Experiences in Extemporaneous Incorporation of Real Objects in Immersive Virtual Environments

Benjamin Lok
University of Florida
lok@cise.ufl.edu

Samir Naik
Disney Corporation
Samir.D.Naik@disney.com

Mary Whitton, Frederick P. Brooks Jr.
University of North Carolina at Chapel Hill
whitton@cs.unc.edu, brooks@cs.unc.edu

Abstract

We present our experiences in developing, and user responses in experimenting with, virtual environments that can dynamically incorporate (visualize and interact with) real objects. We have found hybrid environments can enable natural interaction, extremely rapid development, and an effective platform for conducting tasks with a substantial manual component. We hope this article will encourage more research into incorporating real objects into virtual environments. This would improve interaction, enhance effectiveness, and increase the applicability of virtual environments to a broader range of tasks.

1. Introduction

Most virtual environments (VEs) are populated with only virtual objects. This seems like an obvious statement, but it has substantial implications for user interaction. While there are many advantages to having non-real objects, such as cost, logistics, and participant risk, there can be significant disadvantages for interaction, such as lack of haptic feedback, motion constraints, tracking, and affordance matching. So the more virtual objects introduced, the more of a benefit VEs can provide, yet the poorer the quality of interaction. Has this lack of high-fidelity interaction reduced VE's applicability to a wider range of tasks? System developers looking to provide real object interfaces have to spend significant development time modeling, tracking, and simulating, *each* real object that will be included in the VE. This has led to relatively few real objects in most immersive VEs.

We examine a more flexible approach to getting objects into an environment. We developed a hybrid environment (HE) system that uses real-time object reconstruction from multiple cameras inputs. The participant can simply pick up a real object, without prior modeling and tracking, and the system incorporates it in the VE. This allows *extemporaneous* incorporation of real objects, both visually, and as an active participant in simulations. While the algorithms and formal studies have been presented [1,2,3], we believe that our experiences in rapid development of HEs and anecdotes from users would prove of interest to those who grapple with VE interaction challenges.

2. Previous Work

There are many approaches to obtaining shape, appearance, and motion information of a specific real object. Making measurements and then using a modeling package is laborious for static objects, and near impossible for capturing all the degrees of freedom of dynamic objects. Precise models of real objects are not necessary for some applications. Approximations, such as the *visual hull*, are useful substitutes to achieve real-time performance [4].

Ideally, a participant would use his hands, body, and tools to manipulate objects in the environment. As a step toward this goal, some VE systems provide tracked, instrumented input devices. Commercial devices include tracked articulated gloves with gesture recognition or buttons (Immersion's Cyberglove), mice (Ascension Technology's 6D Mouse), or joysticks (Fakespace's NeoWand).

Another approach is to engineer a device for a specific type of interaction to improve affordance. For example, having a toy spider registered with a virtual spider [5] allows the participant to touch a physical object where they see a virtual representation. However, this specialized engineering can be time-consuming and results are often application specific.

3. Reconstruction

We use a HE system that performs image-based object reconstruction to generate real-time virtual representations, *avatars*, of real objects. The system has multiple (four in our current setup), fixed-position cameras that view the scene. At each frame, the visual hulls of objects of interest (e.g. participant, tools, and parts) are computed.

By regenerating the virtual representations at each frame, the system can handle highly dynamic and deformable objects (e.g. clothing, tools, and hands). In the head mounted display, the participant sees real object avatars visually incorporated into the VE. Further, the participant handles and feels the real objects while interacting with virtual objects through specialized collision-detection and collision-response algorithms. This enables the real object avatars to *naturally* affect simulations such as particle systems, cloth simulations, and rigid-body dynamics. We believe this is a natural VE interface.

4. Designing a VE

4.1 User Study

For every potential device and interaction possible, developers must invest substantial development effort in coding, modeling, and tracking. For example, we developed a 'virtual objects only' environment for a user study where participants picked up and manipulate virtual cubes. This required developing software to incorporate tracked pinch gloves to control an avatar, interaction mechanisms between all virtual objects, and models for the avatar and the blocks. Every possible input, action, and model for each object had to be defined, developed, and implemented. This took experienced VE developers three weeks to create. The resulting system established very specific methods in which the participant could interact with virtual objects. Any changes required substantial system modifications.

Figure 1. Participant parts virtual curtains naturally with their hands.

In contrast, the conditions with the HE system to allow the user to manipulate real blocks were developed within hours. The incorporation of real blocks and the participant's avatars did not require writing any code. The reconstruction system obviates prior modeling and incorporating additional trackers. Further, the flexibility of using real objects and the reconstruction system enabled minor changes to be made without much code rework.

4.2 Physical simulations

Beyond just seeing real objects within a virtual world, the user needs to *interact* with the virtual world with the real objects. We used the VE from the popular UNC Pit Walking Experiment. In this VE, one of the rooms has a window with curtains, a basic cloth simulation. What if the participant could naturally interact with the curtains? We replaced the collision detection function call from the traditional collision detection system (Swift++) to a new function that detected collision between real and virtual objects. This allowed the participant to part the curtains with their hands or any other objects that was nearby [Figure 1]. We incorporated real objects into VEs with particle systems, a bouncing ball [Figure 2], and shadows and lighting. The development time required to add the capability of real objects to be active participants in each VE was less than one day.

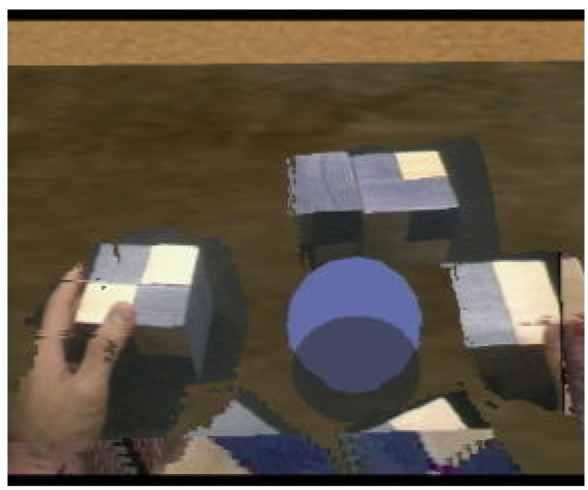

Figure 2. Participant interacts with a rolling virtual ball with various real objects.

4.3 NASA layout evaluation

In [3] we described a case study involving NASA payload designers using HEs to evaluate an abstracted layout task. They were to examine the spacing between virtual payload models using real parts and tools (mocked up with a PVC pipe, power cables, and tongs). This task would typically be difficult to simulate due to the shape complexity (deformability and degrees of freedom) of the objects. The result of the case study was that spacing errors that might slip through in traditional specification and validation methods, could be detected using HEs.

During the case study, NASA engineers had initially validated a layout given a fictional task description (provided using the same validation methods in current use). During the task, they found that they could not complete the task because a cable was too flexible to be inserted into an outlet at the end of a tube. They commented that what they needed was a tool to complete the insertion. We handed them a set of tongs that acted as a set of pliers that allowed easy completion of the task. Unfortunately, this increased the required spacing. This error was found using real tools and parts with virtual models of components.

With the HEs no *a prior* development was needed to try out multiple tools and parts. Total development time was two days, which involved writing a parser for the CAD models that NASA provided.

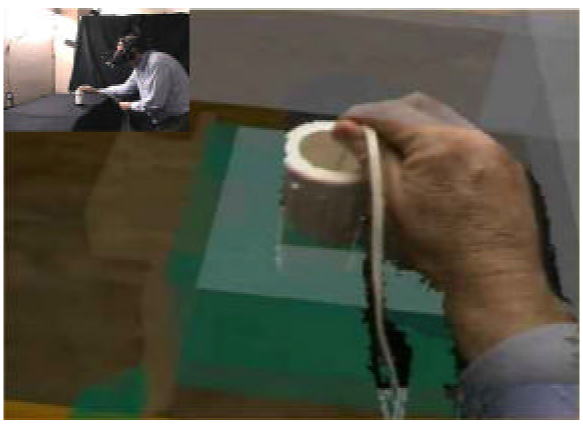

Figure 1. A NASA payload engineer evaluates a spatial layout task using real tools and parts.

5. User Reaction

We review some anecdotal evidence of the importance of real objects for manual tasks.

5.1 Surprise

Many veteran VE users are typically surprised that they can interact with virtual objects by doing the action that most naturally comes to mind. A common discourse when the participant is shown the curtains and told to part them was “With what?” “Try using your hands” “But, you aren’t tracking them”

When they look down and see their hands in the environment, they are surprised. They then immediately reach out to play with the curtains. Afterwards, they say the only thing unusual about the experience was that they *expected* to feel the curtains against their hands as they parted them!

5.2 Extemporaneous Exploration

Following the case study, a NASA engineer noted one solution to the discovered spacing error would be to modify or use a tool that could fit within an allocated space. This highlights another advantage of using real-time reconstruction systems as multiple tools can been tested to see which are capable of fitting within a given layout tolerances. Instead of having to model, track, and code for each possible tool, different tools can be simply handed to the user and be incorporated and tested with the layout dynamically.

6. Advantages

6.1 Minimal Interface Training

We conducted a case study to compare task performance in HE versus traditional VEs [2]. In this study, participants performed a timed block manipulation task in the real world and then in either a HE or VE. As this is the first exposure of many to VR, a period of familiarization with the virtual world and especially the methods of interaction (pinch gloves in the purely virtual condition) were required. We had limited each participant to only fifteen minutes of time in VR to reduce simulator sickness effects. The time used for interface training reduced the time allocated to carrying out the experiment and the number of experiment test patterns. In contrast, the HE conditions used identical blocks to the ones used in the real world test, and the participants remarked they recognized the blocks and immediately understood the task at hand. They also reported (via questionnaires) being more confident in their ability to perform the task and took significantly less time to be acclimated to the VE (than virtual only participants).

6.2 Improved Interaction

Since the system is trying to determine which parts of the working volume is inhabited by real objects, objects that would be typically difficult to model and track could be handled. For example, the participant’s hand, cables, and tools would be difficult, impossible, or time consuming to model and track.

Participants remarked that having the real objects improved affordance (they knew how to use the real objects), and their performance was aided by motion constraints, such as screw threads, and the outlet plug.

Finally, participants felt an anecdotal increased level of immersion due to the familiarity of handling the real blocks, tools, and parts. NASA participants were observed to make concerted efforts to avoid touching the virtual model with their hands and tools. Upon being told this, one participant responded, "That was flight hardware… you don't touch flight hardware."

7. Disadvantages

7.1 Line-of-sight

Camera based approaches to object reconstruction will have line of sight issues. If the task requires many objects, or some that would be obstructed from the camera's view, then the error and uncertainty in the reconstruction will be substantial.

7.2 Model approximation

To achieve real-time performance, we chose to approximate the real objects shape. This approximation will not correctly handle certain shapes, such as a cup, change per frame, and always contain error. Depending on the task (e.g. medical or very precise engineering) this can be a serious drawback.

7.3 No real object properties

During a collision, there is no additional information about objects in intersection. Thus virtual objects can not respond in object specific ways (for example, a virtual ball bouncing differently based on the real object's surface properties). The 'treat everything as a volume', while great for extemporaneous inclusion, limits applications that require additional collision information.

7.4 Haptic feedback from virtual objects

When virtual and real objects collide, the real object is not affected. This limits very complex interactions with the virtual objects. For example, trying to use a virtual object as a brace for a real object or trying to part virtual cables to route a real part requires haptic feedback for effective experiences is not easily accomplished

8. Future Work and Conclusions

We believe there are substantial gains to be made if we spend some time (on the order of a few minutes) with non-real-time reconstruction algorithms, and then incorporate and track the resulting higher quality models. We have done some work with simple marker based tracking [6], and look to combine it with laser scanned virtual representations using a CyberWare 3D scanner. Thus a part or a tool can be scanned, markers affixed to the object, and then used in the HE. Another potential approach is to use volumetric 'matching' algorithms instead of marker based tracking. The system would match a high quality model against the real object's appearance in the live camera images

Extemporaneous inclusion of real objects enhances interaction with the VE by allowing the user to naturally interact with the virtual world. This is vital to extending VE applicability and effectiveness.

9. Acknowledgements

We thank Danette Allen, Marilee Roell, Lloyd Spangler, Kerry Gough, and Glenn Hines from NASA Langley Research Center for their insight. We thank Henry Fuchs, Herman Towles, David Harrison, Kurtis Keller, John Thomas, and the UNC Effective Virtual Environments group. This work was supported by the LINK Simulation and Training Foundation, ONR Virte Program, NIH National Institute of Biomedical Imaging and Bioengineering, and NSF ITR BioGeometry Project.

10. References

[1] Lok, B. "Online Model Reconstruction for Interactive Virtual Environments." In *Proceedings 2001 ACM Symposium on Interactive 3-D Graphics*, 69-72, 248.

[2] Lok, B., Naik, S., Whitton, M., Brooks. F. "Incorporating Dynamic Real Objects into Virtual Environments," *Proceedings of ACM SIGGRAPH 2003*, 701.

[3] Lok, B., Naik, S., Whitton, M., Brooks. F. "Effects of Interaction Modality and Avatar Fidelity on Task Performance and Sense of Presence in Virtual Environments," To appear in *Journal on Presence: Teleoperators and Virtual Environments*.

[4] Laurentini, A. 1994. The Visual Hull Concept for Silhouette-Based Image Understanding. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 16, 2, 150-162.

[5] Hoffman, H., Carlin, A. and Weghorst, S. 1997. Virtual Reality and Tactile Augmentation in the Treatment of Spider Phobia. *Medicine Meets Virtual Reality 5*.

[6] Jackson, J., Lok, B., Kim, J. Xiao, D., Hodges, L., Shin, M. "Straps: A Simple Method for Placing Dynamic Avatars in a Immersive Virtual Environment." Future Computing Lab Tech Report FCL-01-2004. Dept of Computer Science, Univ. of North Carolina at Charlotte.