Projection-Based Augmented Reality in Disney Theme Parks

Mark Mine, David Rose, and Bei Yang, *Walt Disney Imagineering*

Jeroen van Baar and Anselm Grundhöfer, *Disney Research Zürich*

Walt Disney Imagineering and Disney Research Zürich are building a projector-camera toolbox to help create spatially augmented 3D objects and dynamic, interactive spaces that enhance the theme park experience by immersing guests in magical worlds.

Since the first day Disneyland opened its doors in 1955, guests have come to Disney theme parks to immerse themselves in magical worlds created by Walt Disney and the Imagineers at Walt Disney Imagineering, headquartered in Glendale, California. Today, WDI, originally known as WED Enterprises, is the design, engineering, and production arm of the Walt Disney Company, and is responsible for developing Disney theme parks, resorts, cruise ships, and other entertainment venues worldwide.

Imagineers employ an array of tools to create the alternate realities otherwise known as Disney theme parks. Story, music, art, and architecture combine with science, engineering, and advanced technology to take guests to the "World of Yesterday, Tomorrow, and Fantasy."

Recently, Imagineers have begun exploring how to combine Disney storytelling, creativity, and artistry with advanced projection technology and computer graphics. Here, we present an overview of this exploration, describing several loosely related techniques that fall into the category of *projection-based augmented reality* (AR). This work is a collaboration between Imagineers and scientists and technologists at Disney Research Zürich in Switzerland. DRZ is part of the network of research labs founded by Disney in 2008 to closely collaborate with academic institutions such as Carnegie Mellon University and the Swiss Federal Institute of Technology Zürich.

PROJECTION-BASED AR OVERVIEW

The AR community defines projection-based AR as the use of projection technology to augment and enhance 3D objects and spaces in the real world by projecting images onto their visible surfaces. This relates to the Shader Lamps research by Ramesh Raskar and colleagues[1] and falls into the general category of spatial AR as defined by Raskar and Oliver Bimber.[2] The projected images can be computer generated or photographic, and either prerendered or generated in real time.

Projection-based AR typically uses one or more projectors arranged around an object (such as a prop or character in an attraction) or distributed throughout a 3D space (such as an entire scene). If the display uses multiple projectors, their images can be independent of one another or blended together, either manually or via automatic techniques. These displays can casually align images with physical objects or precisely register them to align with surface features. If precisely registered, the projected images can replicate surface colors and features to create a multipli-

cative effect on color, saturation, and contrast. Doing so yields stunning high dynamic range (HDR) results. This effect relates to superimposed dynamic range imagery[3] and can have a dramatic visual impact.

The chief advantage of projection-based AR is that it can create beautiful dynamic environments and bring sets to life in a magical way difficult to achieve with traditional lighting. HDR lighting, per-pixel control, animated media, and interactive content are all exciting new tools for use by theme park designers.

In addition, projection-based AR augments the space around guests and as such creates a shared experience for simultaneous viewing by multiple people. This is a significant advantage in a theme park that hosts tens of thousands of people a day. In contrast, device-based AR techniques, such as head-mounted displays or handheld mobile devices, do not scale as well; they are typically single-user devices, particularly when the augmentation must register (fit precisely) with the real world.

A third advantage of projection-based AR systems is that they have fewer issues with latency than see-through AR devices because the projector and the augmented objects have less relative motion. Many theme park environments consist of only a few slow-moving elements, and the objects typically follow known paths. See-through AR devices, on the other hand, suffer from significant latency artifacts unless the display cannot move.

Finally, projection-based AR is more compatible with Imagineering's design and aesthetic philosophy since it is easier to hide the technology. Items as obvious as handheld mobile devices or computer monitors often conflict with the design and theme of a particular environment. Although modern children might find it difficult to believe, cell phones did not exist on the American frontier, and tablet computers were not on pirate ships.

PROJECTION-BASED AR IN DISNEY PARKS

The most famous early examples of projection-based augmentation in Disney parks were the projected heads in the Haunted Mansion. These projections include Madame Leota, the ghostly head appearing inside a crystal ball in the séance scene, and the quintet of singing busts from the graveyard scene. As Figure 1 shows, by projecting film of actors singing and talking onto physical busts, ghostly characters magically come to life.

Building projection

A more recent use of projection-based AR in Disney parks is building projection. As Figure 2 shows, projection-based AR is an excellent and affordable way to augment and activate existing spaces, transforming them without making significant structural and facility changes. It also facilitates rapidly overlaying new experiences onto existing spaces, important for special or seasonal events.
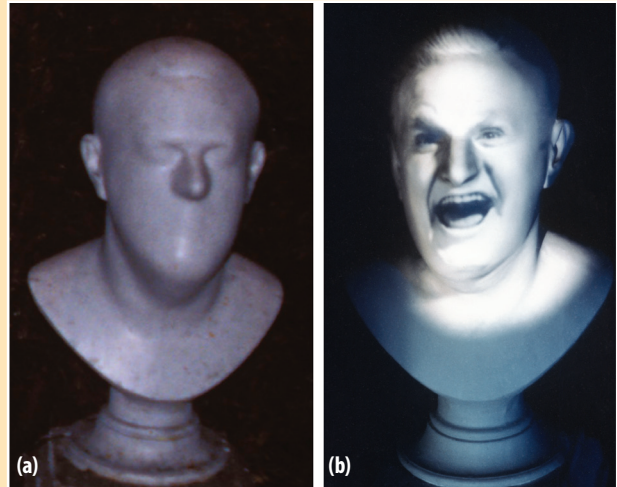
Figure 1. One of the quintet of singing busts in the Haunted Mansion graveyard scene shown (a) without and (b) with projection-based augmentation.

Figure 2. Cinderella's Castle at the Magic Kingdom Park in Orlando, Florida, comes to life during *The Magic, the Memories, and You!*, a show that displays images of park guests taken during that day.

Building projection, for example, helped "unwrap" the new Tower of Terror attraction at Disneyland Paris. The entire building was covered with projected wrapping paper (along with a few bus-sized cockroaches) that was torn off to symbolize the attraction's opening. Similarly, every Halloween, nebulous ghosts fight to burst out of Space Mountain, and Christmas would not be complete without holiday-themed images dancing on the building exterior of "it's a small world."

Scenic projection

Recently, we have begun to explore using projection-based AR, specifically, HDR imagery, to energize scenic elements in attractions.

Figure 3. Upper images, taken during early tests, show a set under (a) normal room lighting conditions and (b) with projected augmentation. Lower image (c) shows the technique applied to figures in Snow White's Scary Adventures. Note that printed images fall short: only in-person viewing conveys the full impact of the HDR results.

In 2010, Disney added projections to several sections of Snow White's Scary Adventures, located in Disneyland's Fantasyland. The idea was to make 50-year-old scenes come alive in ways they never have before.

As Figure 3 shows, HDR illumination and supersaturated colors help scenes pop, making characters stand out far more effectively than with conventional lighting. Furthermore, projecting animated media helps enhance effects such as lightning flashes and magical discharges, dramatically transforming the look of entire scenes, such as when the evil queen turns into the ugly hag.

This work also exemplifies many of the challenges in introducing advanced display technology into a theme park environment. Imagineers must overcome several hurdles when integrating modern technology into a 50-year-old ride system.

Interactive applications

The ability to project real-time imagery onto surfaces and objects enables exciting new forms of play in Disney parks as well. The 2009 D23 Expo (a Disney fan event in Anaheim, California) presented The Storytellers Sandbox, an interactive play space that projected images and effects onto the surface of a table filled with sand. This offered a new multidimensional canvas that guests could literally dig into and modify as they heard stories and participated in interactive activities. As Figure 4 shows, guests could pile sand into a volcano with projected flowing lava or dig

a hole to receive a projected sea turtle egg. Later, the egg hatched to reveal a baby sea turtle that would scramble into the ocean waves (also projected).

We have even explored the potential of projecting on unexpected surfaces such as cakes and water. For example, projected media can enhance cakes with animation, HDR imagery, and interactive techniques. A depth-sensing camera can detect when someone removes a slice from a cake, triggering the release of a swarm of butterflies or a stream of "grim grinning ghosts" (image choice depends on the child) that fly around the cake's surface. At Siggraph 2011, we presented our Thermal Interactive Media display, which combined interactive projected imagery with water.

THEME PARK CHALLENGES AND CONSIDERATIONS

Although projection-based AR affords many exciting opportunities for new entertainment applications, Imagineers must consider numerous challenges before extensively using this technology in Disney parks and elsewhere.

Multiprojector calibration

The number of projectors in Disney attractions has increased dramatically in recent years, and the era of blended multiprojector displays is clearly here to stay. For projected content to accurately correspond to surface features on physical geometry, or for multiple projectors to produce

a single continuous image, it is first necessary to precisely align projectors to content or to each other—a process Imagineers refer to as calibration. Furthermore, projectors inevitably drift over time, so maintenance or some other mechanism must restore this alignment. Multi-projector calibration, therefore, is the first big challenge of integrating projection-based AR into theme park environments.

Manual techniques for projector blending and calibration are imprecise and time-consuming (and therefore expensive). In addition, the difficulty of aligning multiple projectors increases exponentially with the number of projectors, since each projector's proper alignment also depends on the alignment of its neighbors. Most importantly, standards for quality within Disney parks are extremely high; any obvious errors or artifacts will take guests out of the experience.

We believe that achieving and maintaining the required level of precision will require automatic tools that can identify and aid in precisely aligning projected media to match specific surface features. These tools also must align and blend multiple projectors to form a single continuous image. Errors must be smaller than a single projector pixel whenever possible.

Operational and maintenance considerations

Keeping operation and maintenance costs in a theme park as low as possible is essential. This imposes numerous constraints on the design of any technology installed in Disney parks.

Operations 365 days a year leave little time for maintenance, most of which occurs at night. If a ride breaks down, maintenance workers must repair it quickly to avoid guest dissatisfaction. So there is no time to diagnose obscure failures in complex subsystems.

Disney also does not expect its maintenance workers to have expertise in specialized areas such as computer vision. Thus, any automated system must be fully self-reliant, requiring almost no supervision and producing acceptable results from the available input whenever possible. It must automatically identify and report failed components or specific problems in the operating environment such as unexpected lighting changes or obscured line of sight that might have contributed to failure or poor results.

The operating environment poses challenges as well. Although most computer vision experts consider camera calibration a solved problem, traditional calibration techniques do not always work well in a theme park. For example, traditional techniques require capturing several images of a physical pattern (such as a checkerboard) in

Figure 4. The Storytellers Sandbox at the D23 Expo.

a variety of positions and orientations with respect to the camera. This is challenging in an environment where the camera is permanently mounted high above a physical set. Because the calibration pattern must be large enough for the camera to see clearly, moving throughout a complex space becomes unwieldy.

In addition, lighting in theme park attractions can also be uneven or absent altogether. Parts of the set are also inaccessible to maintenance workers. Finally, maintenance workers are unlikely to be experienced with computer vision techniques, and might not accurately identify issues that affect calibration, such as warped boards, poor contrast, motion blur, or poor coverage in the resulting images.

Content authoring

Traditional lighting, set design, and media generation techniques are well-suited for traditional theme park applications, but might not work as well in the new workflows required for projection-based AR. Artists and tools, for example, can have a 2D-projector-centric point of view, working more in terms of the media coming from a single projector. Content that spans multiple projectors, such as a projected butterfly flitting across a scene, for example, might require manual reproduction and alignment of each projector. Consequently, changing the number or layout of projectors in the scene might require regenerating or reworking the artwork.

Many excellent 3D modeling and animation tools can support the notion of a 3D projected space. However, some lack the features and expressiveness of 2D tools. Furthermore, artists can be reluctant to embrace these new 3D tools, media, or workflows.

Though it is possible to envision a day when it will be possible to directly augment the physical 3D space, paint-

ing projected light directly onto 3D surfaces, our near-term goal is to develop systems that let artists work with existing commercial tools. Bridging the gap requires new tools to convert content from traditional media into a form appropriate for AR. For example, artists can use traditional 2D tools to paint a scene as viewed from a single well-defined point of view (that of a camera placed within the scene). They can then automatically distort that content, to apply it to the same scene from a different point of view (that of an overhead-mounted projector). This lets artists continue to use the traditional pipeline as needed, while providing a transition path to a new, more direct workflow in the future.

Life span of theme park attractions

Theme park attractions are extraordinarily long-lived. Typical designs plan for them to be in use for at least 10 to 20 years; some Disneyland attractions have been in operation for more than 50 years.

> **Our goal is to develop a toolbox that provides a suite of methods and algorithms to design and support new projection-based AR installations.**

This extended life span presents unique challenges for any technology-based solution. Obtaining replacement parts can be difficult or impossible after only a few years, as computer systems rapidly become outdated. Techniques that were once standard and well-understood might be unfamiliar to a new generation of designers. Even maintenance tools such as laptops that are commonplace today might become difficult to obtain in the future.

MOTIVATION FOR BUILDING AR TOOLS

ProCams systems can address the challenge of calibrating multiprojector installations.[4] However, the current commercially available systems lack the flexibility and robustness needed to accommodate the variety of scenarios encountered in projection-based AR for theme parks and other entertainment applications.

We have thus elected to develop a ProCams toolbox internally to support projection-based AR installations in theme parks. Developing such a toolbox internally has the following advantages:

- we have full control over the toolbox architecture and can support modular software design,
- a modular toolbox will make it easier to incorporate newly developed algorithms and methods for use in theme park installations, and

- a modular toolbox will make it easier to integrate existing software with products from third-party vendors.

Therefore, our goal is to develop a toolbox that provides a suite of methods and algorithms to design and support new projection-based AR installations. We also hope to motivate and inform the commercial development of AR tools to support the comprehensive systems we envision for the future.

PROJECTOR-CAMERA TOOLBOX

To support automated registration of projectors with respect to the attraction's geometry, our toolbox consists of a collection of techniques and algorithms for developing ProCams systems.

Architecture

Our ProCams toolbox consists of a node-based architecture that includes input devices, such as cameras; processing operations; and output devices, such as projectors or displays. Nodes connect via queues, and connecting a set of appropriate nodes creates applications. The node-based architecture establishes connections at compile time. A node can have zero or more inputs as well as outputs. A camera node, for example, has only a single output, that is, the acquired image. An operation node, such as feature detection, takes an image as input and produces a set of detected features as output. Each node creates its own thread for maximum flexibility and efficient parallel processing.

Figure 5 shows an example application based on our toolbox. Blue boxes represent independent processing nodes, each running in a separate thread. Nodes connect via queues transmitting image data, calibration information, or trigger signals.

The pattern generation node produces and sends an image to the image display node, which communicates with projection devices connected to the system locally or remotely. Sending a trigger signal to the camera capture node guarantees synchronized image capture. The capture node ensures the simultaneous capture of images from all connected camera devices. Its output goes to a trigger-generation node that then passes the output to the image-processing node. At the same time, a trigger signal goes to the synchronization node, triggering the display and acquisition of the next image from the pattern-generation node.

This node-based architecture allows

- easy incorporation of new algorithms and methods as nodes,
- easy comparison of different algorithms without changing the overall graph configuration, and
- selection of specific applications from a variety of nodes, tailored to the specific configuration.
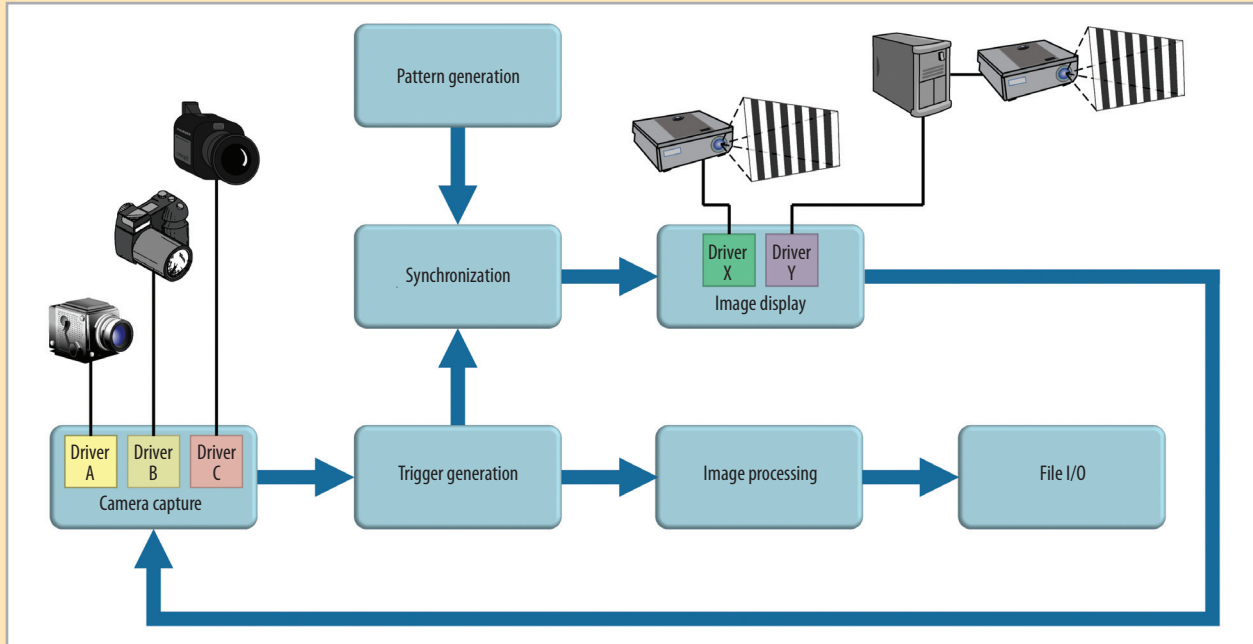
Figure 5. Example of a node-based structured-light-scanning application.

Our toolbox also provides support for feedback to detect hardware or alignment failures.

Techniques and algorithms

Thus far, we have implemented various AR techniques and algorithms that can be used in applications based on the ProCams toolbox.

Robust calibration. ProCams systems can exploit many algorithms and methods from the field of geometric computer vision. Of particular importance is the calibration of pinhole devices, which aim to recover internal (lens) and external (pose) parameters. Lens parameters include focal length, principal point, and lens distortion. Pose parameters include orientation and location. We have implemented a variety of calibration techniques, including stereo and multidevice calibration[5,6] from planar patterns as well as light-emitting diodes. The accuracy of the initial calibration for the cameras and projectors used in a particular configuration can typically be further improved. This can be achieved using an additional global optimization called *bundle adjustment*.[7] This well-known technique simultaneously optimizes the intrinsic and extrinsic parameters as well as the reconstructed 3D points for all devices, with the goal of reducing the overall reprojection errors. Using bundle adjustment, average reprojection errors of less than one pixel can be achieved for the configurations we have considered thus far.

**Structured light scanning and dense correspondences.** Parametric and nonparametric registration methods rely on generating correspondences among the various devices. Structured light scanning has proven effective and accurate, particularly on Lambertian surfaces.[8]

One example of structured light scanning involves presenting increasingly finer binary patterns. By acquiring each pattern with cameras and analyzing the images, an application using the ProCams toolbox can determine a dense set of correspondences between the camera and projector pixels. Combining these patterns with horizontally and vertically shifting line projections enables a precise, subpixel-accurate correspondence calculation.

ProCams applications can also exploit correspondences between a single camera and multiple projectors to determine the correspondences between the individual projectors directly. Doing so is important to generate smooth intensity transitions within the overlapping projections.

3D reconstruction. Given dense correspondences between calibrated cameras and projectors, application designers can use triangulation techniques developed in computer vision to perform a 3D reconstruction of the geometry in the scene.[5,9] Applied afterward, mesh smoothing reduces the noise of surfaces reconstructed from the individual 3D points.

Projector registration. Given a set of reconstructed 3D points with known 2D projector pixel correspondences, application designers can use standard techniques to calibrate the internal and external parameters for the projectors.[5,9] Calibration yields information about how each
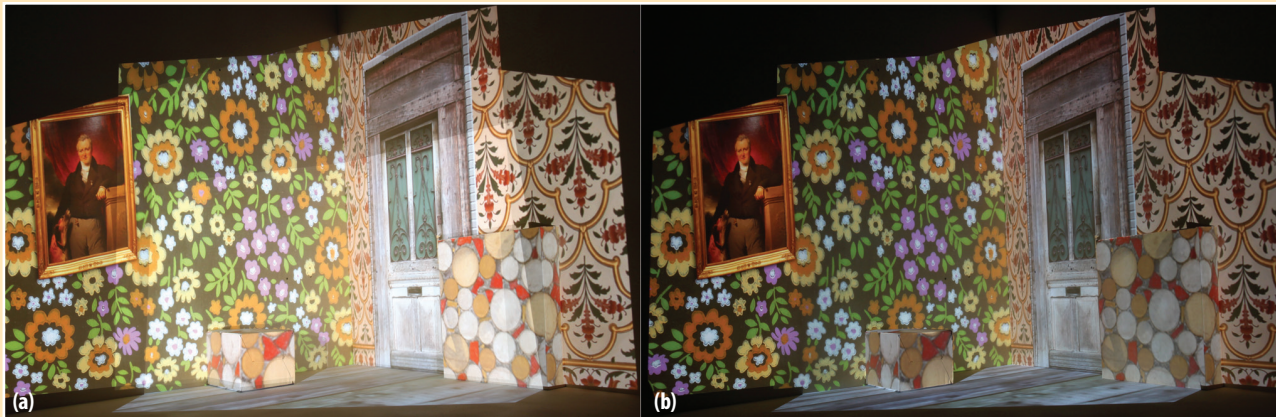
Figure 6. Example scene (a) without luminance compensation and (b) with luminance compensation, generating a perceptually seamless multiprojection surface.

projector relates to other projectors and the geometry. In other words, calibration registers the projectors with respect to the scene, enhancing the final (re-)projection and augmentation.

Luminance compensation. Scenes can require luminance compensation for several reasons:

- when pixels of multiple projectors illuminate the same 3D geometry location, observers will see an increased brightness at that location;
- projection onto (arbitrary) geometry, even with a single projector, can result in nonuniform brightness across the surface; and
- in certain configurations, the projections can cast shadows onto the scene.

Using per-pixel brightness reduction techniques, where the scene geometry and projection configuration determine the required reduction, can compensate for such brightness differences. Figure 6 shows a scene without and with luminance compensation.

We have implemented screen space as well as object space compensation techniques. To avoid the visibility of small misalignments, our compensation techniques exploit the cross-blending of projections whenever possible.[10]

Warping and rendering. Projecting registered content onto a scene can provide reality augmentation. Given the dense correspondences determined earlier, it is possible to prewarp the content and project it onto the real scene.

Suppose that the external parameters for the projectors and a precise 3D model of the scene are also available. If so, it is possible to produce a rendering of a colored and textured virtual scene from the points of view of multiple projectors. We can then project these multiple point-of-view scenes onto the physical scene. Doing so augments a static scene, giving it a dynamic appearance.

Modular design and standards

Projection-based theme park installations require various hardware and software components, often provided by different vendors. This situation makes it necessary to generate interfaces between third-party components and the ProCams toolbox. Building a maintainable system also means making sure that the hardware and software will have a life span far beyond the implementation phase of a product, an important consideration when an entire system must work for 10 to 30 years or more. The modular toolbox makes it easy to add new types of data converters and I/O communications whenever required and helps safeguard the system against future changes.

Currently, most vendors providing projection software tools use their own proprietary data formats for blending maps and image-warping definitions, which makes integrating new tools or replacing specific ones cumbersome because doing so requires additional converters. To simplify this process, we are currently working on a MultiProjector Auto-Calibration Standard (MPACS), providing a comprehensive definition for more convenient data exchange. We are coordinating the development of this MPACS through the Video Electronics Standards Association.

Manual alignment adjustments

Although the ProCams toolbox can generate accurate projector and camera calibrations, local misalignments can occur because of, for example, projection drift or the mismatch between virtual and physical models. For adjusting local misalignment errors, manual adjustment can be faster than a full, automatic recalibration.

Our ProCams toolbox supports manual adjustments by exploiting 3D calibration data and geometry to smoothly warp the projection locally in misaligned areas. The user carries out this operation on the image plane of one projector. User-defined constraints, such as the "pinning" of an

area to prevent its distortion, and local adjustments then go to all projectors. We are extending this tool to impose further constraints, such as depth discontinuities or edges, to semiautomatically guide users.

FUTURE DIRECTIONS

Projection-based AR is a powerful tool for enhancing and energizing theme parks and other entertainment environments. Continuing to expand the utility and encourage the continued adoption of this technology in Disney parks will require addressing several key challenges.

Complex shapes

Current projection-based AR techniques are best suited for regular mathematical surfaces—planes, cylinders, and spheres—and fairly simple, regular 3D shapes, such as buildings, furniture, and automobile exteriors. Although it is possible to project onto more complex shapes, such as rocks, figures, or character faces, surfaces must be fairly smooth and continuous for this to work well. Complex organic shapes with lots of discontinuities, complex surface angles, and self-shadowing—for example, trees or highly caricatured faces—are much more challenging. It is often difficult to cover the entire surface with projected imagery because of discontinuities and shadowing, and pixels are often unevenly distributed across the surface as a result of oblique projection angles.

Opportunities for research and improvement exist in areas such as

- automatic projector placement that optimizes for even pixel distribution and reduced grazing angles,
- improved techniques for per-pixel registration and geometric alignment,
- luminance compensation to adjust for varying levels of projector overlap, and
- advanced optical-system design to allow for complex projector and projection surface configurations.

Advances in these areas are critical to the successful application of projection-based AR techniques in the complex "real world" environments found in Disney theme parks.

Dynamic content

The situation becomes even more challenging when the projected surfaces are in motion or changing shape—for example, trees blowing in the wind or animated figures articulating or emoting. Even something as simple as a moving door can be a challenging projection surface.

If the projection surface's relative motion is known a priori, it is possible to adjust the media to compensate for the motion and synchronize it with the movement. In the real world, however, mechanical devices rarely perform perfectly 100 percent of the time: doors might open slowly or late, and animated figures—especially those driven using pneumatic or hydraulic systems—might move imprecisely. Furthermore, our goal is to develop a toolbox that provides a suite of methods and algorithms to design and support new projection-based AR installations. This requires dynamic tracking of both guests and dynamic objects. Real-time systems like this suffer from latency, which makes it difficult to maintain perfect alignment of the projected imagery.

In addition, displays increasingly use *eye-point-corrected media*: imagery rendered from the guest's current perspective. To work well, eye-point-corrected media requires a precise understanding of the current relationship between the guest's point of view and the projection surface, a precision that is difficult to attain when both the guest and the projection surface are moving.

> **Theme park environments and figures are increasingly interactive, reacting to the presence or actions of guests.**

In most cases, an important component of the solution is the existence of a low-latency, high-precision tracking system to track both guest position and dynamic objects in the environment. This has been, and continues to be, an open and challenging research problem. The need for these systems to be robust, reliable, and easy to maintain compounds the difficulty when we use them in Disney theme parks.

Real-time masking

A subproblem of projecting on dynamic objects is the challenge of *real-time masking*—the ability to selectively project or not project on parts of a scene. Masking is necessary, for example, if an actor must enter a scene augmented by projected imagery. Imagery intended for the scene should not project onto the actor; conversely, if we are also augmenting the actor, the actor's imagery should not project onto the scene.

Although tracking dynamic objects can address this problem, full 3D object tracking might be unnecessary. An understanding of the object's silhouette might suffice to generate a dynamic mask for projected media.[11] The projection-based application can then use this mask to project media only where it should.

Generating and displaying this mask has many of the same issues as does dynamic object tracking, primarily latency. Reducing end-to-end latency is important to ensure the mask does not lag behind the object it is intended to mask.

Figure 7. Concept illustration for the new Goofy's Paint 'n' Play House attraction at Tokyo Disneyland.

Projection-based AR continues to be an important tool in Imagineering's tool chest, helping to make buildings come alive in theme park entertainment shows. We also augment and animate figures in attractions with projected media, and many new attractions currently under construction or development include projection-based AR elements. For example, Goofy's Paint 'n' Play House, scheduled to open at Tokyo Disneyland in autumn 2012, is an interactive environment in which guests work together to help Goofy redecorate his house. As Figure 7 shows, guests will use interactive "paint applicators" to transform Goofy's living room into different themed looks—beach, jungle, or outer space—by painting on the walls, floor, and furniture with projection-based AR techniques. **C**

References

1. R. Raskar et al., "Shader Lamps: Animating Real Objects with Image Based Illumination," *Proc. 12th Eurographics Workshop on Rendering Techniques* (EGWR 01), Springer, 2001, pp. 89-102.
2. O. Bimber and R. Raskar, *Spatial Augmented Reality: Merging Real and Virtual Worlds*, A.K. Peters, 2005.
3. O. Bimber and D. Iwai, "Superimposing Dynamic Range," *ACM Trans. Graphics*, Dec. 2008, p. 150.
4. O. Bimber et al., "The Visual Computing of Projector-Camera Systems," *Computer Graphics Forum*, vol. 27, no. 8, 2008, pp. 2219-2245.
5. R.I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge Univ. Press, 2004.
6. T. Svoboda, D. Martinec, and T. Pajdla, "A Convenient Multi-Camera Self-Calibration for Virtual Environments," *Presence: Teleoperators and Virtual Environments*, Aug. 2005, pp. 407-422.
7. M. Lourakis and A. Argyros, "SBA: A Software Package for Generic Sparse Bundle Adjustment," *ACM Trans. Mathematical Software,* Mar. 2009, pp. 1-30.
8. J. Salvi et al., "A State of the Art in Structured Light Patterns for Surface Profilometry," *Pattern Recognition*, Aug. 2010, pp. 2666-2680.
9. R. Szeliski, *Computer Vision: Algorithms and Applications*, Springer, 2011.
10. A. Majumder and M.S. Brown, *Practical Multi-Projector Display Design*, A.K. Peters, 2007.
11. T.J. Cham, "Shadow Elimination and Occluder Light Suppression for Multi-Projector Displays," *Proc. IEEE Conf. Computer Vision and Pattern Recognition* (CVPR 03), IEEE, 2003, pp. 513-520.

Mark Mine *is the director of the Creative Technology Group at Walt Disney Imagineering. His research interests include theme park design, virtual reality, and interaction techniques. Mine received a PhD in computer science from the University of North Carolina at Chapel Hill. He is a member of the Themed Entertainment Association, the IEEE Computer Society, and ACM. Contact him at mark.mine@disney.com.*

David Rose *is the director of Software Tools Development at Walt Disney Imagineering R&D. His interests include real-time 3D graphics and themed entertainment. Rose received a BA in English from Florida State University and is a member of the Themed Entertainment Association. Contact him at david.r.rose@disney.com.*

Bei Yang *is a technical concept designer at Walt Disney Imagineering. His research interests include computer graphics, art, and visual communication. Yang received an MS in entertainment technology from Carnegie Mellon University and is a member of the Themed Entertainment Association. Contact him at bei.yang@disney.com.*

Jeroen van Baar *is a research scientist at Disney Research Zürich. His research interests include computer graphics and computer vision, primarily the intersection of those fields. Van Baar received an MS in computer science from Delft University of Technology, the Netherlands. He is a member of ACM. Contact him at jeroen@disneyresearch.com.*

Anselm Grundhöfer *is a postdoctoral researcher at Disney Research Zürich, where he works in the fields of projector-camera systems and computer vision. Grundhöfer received a PhD in engineering from Bauhaus University, Weimar, Germany. Contact him at anselm@disneyresearch.com.*