent systems inspired by them [BSS04, EHH04] are notable exceptions and allow dynamic, spatially reconfigurable displays. However, they are restricted to single projectors and scalability is not considered.

2.1. Projection alignment and blending

For appealing projection-based displays, a large variety of problems have to be addressed. An extensive bibliography of projector-related publications is provided by Raskar [Ras05]. Creating high-quality images does not only require geometric alignment, but also smooth blending of multiple units and thorough color adjustment [MS05]. Solutions for projector alignment have mainly been devised for limited, application-specific cases, restricting the display geometry to planar, multi-planar or curved surfaces [CSWL02, HJS02, RP04, AFSR04]. More general results considering projections to arbitrary surfaces have been created by Raskar [Ras03, RBY*99]. While most methods treat the alignment task as a one-time initialization step, Yang and Welch [YW01] have devised a procedure continuously updating parameters during runtime. Similarly, our approach allows dynamic changes of arbitrary display surfaces, but does not rely on time-consuming correlation calculations.

2.2. Calibration

Since projectors can be seen as duals of cameras, projection alignment is naturally related to camera calibration. In wide-area environments, individual cameras often observe only a small fraction of the entire scene, making determination of reliable global relationships very difficult. Approaches based on self-calibration [MF92] resolve this problem by observing a number of features in a set of successive images [LD00, CDS00, IAM04].

Adaptivity of calibration to setup changes has been presented by Zhang and Schenk [ZS97] in the context of robotic systems. Unfortunately, the approach does not consider radial distortion, which is often affecting wide field-of-view cameras. Shen and Menq [SM02] propose a dedicated projector calibration procedure and support projector location changes through rigid-body transformations, while camera positions are restricted to static positions. Similarly, Chen and Li [CL03] present an approach for self-recalibration confined to a single projector-camera pair.

In order to eliminate the drawbacks of the aforementioned methods, we rely on a hybrid approach. First, we extend a self-calibration procedure by Barreto et al. [BD04] to achieve initial calibration of a system consisting of multiple projectors and cameras. At the same time, we preserve the option of simultaneous handling of radial distortion, intrinsic and extrinsic calibration in a reliable way. Second, dynamic changes are continuously accounted for with an efficient rigid-body transformation approach.

2.3. Imperceptible structured light

Our calibration relies on per-pixel surface light control achieved through imperceptible embedding of patterns into projections. In the Office of the Future project, a proof of concept for the embedding of invisible structured light into DLP projections was given [RWC*98]. As major restrictions, the resulting images were greyscale only, and the realization required significant modifications of the projectors.

As opposed to this earlier approach, our research focuses on a novel color-driven embedding procedure having significant advantages. We project structured light imperceptibly with conventional off-the-shelf DLP projectors, while also simultaneously displaying real-time color imagery. Whereas previous results were confined to initial proof-of-concept systems [CNGF04], our adaptive instant displays represent a significantly enhanced setup making extensive use of the color-driven per-pixel light control approach.

2.4. Distributed displays

When handling rendering on a large number of projectors, ways of creating consistent, synchronized image streams have to be found. Depending on the application requirements, different approaches have been pursued. Windows virtual device drivers and OpenGL implementations have been developed for the Scalable Display wall [LCC*00]. WireGL/Chromium software [HHN*02] provides a cross-platform OpenGL implementation for distributed rendering. Similarly, PixelFlex [RGM*03] contains an OpenGL wrap-

per and displays arbitrary X-Window applications using a modified version of VNC [RSFWH98], whereas the ICWall system relies on a parallel implementation of the retained mode Aura API [vdSRG*02]. Unfortunately, none of these systems focuses on easily achieving distributed display of standard Windows applications, a gap we fill with a newly developed approach.

2.5. Human-computer interaction

Projection screens and instant displays require new forms of interaction beyond the usual keyboard and mouse. Discussion of the wide field of human-computer interaction is beyond the scope of this paper. Refer to the papers cited previously for an insight into how various research groups tackle this issue. Solutions include magnetic trackers, 2D and 3D gesture recognition, and speech interfaces. Our laser pointer based user interface is inspired by interaction techniques presented by Sukthankar et al. [RS00, SSM01] and Davis et al. [DC02]. More natural and versatile interaction is achieved through the use of pointer gesture recognition.

3. System overview

The setup used for creating instant displays consists of movable bricks each containing several cameras and a projector connected to a PC (see Figure 1). The networked bricks can be individually oriented to cover the desired working space, usually consisting of several projection walls and desks. A microcontroller unit (MCU) is used as a synchronization source for the cameras and the graphics boards driving the projectors. In our implementation, we rely on three bricks having two color cameras and one greyscale camera each, resulting in a configuration, which allows the bricks to be used for a wide variety of other tasks like stereo reconstruction as well. Figure 2 shows a single brick implementation with its components.

4. Per-pixel surface light control

Our instant display system relies on a technique allowing to imperceptibly control lighting conditions of the projection surface during camera exposure. This is done at the scale of individual projector pixels, thus allowing structured light approaches not noticeable by the user.




Figure 1: Configuration of the instant display setup.



Figure 2: Instant display brick with projector and cameras.

4.1. Imperceptible structured light

Recently, a novel approach to imperceptible structured light has been introduced [CNGF04]. By implementing a proof-of-concept system, the authors have demonstrated its applicability to a wide range of computer graphics and vision algorithms. Likewise, many components of our adaptive instant displays rely on this technique, which will be summarized shortly in the following paragraphs.

In DLP projectors, each displayed pixel is generated by a tiny micro-mirror, tilting towards the screen to project light and orienting towards an absorber to keep the pixel dark. Gradations of intensity values are created by flipping the mirror in a fast modulation sequence (see Figure 3), while a synchronized filter wheel rotates in the optical path to generate colors. By carefully selecting the projected intensities, one can control whether or not the mirrors for the corresponding pixels project light onto the scene during a pre-defined exposure time slot of a synchronized camera.




Figure 3: DLP projector light modulation.

The core idea of the imperceptible pattern embedding is a dithering of the projected images using color sets appearing either bright or dark for the camera, depending on the chosen pattern (see Figure 4). For more details refer to the original publication [CNGF04].

The following sections are devoted to the newly developed embedding enhancements. In order to demonstrate the versatility of the approach, we extend the existing method to low-budget DLP projectors (NEC LT240K) having interdependent color channels and providing no DVI input, but only an analog VGA interface. Additionally, a GPU-based version of the encoding algorithm is sketched, whose integration considerably speeds up projection frame rate.




Figure 4: Projected image and pattern captured by camera.

4.2. Color classification

Projector models having interdependent color channels require an extended version of the camera-based color analysis method described in the original approach [CNGF04]. Since the mirror flips for one primary color (red, green or blue) are not independent of the intensities of the other two primaries (cf. Figure 5), a subspace of the entire color space has to be sampled and analyzed. Its classification into colors appearing bright or dark is then used for the imperceptible embedding of patterns.




Figure 5: Interdependent color channels. Mirror-flips for red are different for varying intensities A, B and C of the remaining primaries (green and blue).

The required equipment for color classification consists of a projector and a camera, which are both synchronized using an external sync generator and a genlock capable graphics board. The procedure is entirely automatic and only needs to be computed once for each new and unknown projector.

Color classification can be split into two parts and has to be repeated separately for each color channel. During the first part, the optimal embedding time slot providing a variety of projection colors has to be found. For this purpose, an image with a gradient from lowest to highest intensity of the channel considered is projected. While cycling through time, we choose the slot resulting in a maximum number of alternating dark and bright stripes perceptible in the camera image. Some projector models dither certain intensities [Uli03], making part of the color values unreliable. These intensities are statistically detected and ignored, resulting in the localization of the time slot shown in Figure 3.

In the second part of color classification, further analysis is required. For all intensities of the primary channel, combinations with a chosen color subset of the two remaining channels have to be classified. In case of projectors showing




Figure 6: Classification of color space. a) Projection and capture of patches. b) Projection and capture of stripes.

no variation in color modulation across the micro-mirrors, grids with colored patches are projected, shown in Figure 6 a). Likewise, projectors showing a spatial dependency along one of the axes can be handled by generating colored stripes, illustrated in Figure 6 b). Color classification is then performed on the captured images by thresholding, while disregarding noisy samples. Colors reliably recognized as bright or dark during camera exposure are then stored in corresponding tables, visualized in Figure 7.




Figure 7: Resulting colors classified as appearing a) bright or b) dark. Computation is performed for an embedding in the red channel. The green and blue primaries are sampled with a resolution of 32.

4.3. Scalability

Multiple projectors are required to cover a large area and achieve an increasingly immersive environment. In combination with the presented embedding approach, light control becomes an issue: Each device projecting its own structured light pattern onto the scene potentially interferes with the illumination from other projectors. For overlapping frusta, we can use different multiplexing approaches.

Three projectors. Scaling up to three devices with a common projection overlap is straightforward: one embedding time slot is available per primary color. During the embedding period of a certain device, we encode entirely dark patterns in the remaining projectors, whose input signals and color wheels are synchronized. Thus, as opposed to the single projector approach, the embedding slots in all the primary colors have to be considered, and the color tables result from intersections of the previously considered sets with the colors appearing dark in the two other primaries.

Four to nine projectors. The spectral response of the primary colors of a DLP projector is shown in Figure 8. By applying band-pass color filters to the cameras, patterns can be simultaneously embedded in all three primaries without noticeably interfering. Due to the filters, the cameras will detect very little light from projectors emitting in different




Figure 8: Measured spectral response of NEC LT240K projector. Colored areas indicate ideal bands of the color filters.

primary colors, even if they are running concurrently in the same overlap area. Thus, we can divide up to nine projectors into three classes, offsetting the synchronization of each class by the duration of one sector in the color wheel (see Figure 9).

More than nine projectors. We present two possibilities for using ten or more projectors. In a first approach, given a particular physical arrangement, each projector has to be assigned one of nine distinct time/color identifiers in such a way, that frusta of projectors with the same identifiers are disjoint. The same must be guaranteed for the cameras. We thus eliminate ambiguity of projector source within any camera image. A detailed investigation of the optimization problem regarding positions, orientations, and working volume is subject to future work. Another approach not imposing any restrictions on the projector placements is based on a token passing scheme. The number of projectors in Figure 9 can easily be expanded by forcing additional projectors to embed completely dark patterns in all time slots. Tokens determining which devices can display arbitrary patterns and which ones need to remain dark are then passed among the projectors. In this case, the price for increased flexibility in projector placement is a decrease in pattern throughput per device. Since this approach is the most general one, we will rely on it in all the following sections.




Figure 9: Scalability using synchronization offsets and camera color filters. The individual camera exposure periods are highlighted. Note that the rate of patterns is limited by the input frequency, i.e. 60Hz, although the illustrated projector internally runs at 120Hz. The clear period (gray) is not used.

4.4. Embedding

Once color classification for a given scalability approach is completed, an embedding and dithering procedure has to be employed. Each possible color value needs to be mapped to the nearest colors appearing either bright or dark in the selected time slot. For projectors with interdependent color channels, straightforward table-based implementations require one mapping entry for each potential display color,

leading to huge memory requirements, which can be lowered significantly by subsampling the input values, analogously to the procedure in Section 4.2. Because of cache misses, software dithering on the main CPU still results in a performance penalty, which can be eliminated by resorting to GPU-based embedding.

Initialization is performed by creating the quantized color mapping tables on the CPU, using kd-trees to speed up nearest-neighbor queries in CIE Lab color space. For each table, we also provide an additional complementary one, containing colors best compensating the color error introduced by the mapping. Subsequently, the data is loaded into the GPU as 3-dimensional volume textures. To complete initialization, a dither matrix and an alpha mask are transferred (for projection blending, cf. Section 6.2).

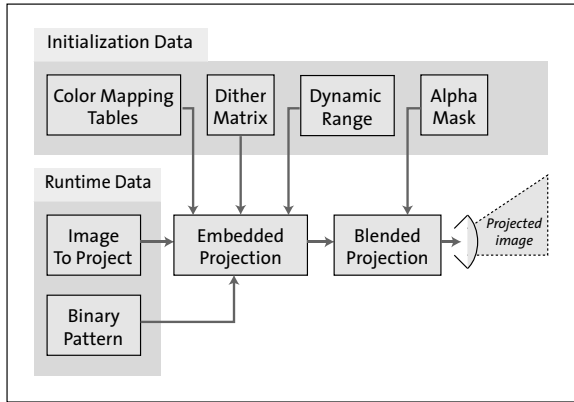


Figure 10: Runtime pipeline in the GPU fragment shader.

During runtime (see Figure 10), both the image to be projected and the binary pattern are loaded into the graphics board as rectangular textures. For each pixel, a fragment shader program adapts the resulting dynamic range [CNGF04], and depending on the corresponding pattern value, two colors are extracted from the bright or dark volume textures. Both colors have previously been created to mutually compensate for the color error induced by the mapping. We therefore mix these two colors by switching between them, as defined by the dither matrix. Finally, if several projections overlap, a blending mask is applied to the projected image. Note that since blending modifies the projected colors, this last step results in removal of the imperceptible pattern in overlap regions.

5. Calibration

To achieve seamless alignment of instant display projections, at any time all devices must be calibrated intrinsically and extrinsically with relation to each other. To achieve this goal, we have chosen an approach, which can be split into two distinct parts: An initial calibration of the setup and a continuous dynamic recalibration of modified components.

5.1. Initial calibration

During initial calibration, which is run only once at the beginning of operation, projection properties of all compo-

nents are determined in a static setting. While traditional calibration procedures [Tsa89, Bou05] provide fast and reliable parameter estimation, they are quite tedious for wide-area setups and are not designed for simultaneous handling of projectors, which are better included in self-calibration approaches. In an effort to combine the advantages of both traditional calibration and self-calibration, we extend an approach presented by Barreto and Daniilidis [BD04] to setups containing multiple projectors and cameras. The unified procedure, which propagates the Euclidean structure of traditionally calibrated cameras by self-calibration, is described in the following sections.

Precalibration. Propagation of Euclidean structure requires precalibrated components as a basis. Therefore, two cameras belonging to a single brick are calibrated using conventional methods by capturing common views of a reference pattern [Tsa89, Bou05]. At the same time, this practical procedure registers calibration to a predefined physical environment, therefore avoiding the stratification process commonly used in self-calibration approaches and leading to more accurate results in general.

Imperceptible point correspondences. In the subsequent computation steps, similarly to many self-calibration approaches, the procedure by Barreto and Daniilidis [BD04] is based on a set of point correspondences between views. In contrast to the original method, where the correspondences are generated by moving an LED in hundreds of unknown positions in front of the cameras, our approach relies on imperceptible projection of structured light. Binary patterns introduced by Vuylsteke and Oosterlinck [VO90] are embedded in a fast sequence, both to encode column and row indices of a discrete grid of points (see Figure 11). This approach not only allows us to create reliable point correspondences, but also permits extraction of the geometry as soon as the system is calibrated. As opposed to the LED approach, our system has many advantages: It is not restricted to dark environments, includes the projectors into a unified calibration approach, and can run continuously during system operation in a non-interfering way.




Figure 11: Patterns for encoding column and row indices.

Propagation of Euclidean structure. The point correspondences generated imperceptibly are used to propagate calibration parameters from already calibrated components to the remaining cameras and projectors. In order to achieve a

complete, consistent calibration, all devices must be connected to the precalibrated cameras over a chain of point correspondences (cf. Figure 12). Projections do not necessarily have to overlap, as long as certain cameras have a common view of the corresponding patterns. For that purpose, cameras need to alternately switch to the various multiplexed projectors.




Figure 12: Propagation using point correspondences.

Computation of propagation is performed by first calculating the depths of points relative to calibrated components, and then determining the projection matrices of uncalibrated components by taking the depths and the imaged point coordinates into account. Both steps are formulated as null subspace problems of matrices derived from known data according to Barreto and Daniilidis [BD04]. Solutions are given by the Eckart-Young-Mirsky theorem [GL83] as the right singular vectors corresponding to the smallest singular values of the matrices considered. It is remarkable, that the system does not rely on a single non-linear minimization, leading to reliable solutions without convergence issues. Using proper equilibration of the design matrices [Har98, Har95, MM99], reprojection errors approaching the lower theoretical bound in the order of the measurement errors are achieved. During calibration computation, outliers are continuously removed from the data set, leading to a procedure, which is robust against erroneous point locations. In our instant display setting, we have found error accumulation of propagation across multiple bricks to be negligible. To ensure reliable initial calibration, a temporary projection surface is moved during pattern projection to fill the entire working volume with point correspondences. If no movable surface is desired, a person walking through the projection volume can serve as a reflector for the patterns.

Radial distortion. An extended version of the propagation algorithm by Barreto and Daniilidis [BD04] also simultaneously handles radial distortion in an elegant and unified way. As an alternative, radial distortion can also either be ignored if negligible or precomputed using traditional calibration procedures.

5.2. Continuous adaptivity

After the initial calibration has been computed, components can be moved and reoriented to accommodate new user requirements. Our recalibration procedure detects such

changes during system operation and recalibrates accordingly in a very efficient manner. For simplicity, we expect intrinsic parameters to remain constant, an assumption which is often valid in practice, where components are mostly moved or reoriented.

Detection of changes. With the help of imperceptible point correspondences collected during system operation on the projection surface, the validity of calibration parameters is verified continuously. For that purpose, we rely on a pairwise comparison of components by checking whether common point correspondences are still consistent with calibration. In each camera we compute the distance between a point and the epipolar line corresponding to the respective point in the other camera. If for a component and all remaining components the distance exceeds a predefined threshold for a certain percentage of pairwise common points, we assume this component to be modified. In case an entire brick is moved at a time, the procedure described above will still recognize the brick as calibrated. Therefore, we have to apply the same distance test on a brick level to correctly detect displacements of entire bricks. As long as a single component or brick is moved at a time and multiple elements remain static, the detection is unambiguous. Note that changes in projection geometry do not affect calibration and will properly be ignored by the verification procedure.

Recalibration. Components, which have been detected as displaced or reoriented, need to be recalibrated extrinsically. For that purpose, imaged points are reconstructed in 3D, provided that they are at least visible in two components with valid calibration. Using this information a gradient descent method is applied to a 6-dimensional reprojection error minimization problem taking into account translational and rotational degrees of freedom. For all practical cases, previously valid extrinsic parameters have proven to be good starting values for the minimization process. Note that this procedure works well even in case of coplanar 3D point correspondences, which often occur in projections onto flat surfaces. Our efficient recalibration procedure can also be used as a substitute for the initial calibration as long as the intrinsic parameters have not been changed from the last execution of Euclidean calibration propagation.

6. Adaptive instant displays

Instant displays, which can be activated anywhere on demand, are our solution to the restrictions imposed by current computer and home entertainment screens, which are generally fixed in size and position. The following sections present our techniques to project and blend on-demand screens generated by multiple distributed nodes.

6.1. Distributed display

In our system, instant displays are created using multiple projectors in a distributed environment. We present a new method for distributed display, relying on a very efficient and scalable transmission based on the Microsoft RDP pro-

TOCOL. Benchmarks of various remote display mechanisms can be found in the literature [YNST02, NYN03].

Our approach is able to accommodate multiple modules, without imposing any additional network load. This is achieved by emulating multicast functionality. RDP clients are modified not to use dedicated connections, but to passively intercept traffic of a single dedicated link between an RDP server and an active RDP client connected to a common non-switching hub (see Figure 13). Interception is performed by replacing standard TCP/IP communication with an emulation layer integrating its own TCP/IP stack [Woj05] and a packet capture driver [DBRV03]. Note that no modification is required on the server, which can run arbitrary Windows applications. The fact that rendering of 3D geometry is not accelerated does not affect our applications, which mainly focus on standard office scenarios.

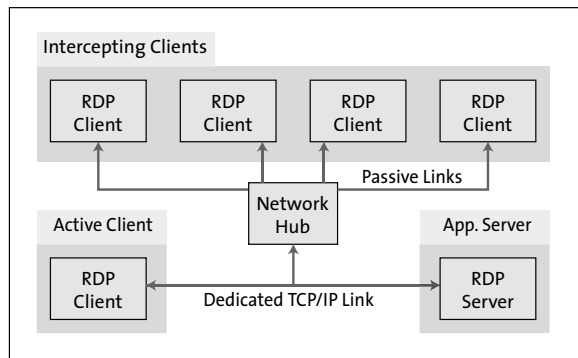


Figure 13: Distributed display architecture.

6.2. Projection blending

Images from multiple projectors need to be blended to create appealing displays at frustum intersections. We have implemented an OpenGL-based GPU-procedure resulting in pixel-level alpha masks, which can directly be applied to projections onto arbitrary surfaces. Prerequisite for successful mask computations are valid calibration parameters and knowledge of the display geometry, which can be retrieved using projected patterns (cf. Section 5.1) or the laser pointer path defining the 3D screen geometry (cf. Section 6.3).

In our procedure, the real physical setup with its display surface and projectors is simulated using projective texturing. In a first step, all projectors are loaded with an initial mask shown in Figure 14 a). The intensity $I_k(x, y)$ for pixel (x, y) of projector k is largest in the center and fades smoothly to zero at the borders of the image, therefore gradually decreasing the weighting of the projector. A simulation of the overlapping projection is then performed, leading to illumination intensity values on the projection surface, shown in Figure 14 b). Ideally, the masks should generate projections with uniform intensities. To eventually achieve this property, the masks have to be modified. For this purpose, the illuminated scene is projected back to the various projectors in order to extract the brightnesses $B_k(x, y)$ of the scene points corresponding to the initial mask intensities $I_k(x, y)$. With this information, a mask normalization step

can be performed, leading to a uniform projection. Assuming the maximum mask intensity is 1, the final mask can be computed as

$$M_k(x, y) = \left(\frac{I_k(x, y)}{B_k(x, y)} \right)^{1/\gamma}, \quad (1)$$

where γ is the projector gamma correction. An example of a computed final blending mask is depicted in Figure 14 c).

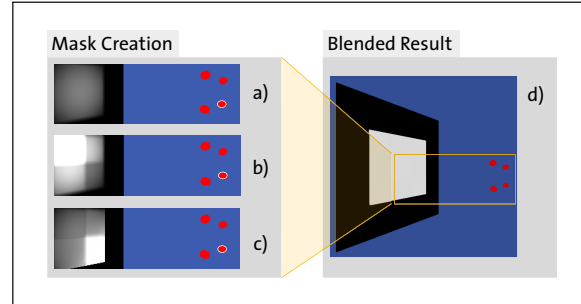


Figure 14: Simulation for mask creation of four projectors (red dots). For convenience, the approach is illustrated on a planar wall for the lower right projector. a) Initial mask. b) Overlap created with initial masks. c) Resulting normalized mask. d) Overlap created with normalized masks.

For simplicity, we confine ourselves to projectors having identical dynamic response curves, and due to diffuse reflection assumptions, the display surface at overlap should be oriented approximately at the same angle with respect to all the projectors. Additionally, we currently ignore gamut matching and variations of colors over the display surface.

6.3. Interaction

Interaction is an important factor influencing how the quality and usefulness of a system is perceived. To accommodate multiple users and to allow for encumbrance-free interaction, instant displays can no longer rely solely on traditional input devices like keyboards and mice. For this reason, we have developed an appealing interaction technique, relying on laser pointer gesture recognition.

Previous papers have considered laser pointers for interaction with projected displays [RS00, SSM01, DC02]. We extend these ideas by adding intuitive command events, which can be defined by pointer gestures. For this purpose, Kalman-filtered laser pointer paths are reconstructed in 3D and mapped to events using a previously trained neural network. In an intuitive way, the displays can be rotated, scaled and moved to new locations.

7. Results

Based on the methods presented in this paper, we have implemented an instant display system using off-the-shelf components in a standard office environment. Our setup operates with three bricks (see Section 3), each consisting of a standard PC with an NVIDIA Quadro FX3000G graphics board, a projector, a greyscale camera and two color cameras. The system is complemented by a custom-made syn-

chronization device and a network hub. The projectors and cameras are mounted on a portable aluminium rig as shown in Figure 2. We currently use NEC LT240K projectors with XGA resolution (1024×768) and Point Grey Dragonfly cameras with the same resolution.

Various projections containing imperceptible patterns and the corresponding camera captures are shown in Figure 15. Although the cameras show smearing artifacts due to short shutter times, the patterns are clearly visible.




Figure 15: Projected images and captured patterns.

With the help of imperceptible point correspondences extracted from the patterns, the initial calibration computes the arrangement of the instant display setup and the intrinsic parameters of its components. In our example, the radial distortion of the camera lenses is precomputed using traditional calibration procedures. The resulting brick layout and the reconstructed calibration points in the working volume are shown in Figure 16. For all components, root-mean-square reprojection errors below three pixels are achieved. In our current implementation, which has not yet been optimized for speed, scalability of embedding is achieved using the token passing approach and our point collection for initial calibration takes approximately three minutes. We expect at least a tenfold increase in speed in a tuned system. Once the initial calibration has been computed, an instant display can be created over multiple overlapping projections. Figure 17 a) illustrates the quality of initial calibration using imperceptible point correspondences and the capabilities of our distributed rendering of unmodified Windows




Figure 16: Computed 3D layout of the brick setup and its components. Three projectors were imaging patterns for imperceptible point correspondences.




Figure 17: Instant display using three projectors. Artifacts in the upper left area are due to surface discontinuities. a) Without blending masks. b) After applying blending masks.

applications. The same scene is displayed with blending masks and a gamma correction value of 2.2 in Figure 17 b). The continuous calibration verification and the subsequent recalibration step are illustrated in Figure 18. After detecting displaced components, in our case a projector, the prototype system automatically recalibrates using the imperceptibly generated point correspondences.




Figure 18: Adaptivity of calibration. a) Misaligned display resulting from projector displacement during system operation. b) Projection after automatic recalibration.

Various display reconfiguration possibilities of our laser pointer based interaction technique are shown in Figure 19. Since complete calibration information and knowledge of the brick layout are available, screen redefinition is not restricted to a single projection surface. Therefore, displays can also be moved to new surfaces as shown in Figure 20.

8. Conclusion

In this paper, we have presented a novel approach for achieving scalable on-demand displays in setups containing multiple cameras and DLP-projectors. Relying on self-calibration, our system is easy to install, handles all components in a unified way and allows continuous dynamic setup changes. Users cannot only define arbitrary screen locations, but also interact with the displayed applications in an intuitive way. Unmodified applications are rendered on the distributed nodes in an efficient, scalable way, while seamless imagery is achieved by an alignment and blending procedure supporting arbitrary projection geometry. Generation of point correspondences for calibration are performed concurrently with the main application of the system in a non-interfering way. This is achieved by using an imperceptible structured light approach enabling per-pixel surface light control. We strongly believe that our framework for realizing interactive instant displays is an appealing solution to



Figure 19: Interaction with instant displays projected using three bricks. Screens can be moved, scaled, rotated and changed in aspect-ratio at any time by the user with intuitive laser-pointer commands.

lift the placement and interaction restrictions imposed by current computer and home entertainment systems.

9. Future work

We intend to explore further interaction techniques in order to provide attractive and natural interfaces to the end-user. For that purpose, imperceptible structured light cannot only scan the environment, but also the user, thus permitting gesture recognition on the approximate 3D reconstruction. In the future, our seamless blending can be extended to colored surfaces by taking color balancing into account. Projector gamut matching and direction-dependent projection masks further improve image quality. Additionally, we intend to temporarily disable embedding when no dynamic adaptation is desired or required by the user. We also plan to explore new ways of scalability by replacing the token passing scheme with a carrier sense multiple access / collision detection approach.

Acknowledgments

We thank Joao P. Barreto for sharing his implementation of Euclidean structure propagation, Jim Mahaney for constructing the prototype brick, Herman Towles for the spectral response measurements of the projectors, and Anton Hafner for providing samples of color filters.



Figure 20: Instant displays can be redefined with a laser pointer, allowing a user to move projections to new surfaces.

References

- [AFSR04] Ashdown M., Flagg M., Sukthankar R., Rehg J. M.: A Flexible Projector-Camera System for Multi-Planar Displays. In Proc. of CVPR 2004 (2004), pp. 165–172.
- [BD04] Barreto J., Daniilidis K.: Wide area multiple camera calibration and estimation of radial distortion. In Proc. of the 5th Workshop on Omnidirectional Vision, Camera Networks and Non-classical cameras (2004).
- [Bou05] Bouguet J.-Y.: Camera Calibration Toolbox for Matlab, http://www.vision.caltech.edu/bouguetj/calib_doc, 2005.
- [BSS04] Butz A., Schneider M., Spassova M.: SearchLight - A Lightweight Search Function for Pervasive Environments. In Proc. of Second Intl. Conf. on Pervasive Computing (PERVASIVE 2004) (2004), Ferscha A., Mattern F., (Eds.), vol. 3001 of Lecture Notes in Computer Science, Springer, pp. 351–356.
- [CDS00] Chen X., Davis J., Slusallek P.: Wide area camera calibration using virtual calibration objects. In IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'00) (2000), pp. 2520–2527.
- [Che01] Chen Y.: Building a Scalable High-Resolution Display Wall. PhD thesis, Princeton University, 2001.
- [CL03] Chen S., Li Y. F.: Self-recalibration of a colour-encoded light system for automated three-dimensional measurements. Meas. Sci. Technol. 14 (2003), 33–40.
- [CNGF04] Cotting D., Naef M., Gross M., Fuchs H.: Embedding Imperceptible Patterns into Projected Images for Simultaneous Acquisition and Display. In Proc. of Third IEEE and ACM Intl. Symposium on Mixed and Augmented Reality (ISMAR 2004) (2004), IEEE Computer Society Press, pp. 100–109.
- [CSWL02] Chen H., Sukthankar R., Wallace G., Li K.: Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees. In Proc. of the Conf. on Visualization '02 (2002).
- [DBRV03] Degioanni L., Baldi M., Risso F., Varenni G.: WinPcap: una libreria open source per l'analisi di rete. In AICA 2003 Annual Congress (2003).
- [DC02] Davis J., Chen X.: LumiPoint: Multi-User Laser-Based Interaction on Large Tiled Displays. In Displays 23:5 (2002), Elsevier Science, pp. 205–211.
- [EHH04] Ehnes J., Hirota K., Hirose M.: Projected Augmentation - Augmented Reality using Rotatable Video Projectors. In Proc. of Third IEEE and

- ACM Intl. Symposium on Mixed and Augmented Reality (ISMAR 2004) (2004), IEEE Computer Society, pp. 26–35.
- [GL83] Golub G., Loan C. V.: *Matrix Computations*. The John Hopkins University Press, 1983.
- [GSSB03] Germans D., Spoelder H., Schaaf T., Bal H.: Realizing a Stereo Tiled Display Using Commodity Components. In Proc. of the ASCI 2003 conference (2003).
- [Har95] Hartley R. I.: In defence of the 8-point algorithm. In Proc. of the Fifth Intl. Conf. on Computer Vision (1995).
- [Har98] Hartley R. I.: Minimizing Algebraic Error. In Proc. of the Sixth Intl. Conf. on Computer Vision (1998), IEEE Computer Society, p. 469.
- [HHN*02] Humphreys G., Houston M., Ng R., Frank R., Ahern S., Kirchner P. D., Klosowski J. T.: Chromium: a stream-processing framework for interactive rendering on clusters. In Proc. of ACM SIGGRAPH 2002 (2002), pp. 693–702.
- [HJS02] Hereld M., Judson I. R., Stevens R.: Dottyoto: A measurement engine for aligning multi-projector display systems. Argonne National Laboratory preprint ANL/MCS-P958-0502, 2002.
- [IAM04] Ihrke I., Ahrenberg L., Magnor M.: External camera calibration for synchronized multi-video systems. *Journal of WSCG* 12, 1-3 (Jan. 2004), 537–544.
- [LCC*00] Li K., Chen H., Chen Y., Clark D. W., Cook P., Damianakis S., Essl G., Finkelstein A., Funkhouser T., Housel T., Klein A., Liu Z., Praun E., Samanta R., Shedd B., Singh P. J., Tzanetakis G., Zheng J.: Early Experiences and Challenges in Building and Using A Scalable Display Wall System. *IEEE Computer Graphics and Applications* 20 (2000), 671–680.
- [LD00] Lourakis M., Deriche R.: Camera self-calibration using the Kruppa equations and the SVD of the fundamental matrix: The case of varying intrinsic parameters. Research report, INRIA Sophia-Antipolis (2000).
- [MF92] Maybank S. J., Faugeras O. D.: A theory of self-calibration of a moving camera. *The Intl. Journal of Computer Vision* 8(2) (August 1992), 123–152.
- [MM99] Muhlich M., Mester R.: Subspace methods and equilibration in computer vision. Tech. Rep. XP-TR-C-21, Institute for Applied Physics, Johann Wolfgang Goethe-Universitat, Frankfurt am Main, Germany (1999).
- [MS05] Majumder A., Stevens R.: Perceptual photometric seamlessness in projection-based tiled displays. *ACM Trans. Graph.* 24, 1 (2005), 118–139.
- [NYN03] Nieh J., Yang S. J., Novik N.: Measuring thin-client performance using slow-motion benchmarking. *ACM Transactions on Computer Systems (TOCS)* archive 21 (2003), 87–115.
- [Pin01] Pinhanez C.: Using a Steerable Projector and a Camera to Transform Surfaces into Interactive Displays. In *ACM Conf. on Human Factors in Computing Systems (CHI 2001)* (March 2001), pp. 369–370.
- [Ras03] Raskar R.: iLamps: Geometrically aware and self-configuring projectors. In *SIGGRAPH 2003 Conf. Proc.* (2003), ACM SIGGRAPH Annual Conf. Series.
- [Ras05] Raskar R.: Projector-related papers, <http://www.cs.unc.edu/%7Eraskar/Projector/proj-bib.html>, 2005.
- [RBY*99] Raskar R., Brown M. S., Yang R., Chen W.-C., Welch G., Towles H., Seales W. B., Fuchs H.: Multi-Projector Displays Using Camera-Based Registration. In *IEEE Visualization* (1999), pp. 161–168.
- [RCWS98] Raskar R., Cutts M., Welch G., Sturzlinger W.: Efficient Image Generation for Multiprojector and Multisurface Displays. In *Proc. of the 9th Eurographics Workshop on Rendering* (June 1998).
- [RGM*03] Raij A., Gill G., Majumder A., Towles H., Fuchs H.: PixelFlex2: A Comprehensive, Automatic, Casually-Aligned Multi-Projector Display. In *Proc. of IEEE Intl. Workshop on Projector-Camera Systems (PROCAMS)*, held in conjunction with the 9th IEEE Intl. Conf. on Computer Vision (ICCV) (2003).
- [RP04] Raij A., Pollefeys M.: Auto-Calibration of Multi-Projector Display Walls. In *17th Intl. Conf. on Pattern Recognition (ICPR2004)* (2004).
- [RS00] R. Sukthankar R. Stockton M. M.: Self-Calibrating Camera-Assisted Presentation Interface. In *Proc. of Intl. Conf. on Automation, Control, Robotics and Computer Vision* (2000).
- [RSFWH98] Richardson T., Stafford-Fraser Q., Wood K. R., Hopper A.: Virtual Network Computing. *IEEE Internet Computing* 2(1) (1998), 33–38.
- [RWC*98] Raskar R., Welch G., Cutts M., Lake A., Stesin L., Fuchs H.: The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. *Proc. of SIGGRAPH 98* (July 1998), 179–188.
- [SM02] Shen T.-S., Menq C.-H.: Digital Projector Calibration for 3D Active Vision Systems. *Transactions of the ASME* 124 (February 2002), 126–134.
- [SSM01] Sukthankar R., Stockton R., Mullin M.: Smarter Presentations: Exploiting Homography in Camera-Projector Systems. In *Proc. of Intl. Conf. on Computer Vision* (2001).
- [Sur99] Surati R. J.: Scalable Self-Calibration Display Technology for Seamless Large-Scale Displays. PhD thesis, Massachusetts Institute of Technology, 1999.
- [Tsa89] Tsai R.: Synopsis of recent progress on camera calibration for 3D machine vision. In *The Robotics Review* (1989), MIT Press, pp. 147–159.
- [Uli03] Ulichney R.: Method for Increasing Apparent Amplitude Resolution and Correcting Luminance Nonuniformity in Projected Displays. In *Proc. of IEEE Intl. Workshop on Projector-Camera Systems (PROCAMS)*, held in conjunction with the 9th IEEE Intl. Conf. on Computer Vision (ICCV) (2003).
- [vdSRG*02] van der Schaaf T., Renambot L., Germans D., Spoelder H., Bal H.: Retained Mode Parallel Rendering for Scalable Tiled Displays. In *Proc. 7th annual Immersive Projection Technology (IPT) Symposium* (2002).
- [VO90] Vuylsteke P., Oosterlinck A.: Range Image Acquisition with a Single Binary-Encoded Light Pattern. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 2 (Feb 1990), 148–164.
- [Woj05] Wojtczuk R.: Libnids, <http://www.packetfactory.net/projects/libnids/>, 2005.
- [YNST02] Yang S. J., Nieh J., Selsky M., Tiwari N.: The Performance of Remote Display Mechanisms for Thin-Client Computing. In *Proc. of the 2002 USENIX Annual Technical Conf.* (2002).
- [YW01] Yang R., Welch G.: Automatic Projector Display Surface Estimation Using Every-Day Imagery. In *9th Intl. Conf. in Central Europe on Computer Graphics, Visualization and Computer Vision* (2001).
- [ZS97] Zhang Z., Schenk V.: Self-Maintaining Camera Calibration Over Time. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'97)* (1997).