

Laser Pointer Tracking in Projector-Augmented Architectural Environments

Daniel Kurz, Ferry Häntsch, Max Große, Alexander Schiewe, and Oliver Bimber
Bauhaus-University Weimar*




Figure 1: Basic object manipulation techniques such as translation (a) and rotation (b) are illustrated in long exposure photographs. Augmentation can be projector-based (a-c) or via video-see-through (d). These application examples show basic augmentations of building structures (a,b,d), distance measurements (c) and material-color simulations (c).

ABSTRACT

We present a system that applies a custom-built pan-tilt-zoom camera for laser-pointer tracking in arbitrary real environments. Once placed in a building environment, it carries out a fully automatic self-registration, registrations of projectors, and sampling of surface parameters, such as geometry and reflectivity. After these steps, it can be used for tracking a laser spot on the surface as well as an LED marker in 3D space, using inter-playing fisheye context and controllable detail cameras. The captured surface information can be used for masking out areas that are critical to laser-pointer tracking, and for guiding geometric and radiometric image correction techniques that enable a projector-based augmentation on arbitrary surfaces. We describe a distributed software framework that couples laser-pointer tracking for interaction, projector-based AR as well as video see-through AR for visualizations with the domain specific functionality of existing desktop tools for architectural planning, simulation and building surveying.

Index Terms: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Tracking; J.5 [Computer Applications]: Arts and Humanities—Architecture

1 INTRODUCTION AND MOTIVATION

Immersive and semi-immersive projection displays, such as CAVEs, walls, workbenches, cylinders, and domes are being used to support virtual reality applications in many professional domains. The visualization of data with these displays, however, requires dedicated rooms for setting up non-mobile screens, and allows the interaction with purely synthetic information only.

The aim of the sARc project¹ is to investigate and develop the conceptual and technological fundamentals for realizing such visu-

*e-mail: {Daniel.Kurz, Ferry.Haentsch, Max.Grosse, Alexander.Schiewe, Oliver.Bimber}@medien.uni-weimar.de

alizations in real world environments. It strives for enabling immersive and semi-immersive virtual reality, as well as augmented reality experiences without the need for special display surfaces or permanent screen configurations. The main focus lies on the visualization of and the interaction within existing buildings in the course of building surveying and planning processes, and supporting early architectural design phases. Unstructuredly aligned projectors and cameras are applied for an ad hoc visualization of interactive data on arbitrary surfaces within real-world indoor environments.

We have engineered a computer-controlled pan-tilt-zoom (PTZ) camera equipped with a laser module that automatically captures geometry and reflectance properties of the surrounding environment. After self-registration, the system is used for calibration the arbitrarily aligned projectors fully automatically. During run-time it enables basic laser pointer interaction in such real environments. Using a fish-eye context camera the device constantly captures a full 180deg hemispherical angle of its environment in low resolution. When detecting a laser-spot, it aligns a high-resolution zoom camera for a precise estimation of its position. During calibration, the system will analyze the scanned geometry and reflectance properties of the environment, and automatically detects and masks out areas in which a laser pointer tracking is not possible (e.g., light emitting surfaces, such as windows or lamps).

2 RELATED WORK

In this section we summarize the relevant related work in two areas: general laser pointer interaction as well as approaches that apply mechanically controlled cameras and projectors for calibration, interaction and scene acquisition. Note that discussing the corpus of general projector-camera systems that apply mechanically static cameras for calibration is out of the scope of this paper. Good overviews can be found in [9] for multi-projector systems and mainly optimized screen surfaces, and in [6, 8] for projector-based augmentations of complex everyday surfaces.

2.1 Laser Pointer Interaction

Laser pointer interaction is an intuitive technique for interacting over distances. In the simplest case, a laser dot is tracked by a single calibrated camera on a projection screen. For this, the laser position must be localized in the camera's image and has then to be

¹<http://www.sARc.de>

transformed into screen coordinates to -for instance- emulate mouse movements. The optical appearance and disappearance of the dot can be mapped to simple mouse click events [18]. Other techniques apply additional laser dots for communicating interaction parameters optically [20].

More efficient mapping strategies have been introduced for triggering events by applying special temporal or spatial gestures [13]. Optimized gesture parameters and device form factors are derived from user studies [23][21]. To avoid encoding trigger events optically, other approaches simply equip laser pointer devices with physical buttons and wireless communication electronics[22] [5].

Since the jitter while pointing over large distances can be enormous, the tracking data is usually smoothed. This together with the camera latency and additional mis-registration, however, leads to certain inconsistencies between laser dot and estimated position. This is in particular irritating if both - laser dot and cursor are visible simultaneously. The application of infrared laser pointers is one possible solution to this problem [12]. However, user studies have shown that -despite mis-registrations- interaction with visible laser dots outperform the interaction with invisible (e.g., infrared) laser dots [10].

More recent work focusses on laser-pointer interaction together with large screen surfaces and tiled displays. A method using an unstructured set of calibrated cameras was described in [1]. In particular for large area displays multi-user interaction becomes desirable. Thus different laser pointers must be distinguishable from each other. This can be achieved by multiplexing the laser light in time [22] or in wavelength [14]. More degrees of freedom can be reconstructed for laser-pointer tracking by projecting multiple laser dots (e.g., two [19], three [20]) or entire patterns [27].

While the related work described above supports laser pointer tracking on single optimized display screens (i.e., regular, flat and white projection screens), other approaches extend this to multiple (to our knowledge no more than two [25]) planar -yet still optimized- screen surfaces.

2.2 Steerable Cameras and Projectors for Calibration, Interaction and Scene Acquisition

Originally, PTZ cameras are utilized mainly for video surveillance applications. Recently, however, they are also applied for the calibration of multi-projector systems, for vision based interaction, and for scene acquisition tasks.

An un-calibrated PTZ camera, for example, supported the automatic alignment of high-resolution, tiled multi-projector displays [11]. In contrast to the application of wide-angle cameras, the adjustable camera allows for sequentially covering the large display area in an acceptable resolution.

A manually controlled (not motorized) PTZ camera equipped with a laser module was used to support the mobile acquisition of geometric building data. Thereby, the relative orientation of the camera was detected by shaft encoders mounted at the rotation axes. Mapping captured textures composed from several single images onto the geometry leads to depth enhanced panoramas (DEP) [4]. Multiple DEPs can be merged to model large scale building interiors [2, 3].

Steerable projector-camera systems have been used to convert optimized everyday surfaces into interactive displays (e.g., [24] or [17]). The system calibration for such approaches is performed manually for more complex surfaces by assuming that display surfaces and interaction widgets are known, or are defined during the calibration process (manually or by detecting aligned marker tags). For planar surfaces, homography matrices can be estimated automatically. After calibration, hands, heads, fingers, or tools with attached marker tags can be recognized and tracked by such systems if the environment illumination is high enough - which is generally disadvantageous for projector-based augmentations.

The contribution of our work is a system that supports laser pointer tracking together with projector-based augmentations within arbitrary real environments. Thereby, the screen surfaces are neither optimized for projections nor for laser pointer tracking. They can be geometrically and radiometrically complex. The system performs a fully self-controlled calibration of unstructuredly aligned PTZ camera and projectors, and automatically acquires scene properties that improve visualization and tracking quality. To our knowledge, this has not been achieved so far. Furthermore, we present a distributed software framework that couples laser-pointer tracking for interaction, projector-based AR as well as video see-through AR for visualizations with the domain specific functionality of existing desktop tools for architectural planning, simulation and building surveying.

3 SYSTEM OVERVIEW

This section provides details on our custom-built PTZ camera prototype and on the distributed software architecture that links all soft- and hardware components together.

3.1 Custom-Built PTZ Camera

Off-the-shelf PTZ cameras that are applied for video surveillance or video conferencing do not satisfy the demands of the laser pointer tracking application described in this paper. They lack in performance, quality and control possibilities. This gave motivation for designing and developing the mechanics and electronics of a customized PTZ camera prototype that satisfies our needs (cf. figure 2a).




Figure 2: Current PTZ camera prototype (a) and system overview (b).

A circuit-diagram is illustrated in figure 2b. Our current prototype consists of two video-cameras - a low resolution wide angle *context camera* (2) and a high resolution PTZ *detail camera* (4). A micro-controller (3) and two stepper-motors (8,9) with their controllers (6,7) are applied for rotating the detail camera and an attached laser module (5). Both cameras are directly connected to a PC (1) that analyzes the video images and controls motors, camera settings, and the laser module over the micro-controller (Atmel ATmega32) that is connected through an UART-USB converter.

The micro-controller functions as I²C-master for communication with the two stepper-controllers. As stepper-controllers, we chose Trinamic TMC222. This is a small but powerful chip which integrates 16× micro-stepping, 16bit position-counter, linear ramp-generator, in-progress programming and several diagnostic functions. A 29V/800mA H-Bridge driver stage is also included.

The stepper-motors have a resolution of 1.8deg per step. By using 16× micro-stepping and a gear-ratio of 1:6 the resulting resolution is 19000 micro-steps/360deg on both axes. Around the y-axis (pan), the camera can perform 3.4 full turnarounds (limited by the counter-depth of the TMC222 (16bit)). A full 360deg turn is possible around the x-axis (tilt).

The low resolution (320×240px) context camera applies a fish-eye lens [15] that produces 180 degree full circle view. It is mounted at the top of the camera-system to avoid occlusions with

the detail camera. The detail camera supports full PAL-resolution (720×576px) images, auto- and manual focus, zoom (f: 3.9-85.8mm), as well as variable shutter and aperture. All parameters can be controlled by the PC.

The 80mW, 650-658nm, laser module (5) on top of the detail camera is used for distance measurements during self-calibration. It can be triggered by a command from the PC through the microcontroller.

On startup of the PTZ camera the micro-controller initializes the motor parameters in the TMC222 and moves the motors to a reference position that is encoded by a switch and a reflex-coupler.

3.2 Distributed Software Architecture

For handling all hard- and software components, such as projectors, cameras and applications we decided to apply an existing distributed software architecture that is based on a communication framework used to link networked architectural applications [16]. A central lightweight server manages a database which stores raw architectural data (such as geometry and surface properties) that can be accessed by clients through shared software kernels. The distributed architecture (cf. figure 3) allows existing architectural *application clients* to access the functionality provided by new *service clients*. Shared libraries (called *kernels*) act as interfaces to the server and provide basic functionalities to all clients. This structure makes the system extremely flexible and extendable.

The different clients operate on the centrally stored building data. Application clients provide architectural functionalities, such as tachymeter-based, photogrammetry-based, or sketch-based geometry acquisition of building structures, material and lighting simulations, or modeling and inventory management. These clients represent a set of well-established working tools for architects - but support only desktop-based visualizations and interactions (cf. figure 4). To enable projector-based augmentations and laser-pointer interaction, individual service clients and kernels have been developed and integrated into this system. They provide specific functionalities, such as real-time image corrections for projecting onto complex surfaces, projector-camera calibration, synchronization management of different devices and laser pointer tracking.

The interaction among the different clients and the server can be explained best by describing an application case: A first step of the architectural surveying and planning process within existing buildings is the inventory of the building structures. Optionally, tachymeters or photogrammetry are used for exact geometry measurements. Alternatively, rough measurements, sketches and numerical constraints-optimization can be applied for less precise geometric estimates. For both situations, individual *manual geometry*

acquisition clients are used by architects. The building structure is stored in the database of the server and is available to all clients. Other application clients, such as the *material&lighting simulation* client accesses the building structure from the database to perform global lighting or material simulations based on the acquired geometry. Results are normally displayed on desktop screens (cf. figure 4).




Figure 4: Desktop application client example used for lighting and material simulations [26].

For supporting a projector-based augmentation of such simulation data within the actual building environment, it is necessary to calibrate projectors and cameras to the building structure initially. The *camera calibration* service client is used first to register the PTZ camera to the surrounding environment, as explained in section 4. This requires access to the building structure previously acquired and stored on the server. Furthermore, the *automatic scene acquisition* service client samples scene properties, such as surface reflectance and light sources, as explained in section 6. These parameters are stored in textures and have to be mapped correctly onto the acquired building geometry. This is handled by the *mapping* service client that ensures the correct mapping and storage of the image-based parameter information with respect to possibly low-resolution building geometry. As the geometry, the detailed scene properties are stored in the database and can be accessed by every client.

After the PTZ camera is calibrated and scene parameters have been acquired, it is used to calibrate the projectors. Since each device can be connected to a different PC, a spatial distribution of hardware components is possible. Structured light projection and camera feedback (produced and captured by multiple instances of the *projector calibration* service client) allow to calibrate the projectors with respect to the building structure, as described in section 5. For this, camera and projectors need to be synchronized. This is managed by a special *synchronization* kernel that triggers projections and capturing steps. Again, this requires access to the building structure and the scene parameters stored in server's database.

Finally, the system is ready to be used by application clients. If simulation data needs to be visualized within the real environment, the corresponding application client uses possibly multiple instances of the *projection* service client (one for each projector). These clients can carry out multiple image correction steps to compensate for geometric and radiometric distortions caused by the underlying surfaces before the images are projected out. The surface properties that are necessary for these computations can be accessed in the database through the correct mapping. Both have been acquired during the system calibration. Laser pointer-based interac-

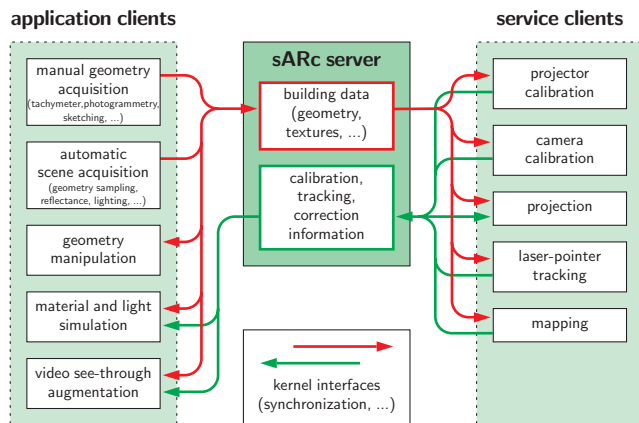


Figure 3: Distributed software architecture overview.

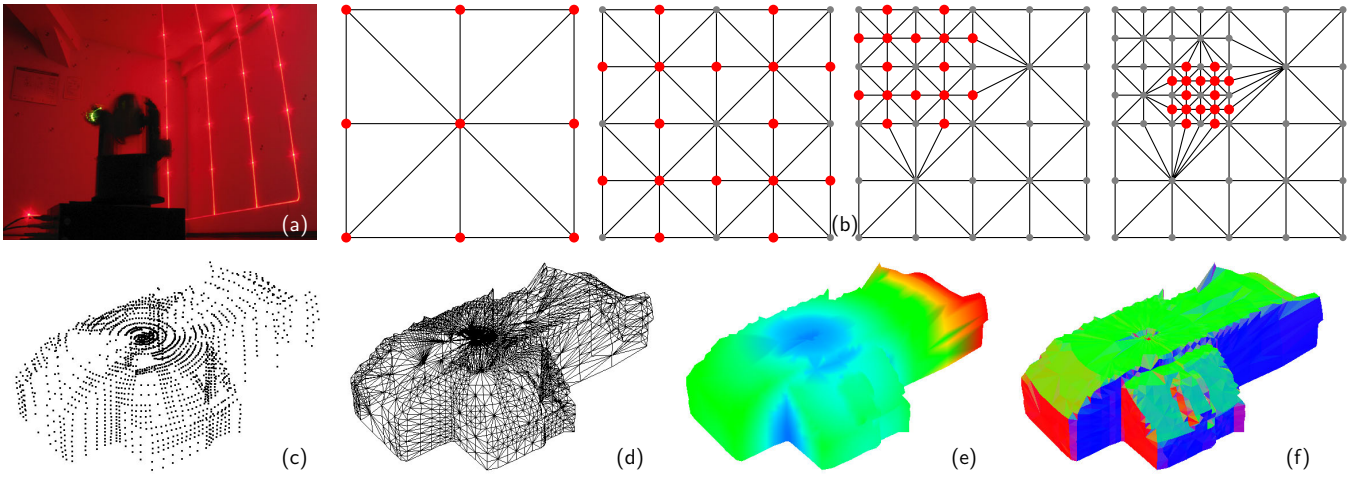


Figure 5: Sampling surface geometry: long exposure recording of the sampling process of the PTZ camera (a), adaptive multi-resolution sampling examples with vertex connectivity for the first four sample levels (b), sampled point cloud of entire room (c), triangulated surface geometry (d), depth map (e) and normal map (f).

tion techniques can be supported by the *laser-pointer tracking* service client that delivers the tracking data directly to the application clients that map them to their individual functionalities.

Besides projector-based augmentations, video see-through augmentations are supported through the *video see-through augmentation* application client that receives the live video image from the PTZ camera. In contrast to projector-based augmentations, video see-through augmentations enable the visualization of floating building structures. Note again, that synchronization between all application and service clients is managed by the *synchronization* kernel.

4 PTZ SYSTEM CALIBRATION AND REGISTRATION

This section explains the internal calibration of the PTZ camera system, the transformations between different local and global coordinate systems, as well as the steps of the automatic self-registration that has to be carried out before the system can be used.

4.1 Internal Pre-Calibration

Internal system parameters of the PTZ camera are pre-calibrated initially (one time) and remain independent from the actual application of the system. These parameters include all intrinsic values of the context camera (c) and the detail camera (d), the matrices (T_{c2s} , T_{d2s}) that describe the transformations between the individual local coordinate systems of both cameras (considering the rotations of the stepper-motors) relative to the fixed device coordinate system (s). This is illustrated in figure 6.

The intrinsic parameters (principal point u_0, v_0 , focal lengths f_u, f_v , and lens distortion parameters k_1, k_2, p_1, p_2) of the detail camera change for different zoom levels. Since it is difficult to find a precise enough analytical mapping, the parameters are measured for 22 discrete zoom levels f_i within a range of 3.9mm-13.9mm (larger zoom levels were not required), and intermediate (arbitrary) zoom steps f are piecewise linear interpolated. Using these parameters, rays that pass through pixels (u, v) on the image planes of either detail or context camera can be computed for their local coordinate systems, and are then transformed to the device coordinate system using T_{d2s} or T_{c2s} .

The laser module is calibrated relative to the detail camera by measuring a set of pixel projections (u, v) of the laser spot at different distances relative to the detail camera. This allows computing a continuous depth function $D_i(u, v)$ for each calibrated zoom level i .




Figure 6: Transformations between local and global coordinate systems of the PTZ camera.

For the same reason as for the intrinsic parameters, the estimated depth values at intermediate zoom levels are linear interpolated. Given T_{d2s} and the estimated depth $D_i(u, v)$ of a sampled surface point relative to d , the coordinates x, y, z of this point in s can be computed. Note, that lens distortions of the detail camera are corrected for all zoom levels initially.

4.2 Automatic Self-Registration

When the PTZ camera is positioned at the architectural site, a full automatic self-registration is carried out. This leads to a matrix T_{s2w} that transforms the device coordinate system (s) to the world coordinate system (w). The building structure as well as the origin of the world coordinate system are known. They have been measured and defined during the early steps of the architectural survey process, and stored in the database (see section 3.2). Thus a low-resolution reference model of the environments geometry can be accessed.

For registration, the PTZ camera samples the depth of the environment using the attached laser module (cf. figure 5a). To avoid unnecessary motor movements (and resulting long registration times), the sampling is not carried out uniformly along the hemispherical sampling space of the PTZ camera. Instead, an adaptive multi-resolution sampling is carried out: only if the divergence

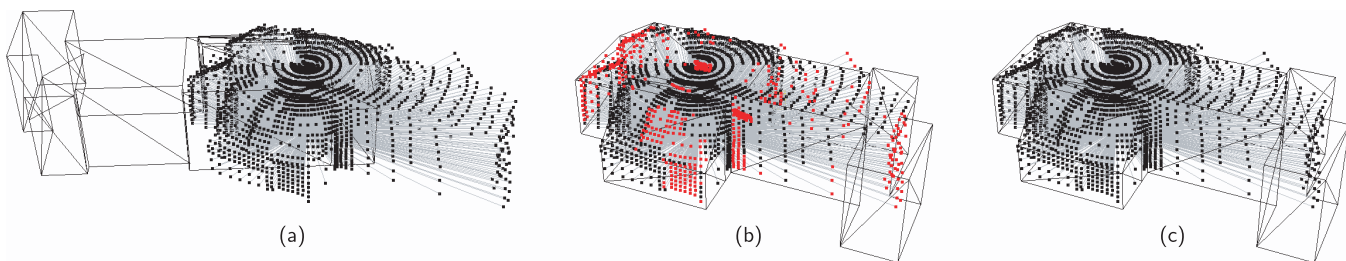


Figure 7: Unregistered sampled points and reference model (a), first registration step with identified outliers marked in red (b) and a mean registration error of 3.6 cm (in world coordinates), and final solution with minimal error of less than 1cm (c).

of the sampled world coordinates of nine points (eight corners and the center point of a quad) from their mean-tangential plane is above a pre-defined threshold, the area enclosed by the quad is further refined (cf. figure 5b). This is recursively repeated until the planarity condition is met or a predefined maximum recursion depth is reached. This results in a low-resolution sampling for largely planar environmental sections (such as walls), and in a high-resolution sampling for geometrically more complicated sections (such as corners).

The resulting point cloud (cf. figure 5c) can be triangulated (cf. figure 5d), intermediate depth values can be interpolated (cf. figure 5e) and normal vectors can be computed (cf. figure 5f). Since the sampled points are measured in device coordinates, the rigid transformation matrix T_{s2w} can be estimated numerically by finding the best geometric match between the point cloud and the measured reference model stored in the database. This is achieved with a variation of the iterative closest point algorithm. The error function, however, computes the least square error of the minimal geometric distance between the point cloud and the planes spanned by each triangle in the geometry of the reference model. The matching process is illustrated for an example in figure 7.

Extremely reflecting or absorbing surface areas cannot be sampled correctly. Furthermore, the scanned environment might contain geometric features that have not been measured in the reference model (such as furniture). All of this leads to a certain inconsistency between scanned point samples and the reference model. To avoid large registration errors during the numerical minimization, extreme outlier points are automatically identified and deleted during the iterative matching phases. Once a first minimum was found, all sampled points whose distance to their reference planes are above the mean distance over all points are identified as outliers and removed (cf. figure 7b). A second minimization pass is applied to achieve the final transformation (cf. figure 7c).

Besides sampling geometry, other surface parameters can be captured during this process. The diffuse surface reflectance, for instance is captured for each discrete perspective of the detail camera under a short and under a normal exposure. The individual image patches (cf. figure 8a) are later blended using linear ramps to create seamless overlap (cf. figure 8b) and stored in two spherical environment maps (cf. figure 9a+b). In our implementation, these environment maps have a resolution of 1000×1000 entries and are computed from $9 \times 9 = 81$ normal exposure and from the same amount of short exposure images. The acquisition of the surface reflectance for this resolution takes about 5 minutes and is part of the camera calibration process, as explained in section 4. The lion's share in this process is the relatively long duration that it takes to auto-focus the detail camera sufficiently to different distances.

These information are required for the subsequent scene analysis part (see section 6) and for radiometric compensation techniques, as described in section 5. An alternative to scanning high resolution patches from the detail camera during calibration is to map the context camera's low quality omnidirectional image to the reference

geometry in real-time. For static scenes, however, the high resolution map is clearly preferred over the low resolution one.

After registration, the PTZ camera delivers world coordinates. Furthermore, the detailed surface properties, such as scanned reflectance maps, depth maps, and normal maps can be registered to the low-resolution reference model and are stored in the database on the server to make them available to all application and service clients.

For sampling one single point during camera self-registration requires approximately 800ms on average (including camera movements, capturing and image analysis). With a maximum of 2000 samples we achieved an average registration precision of 0.19 degrees in our experiments. As explained above, scanning the surface reflectance in addition for the full 180 degree hemisphere in a resolution of $9 \times 9 = 81$ patches requires approximately 5 minutes. Thus, both processes are carried out fully automatic and require together about 30 minutes for a maximum sampling resolution – and less for lower sampling rates.

5 PROJECTOR CALIBRATION

After the camera is registered to the surrounding scene geometry, each vector (α, β, f, u, v) within the camera coordinate system can be correlated to a 3D point (x, y, z) within the world coordinate system. Thereby, x, y, z are the surface coordinates that are projected to the sub-pixel u, v on the image plane of the detail camera under pan-tilt-zoom settings α, β and f .

Given this mapping, unstructuredly aligned projectors can be registered automatically with respect to the common world coordinate system. Therefore, all projectors are calibrated sequentially. This is achieved by projecting coded patterns, detecting these patterns with the PTZ camera, and establishing a correlation between image coordinates of the projected patterns on each projectors' image plane and their 3D coordinates in the world coordinate system at which they are projected. The projected patterns are detected by the PTZ camera in the same way as detecting laser spots (see section 7.1). Finally, the intrinsic (i.e., fov, lens offset, aspect ratio, and radial distortion) and extrinsic (i.e., position and orientation in the world coordinate system) parameters of each projector can be estimated numerically. Calibrating a single projector by displaying, capturing and processing 49 structured code images takes about 1.5 minutes. For this, we achieve a registration accuracy of 10.7 mm, when comparing projected points of the reference model to known points on the physical building surfaces.

Furthermore, other surface properties, such as local reflectance, global light modulations (e.g., inter-reflection), or a pixel-precise geometric mapping can be measured in addition. All these detail parameters are stored in textures in the database that are correlated to the scene geometry. This ensures that multiple entries from different camera and projector settings are addressed correctly through unique look-up operations over the scene geometry. These surface parameters can then be used for pixel-precise geometric and radiometric compensation, consistent photometric projection and inten-




Figure 8: Sampling surface reflectance: 81 partially overlapping normal exposure images capture the surface reflectance at all possible angles (a), blended images that lead to a single consistent environment map (b), a close-up on a fraction of the blended environment map (c), and a photograph from the same view (d). Slight misregistrations are due to small calibration errors.

sity blending of overlapping image regions, as explained in related camera-projector approaches (see overviews in [9, 6, 8]).

After projector calibration, it is also known which surface areas are covered by which projector, in both - the world coordinate system and the device coordinate system of the PTZ camera.

6 SCENE ANALYSIS

Besides the optional acquisition of scene parameters through structured light sampling that enables the correct projection onto complex surfaces (see section 5), the video images that are captured during camera calibration (see section 4 and figure 8) can be analyzed. They reveal information about surface areas that are not appropriate for laser pointer tracking, such as non-diffuse surfaces and in particular light emitting areas, such as lamps or windows.

Since intensity thresholding is used to identify the laser spot in the camera image as explained in section 7.1, bright light emitting scene areas lead to wrong classifications. Applying additional pattern matching to differentiate the laser spot from such areas does not lead to satisfying results, because the shape of the laser spot can vary when being reflected from different surface areas (e.g., due to inter-reflections) or when captured under diverse motion speeds.

Instead, we identify and mask out light emitting surfaces during the camera calibration process. They will not be considered during




Figure 9: Spherical environment texture: captured under normal exposure (a) and under short exposure (b). Binary mask of invalid areas unsuited for laser pointer tracking (c).

intensity thresholding. Note, that non-diffuse surfaces do not have to be masked, since they usually do not reflect the laser spot appropriately, and can consequently not be used for tracking in most situations. In any case they do not influence the detection of the laser spot.

During camera calibration, bright light emitting surface areas are detected via intensity thresholding of the short exposure environment map that was captured with the detail camera (see section 4.2). The coordinates (α, β, f, u, v) of the remaining high intensity pixels are mapped into a spherical binary texture that encodes valid and invalid surface areas via unique spherical coordinates. This is illustrated in figure 9. After camera calibration, the final binary sphere map covers the entire upper hemisphere of the PTZ camera in high resolution and is registered to the world coordinate system.

The masking process takes place before the camera image is analyzed for detecting the laser spot: For a particular position of the detail camera (α, β, f) during run-time, a perspective image of the binary sphere map is off-screen rendered with the same camera settings. The resulting stencil mask matches the captured video image geometrically. Both images are multiplied for masking out invalid surface areas (that are zero-coded in the sphere map). The final result is then analyzed for detecting the laser spot during laser pointer tracking, while the invalid areas (i.e., the ones that would make the intensity thresholding fail) are ignored. Note, that such a masking is applied to the context and to the detail camera's images. For the context camera, the acquisition of the mask is a single threshold-binarized image which is rotated and translated with respect to the actual rotation of the camera and multiplied with the original image.

7 LASER POINTER INTERACTION

After calibration, the PTZ camera can be used for tracking a single visible laser spot that appears on a scene surface and is within the viewing range of either the detail camera or the context camera. We use an off-the-shelf 35nW, 532NM (green) laser pointer. In the following section, we first explain how the laser spot is tracked. Simple interaction techniques are finally outlined.

7.1 Tracking

Assume the detail camera is aligned in a direction so that the laser spot is within its viewing frustum. In this case we identify the laser spot in the camera image through intensity-thresholding. The exposure time of the detail camera is low to capture only bright luminances. Furthermore critical areas are masked out, as described in section 6. Computing the center of gravity of the remaining pixels produced by the laser spot leads to a sub-pixel precise identification of the laser spot's image (u, v) . Since the camera parameters (α, β, f) are known, a ray is computed in the world coordinate system using the known matrices T_{d2s} and T_{s2w} (see section 4) that is intersected with the registered scene geometry. This leads to the spot's three-dimensional surface coordinates (x, y, z) .

The detail camera is mechanically re-aligned to center the laser spot only if distance between the (u, v) coordinates and the image center is above a pre-defined threshold. This avoids unnecessary rotation of the steppers motors. Thus, tracking the laser spot solely with the detail camera is possible until it leaves the detail camera's

viewing frustum entirely. There are two reasons when this can happen: Either the laser pointer is turned off and turned on again - pointing somewhere outside the current perspective of the detail camera. Or the movement of the laser spot is too fast for stepper-motors and detail camera to track. In this case, the context camera is used to get a coarse position of the laser spot and the detail camera is aligned to capture the spot again (cf. figure 10). This happens also in the initial case, when the spot has to be searched.

Note, that panning and tilting the detail camera is processed on the micro-processor. Thus update commands can be sent in parallel - while the camera is in motion. Detecting the laser spot with the context camera and aligning the detail camera initially through a 180 degree rotation (in the worst case) takes approximately 1.5 seconds - but usually less if a shorter rotation is sufficient. We achieve 30 samples per second for tracking the detected laser spot with the detail camera (including camera movements, auto focus adjustments, capturing, and processing). The latency in this case (comparison between a projected cursor point and the actual laser spot) is approximately 300 ms (including camera latency, network latency, latency of involved software clients, such as projection client and laser-pointer tracking client). The deviation of the tracked laser spot compared with known surface points is on average 8 mm. The deviation between the laser spot and its projected counterpart is on average 14 mm (this includes the registration deviations of the projectors - see section 5).




Figure 10: Tracking a laser spot in a 3D environment: If visible in the detail camera (a), the position of the laser spot can be estimated. If the spot is lost (b), the context camera delivers information for re-aligning the detail camera (c).

7.2 Interaction and Visualization Examples

Receiving continuous information about the laser spot's 3D world coordinates on the surrounding surfaces and its current state (on/off), simple spatial and time encoded gestures are currently used to trigger events, such as changing the interaction mode or for picking and dropping objects. A cursor is rendered to display the current tracking position and acoustic signals are used to indicate different states (such as the beginning of a motion gesture). Neural networks and context free grammar parsing are applied as described in [7] for recognizing and interpreting motion gestures to carry out object selection, translation, rotation and scaling transformations (cf. figure 1a,b).

While a projector-based augmentation is best suited for the visualization information directly on the surfaces (e.g. figure 1c), such as lighting, color and material simulations or geometric structures with small depth variations, video see-through augmentations supports visualizations of free-floating 3D structures from the (controllable) perspective of the PTZ camera (e.g. figure 1d). Although both visualization methods are supported, our laser-pointer based interaction is currently constrained to the surfaces. Thus, points in free-space can only be selected through simple ray-casting in the video see-through mode (computing the ray from the camera's center to the laser dot on the surface). An arbitrary interaction in free-space is not possible at the moment. This requires either the tracking of




Figure 11: Two perspective views projected on the same surface. The viewer perspective is indicated with an active three-color LED marker.

more degrees of freedom or the investigation of adapted picking and transformation techniques - which belongs to our future work.

Yet, the PTZ camera itself can be used for a rough pose-estimation of active LED markers (cf. figure 11). Pressing a button on the marker panel triggers three differently colored LEDs. If no laser pointer is detected during this time, the context camera guides the detail camera towards the LED marker in a similar way as explained for laser spots in section 7.1. Once detected and classified via color thresholding, the full 6DOF pose of the marker in world coordinate system is delivered through conventional pose estimation. This technique can be used for making rough indications of the observers position or other positions and orientations within the world coordinate system. The LED marker can also be continuously tracked. Although the quality and speed of this simple technique (in our current implementation 10 samples per second) does not come close to professional tracking solutions, it proved to be a useful tool for many view-dependent architectural visualizations without notably increasing the system complexity through an additional tracking device. A high precision and tracking performance is not required for many of these cases, but could be achieved by applying better pose estimators. Note again, that the simultaneous tracking of the laser-pointer and the LED marker is not possible with a single PTZ camera system. Only a sequential tracking is currently supported.

8 SUMMARY AND FUTURE WORK

In this paper we have presented a system that applies a custom-built pan-tilt-zoom camera for laser-pointer tracking in arbitrary real environments. Once placed in a building environment, it carries out a fully automatic self-registration, registrations of projectors, and sampling of surface parameters, such as geometry and reflectivity. After these steps, it can be used for tracking a laser spot on the surface as well as an LED marker in 3D space, using inter-playing fisheye context and controllable detail cameras. The captured surface information can be used for masking out areas that are critical to laser-pointer tracking, and for guiding geometric and radiometric image correction techniques that enable a projector-based augmentation on arbitrary surfaces. We described a distributed software framework that couples laser-pointer tracking for interaction, projector-based AR as well as video see-through AR for visualization with the domain specific functionality of existing desktop tools for architectural planning, simulation and building surveying.

We see projected-based augmented reality as a potential interface for architectural applications, that might offer more flexible on-site visualizations of light and material simulations. With our laser-pointer tracking approach we have presented a first basic interaction component. This module has to be extended in future to support the tracking of more degrees of freedom (such as in [19], [20], or [27]). In particular the performance and precision of the PTZ camera have to be improved. This will have a direct impact on calibration and tracking speed and quality. We can also imagine the application of multiple synchronized PTZ cameras. This would allow covering a

larger area and to simultaneously track multiple entities (e.g., laser spots or LED markers).

Another interesting aspect that has to be investigated are the possibilities and limitations of adapted interaction techniques that can support new architectural tasks. Examples could include *material copy-and-paste* scenarios in which selected material samples are scanned, analyzed, enlarged via texture synthesis techniques, and finally re-produced at other surface portions via projector-based augmentations. Furthermore, the investigation of appropriate placement techniques for menus and other interaction items with respect to the surface geometry and reflectivity, as well as under consideration of constraints related to the laser-pointer tracking (e.g., surface visibility, hand jittering, etc.) will be part of our future work.

Finally, an in-depth evaluation has to be carried out together with professional architects to investigate whether or not such a spatial augmented-reality interface is useful. Throughout an initial informal user study, 25 professional participants (architects) gave generally positive feedback, but also pointed out limitations and required improvements. To address them will be our main focus.

ACKNOWLEDGEMENTS

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) under contract number PE 1183/1-1. We would like to thank Anselm Grundhöfer and Christian Tonn for help with the implementation.

REFERENCES

- [1] B. A. Ahlborn, D. Thompson, O. Kreylos, B. Hamann, and O. G. Staadt. A practical system for laser pointer interaction on large displays. In *VRST '05: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 106–109, New York, NY, USA, 2005. ACM Press.
- [2] G. Bahmutov, V. Popescu, and M. Mudure. Efficient large scale acquisition of building interiors. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Sketches*, page 127, New York, NY, USA, 2006. ACM Press.
- [3] G. Bahmutov, V. Popescu, and M. Mudure. Efficient large scale acquisition of building interiors. In *Proc. Eurographics*, pages 655–662, 2006.
- [4] G. Bahmutov, V. Popescu, and E. Sacks. Depth enhanced panoramas. In *VIS '04: Proceedings of the conference on Visualization '04*, page 598.11, Washington, DC, USA, 2004. IEEE Computer Society.
- [5] X. Bi, Y. Shi, X. Chen, and P. Xiang. Facilitating interaction with large displays in smart spaces. In *sOc-EUSAI '05: Proceedings of the 2005 joint conference on Smart objects and ambient intelligence*, pages 105–110, New York, NY, USA, 2005. ACM Press.
- [6] O. Bimber. Projector-based augmentation. *Book Chapter in Emerging Technologies of Augmented Reality: Interfaces and Design*, pages 64–89, 2006.
- [7] O. Bimber, L. Encarnao, and A. Stork. A multi-layered architecture for sketch-based interaction within three-dimensional virtual environments. *Computers and Graphics*, 24(6):851–867, 2000.
- [8] O. Bimber, D. Iwai, G. Wetzstein, and A. Grundhöfer. The visual computing of projector-camera systems. In *Proc. of Eurographics (State-Of-The-Art Report)*, 2007.
- [9] M. Brown, A. Majumder, and R. Yang. Camera Based Calibration Techniques for Seamless Multi-Projector Displays. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):193–206, 2005.
- [10] D. Cavens, F. Vogt, S. Fels, and M. Meitner. Interacting with the big screen: pointers to ponder. In *CHI '02: CHI '02 extended abstracts on Human factors in computing systems*, pages 678–679, New York, NY, USA, 2002. ACM Press.
- [11] Y. Chen, D. W. Clark, A. Finkelstein, T. C. Housel, and K. Li. Automatic alignment of high-resolution multi-projector display using an un-calibrated camera. In *VIS '00: Proceedings of the conference on Visualization '00*, pages 125–130, Los Alamitos, CA, USA, 2000. IEEE Computer Society Press.
- [12] K. Cheng and K. Pulo. Direct interaction with large-scale display systems using infrared laser tracking devices. In *APVis '03: Proceedings of the Asia-Pacific symposium on Information visualisation*, pages 67–74, Darlinghurst, Australia, Australia, 2003. Australian Computer Society, Inc.
- [13] J. Dan R. Olsen and T. Nielsen. Laser pointer interaction. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–22, New York, NY, USA, 2001. ACM Press.
- [14] J. Davis and X. Chen. Lumipoint: Multi-user laser-based interaction on large tiled displays. *Displays*, 23(5), 2002.
- [15] H. G. Dietz. Fisheye digital imaging for under twenty dollars. Technical report, University of Kentucky, Apr 2006.
- [16] D. Donath and T. Thurow. Integrated architectural surveying and planning. methods and tools for recording and adjusting building survey data. In *Automation in Construction, Volume 16, Number 1*, pages 19–27, 2007.
- [17] J. Ehnes, K. Hirota, and M. Hirose. Projected augmentation augmented reality using rotatable video projectors. In *IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR04)*, pages 26–35, 2004.
- [18] C. Kirstein and H. Mueller. Interaction with a projection screen using a camera-tracked laser pointer. In *MMM '98: Proceedings of the 1998 Conference on MultiMedia Modeling*, page 191, Washington, DC, USA, 1998. IEEE Computer Society.
- [19] S. V. Matveyev and M. Göbel. Direct interaction based on a two-point laser pointer technique. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Sketches & Applications*, pages 1–1, New York, NY, USA, 2003. ACM Press.
- [20] S. V. Matveyev and M. Göbel. The optical tweezers: multiple-point interaction technique. In *VRST '03: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 184–187, New York, NY, USA, 2003. ACM Press.
- [21] B. A. Myers, R. Bhatnagar, J. Nichols, C. H. Peck, D. Kong, R. Miller, and A. C. Long. Interacting at a distance: measuring the performance of laser pointers and other devices. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 33–40, New York, NY, USA, 2002. ACM Press.
- [22] J. Oh and W. Stuerzlinger. Laser pointers as collaborative pointing devices. In *Graphics Interface 2002*, pages 141–149, 2002.
- [23] C. H. Peck. Useful parameters for the design of laser pointer interaction techniques. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 461–462, New York, NY, USA, 2001. ACM Press.
- [24] C. Pinhanez. Using a steerable projector and a camera to transform surfaces into interactive displays. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 369–370, New York, NY, USA, 2001. ACM Press.
- [25] J. Rekimoto and M. Saitoh. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 378–385, New York, NY, USA, 1999. ACM Press.
- [26] C. Tonn and D. Donath. Color, material and light in the design process – a software concept. In *Proceedings of the Joint International Conference on Computing and Decision Making in Civil and Building Engineering, ICCCB 2006*, pages 1467–1476, Montral, Canada, 2006. IEEE Computer Society.
- [27] C. Wienss, I. Nikitin, G. Goebels, K. Troche, M. Göbel, L. Nikitina, and S. Müller. Sceptre: an infrared laser tracking system for virtual environments. In *VRST '06: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 45–50, New York, NY, USA, 2006. ACM Press.