Mirror Worlds: Experimenting with Heterogeneous AR

Alex Hill, Evan Barba, Blair MacIntyre
School of Interactive Computing
Georgia Institute of Technology
{ahill,ebarba3,blair}@gatech.edu

Maribeth Gandy, Brian Davidson
Interactive Media and Technology Center
Georgia Institute of Technology
{maribeth.gandy,briand}@imtc.gatech.edu

***Abstract*—Until recently, most content on the Internet has not been explicitly tied to specific people, places or things. However, content is increasingly being geo-coded and semantically labeled, making explicit connections between the physical world around us and the virtual world in cyberspace. Most augmented reality systems simulate a portion of the physical world, for the purposes of rendering a hybrid scene around the user. We have been experimenting with approaches to terra-scale, heterogeneous augmented reality *mirror worlds*, to unify these two worlds. Our focus has been on the authoring and user-experience, for example allowing ad-hoc transition between augmented and virtual reality interactions for multiple co-present users. This form of ubiquitous virtual reality raises several research questions involving the functional requirements, user affordances and relevant system architectures for these mirror worlds. In this paper, we describe our experiments with two mirror world systems and some lessons learned about the limitations of deploying these systems using massively multiplayer and dedicated game engine technologies.**

Keywords-Ubiquitous Computing; Augmented Reality; Virtual Reality; World Wide Web; 3D User Interfaces

I. INTRODUCTION

Most of the content on the Internet today is information-oriented (i.e. documents, images and e-mail), although this information is increasingly tied to the physical world (e.g. geo-coding of images and social media, aggregating information about people, inferring relationships to objects and places). Alongside this semantically meaningful virtual information is a growing collection of knowledge about physical phenomena, such as locations of people, buildings and vehicles along with the descriptions of all manner of objects bought and sold. With the growing popularity of social media and the power of mobile devices, many see cyberspace becoming an almost real-time simulation of the world we live in, with each physical object having a virtual antecedent. We have been experimenting for the past five years with the implications of such a mirror world on virtual and augmented reality system design and implementation.

Mirror worlds are a structuring of the virtual world that aligns closely with the physical world, allowing both localized in-situ users and non-local users to interact through a heterogeneous mixture of augmented and virtual reality. Our interest is specifically in mirror world systems having three main qualities; a) inherent massively multiuser online (MMO) scale and persistence, b) technology and affordances to support facile transitions between augmented reality (AR) and virtual reality (VR) on heterogeneous platforms, and c) data formats that support individual authoring and distribution of content much like that of the modern web. We have been experimenting with a mix of technologies and affordances that satisfies these broad requirements.

In this paper we will discuss our handheld AR Second Life (SL) system that was used to explore the interfaces for and the practicalities of delivering Mirror World experiences on mobile devices. Then we will describe the more flexible successor system we developed using a combination of the Unity 3D game engine and Sun Microsystems DarkStar game server. Lastly, we will discuss some lessons we learned from these two architectures and describe their influence on subsequent projects.

II. RELATED WORK

The mirror world concept of bridging the real and virtual, allowing a variety of interfaces to connect to common spaces, has been explored in many related projects, which we cannot fully summarize here. For example, in 2002, Piekarski & Bruce Thomas created an AR version of the popular Quake game on their campus [5]. Similarly, "Human Pacman" was a campus-scale Pacman game, which allowed remote players to observe and participate remotely in the live AR game play [2].

A variety of mixed reality projects have allowed collaboration between mobile AR users and remote visitors via full virtual reality systems, browser based 3D environments, mixed reality systems and ubiquitous computing installations. Stafford et al developed a method of presenting a "god-like" view of a physical scene via a tabletop projection that enabled indoor users to communicate with outdoor AR users [3]. Most recently, Murphy et al Image Space allows users to share location aware media via a web based interface as well as a mobile AR system [4].

III. HANDHELD AR SECOND LIFE

Our prior experience with developing an AR Second Life client (similar to others [5]) highlighted the significant difference in affordances available to VR users over those using an AR HMD [6]. While SL users in a VR mode could navigate and manipulate their camera view at will, AR users could only view SL avatars and content from their first-person perspective. Our first SL system used high-performance 6D tracking hardware to track a user HMD over a wide indoor area. Wanting to experiment with allowing AR users to utilize alternate viewpoints, we developed a handheld version of the SL client using vision-based marker tracking. This provided the user with a “god-like” view of a section of the world, similar to Stafford et al [3].

We used two hardware setups; a Wacom Cintiq tablet display connected to a SL client and a Sony Vaio Micro PC handheld (Figure 1a). The Sony Vaio could not run the SL client at significant frame rates but was able to do marker tracking using the ARToolKitPlus library. We experimented with two thin-client setups on the Vaio; one doing marker tracking on the client with rendering and compositing on a server and a second setup streaming the video feed to a server for both marker tracking and scene compositing.

To provide the “god-like” view of the AR scene we tracked against tabletop-sized fiducial markers with small physical props (i.e. books and Legos) as stand-ins for elements of the physical environment (Figure 1b). We also experimented with two different navigation metaphors both of which were related to “World in Miniature” (WIM) techniques from VR research [7]. In the camera fixed mode, arrow keys on the device changed the position of the user relative to the AR scene (i.e. the character could walk around a subset of the world mapped to the fiducial). The camera following mode allowed the user to unpin the location; as the avatar walked, the world moved underneath it, allowing the user to navigate to and re-pin at the next location, mapping a new area to the fiducial.

The MMO-scale of SL and the accessible authoring of 3D content proved useful when creating a practical mirror world platform. Ultimately, however, the difficulties of working within the rules of the system (e.g. forced updates of the client) and the limitations and tediousness of authoring inspired us to move to a new game engine platform that would provide a mature content pipeline and multi-platform support.

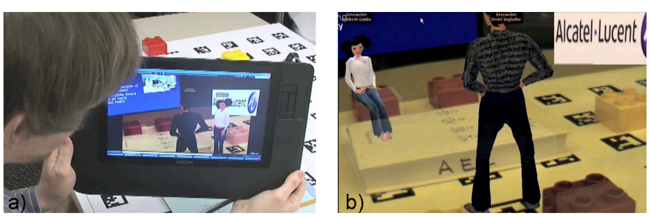

Figure 1. Handheld AR Second Life using a) marker-based camera control and b) props to recreate physical and virtual scenes.

IV. Unity 3D Mirror Worlds

Based on our experiences creating small-scale mirror worlds with AR Second Life, we wanted to experiment with a larger scale mirror world. We were particularly interested in creating a system that could scale more readily to a diverse set of devices and interfaces. To this end, we decided to implement a new prototype using Unity 3D, a high quality game engine that can target desktop and handheld platforms[1].

[1] http://www.unity3d.com

A. *Desktop and Backpack Mirror World*

We created our mirror world on the Georgia Tech campus, leveraging models created by the Imagine Lab in the School of Architecture. The multi-user environment allowed a heterogeneous mix of desktop VR and in situ AR users.

Our goal with this platform was to explore the advanced uses of rich media, served from the cloud, which we imagined would become a reality in the coming years. Our data services were not live; we created pre-loaded content to mimic web data sources. To allow for elegant registration of content and natural interfaces with the environment we used “invisible” versions of the models in the AR view, which provided occlusion between the authored content and the physical world and allowed for selection and interaction by the user. There was 2D web content tightly registered to physical buildings and the ground. Avatars could appear from behind real buildings, and users could touch on physical structures to create and/or interact with content. As a sample application of the mirror world technologies, we allowed users to collaboratively markup the surrounding buildings with graffiti (Figure 2b).

The mobile AR backpack system we employed consists of a Macbook Pro laptop, a MotionNode[2] orientation sensor, a Garmin GPS and a Wacom Cintiq interactive tablet display. Our commercial differential GPS, while accurate to ~1 meter (best case), still resulted in registration errors in the field. To compensate for this we provided a clutching mechanism attached to a button on the tablet display to adjust registration. The Unity 3D system creates both standalone executables and browser-based applications for desktop clients.

Figure 2. Avatar representation of a) remote VR client visible within backpack system AR view and b) in corresponding third person VR view.

B. *iPhone Client and GeoSpots*

The most significant limitation of our initial Unity mirror world system was a lack of state persistence between application executions and the inability to support a large number of simultaneous participants. To solve both problems, we used Sun Microsystems DarkStar server (renamed to RedDwarf[3]) to create a mirror world server; DarkStar is a low latency application server for online games and virtual worlds that allowed us to easily support dozens of simultaneous clients.

A second problem we needed to address was the limitation of GPS and orientation sensor accuracy, not only

[2] http://www.motionnode.com
[3] http://www.reddwarfserver.org/

on our backpack system but also on the newer mobile phones we also hoped to use. (At this point, the iPhone3GS was available and could be targeted by Unity3D.) These commodity sensors in mobile phones are adequate for AR representations of distant objects but too inaccurate for aligning buildings and models near the user. Even when location coordinates are accurate, accelerometer-based orientation sensor noise in commodity mobile phones such as the iPhone 3GS results in content that "swims around" on top of the video. A goal of the mirror world project was to experiment with touch screen interactions that leveraged the underlying building models, and these poor quality sensors made this impractical. Even with current generation phones with gyroscopes, combined position plus orientation accuracy remains a problem and likely will for some time.

In order to compensate for both GPS and orientation accuracy, we shot panoramic photographs at selected locations and allowed the mobile user to indicate their presence at one of these so-called GeoSpots (a similar technique was used in the AR Moon Lander game [8]). Using this panoramic backdrop technique, users can stand at a pre-defined location in situ and view the static image of the scene by changing the phone orientation. Because the registration between background panorama and the surrounding AR content is fixed, we can accurately occlude synthetic content behind buildings and give predictable interactions with the underlying models (Figure 3a). In addition to graffiti painting, we added a second "application" where users could place and manipulate text annotations on buildings.

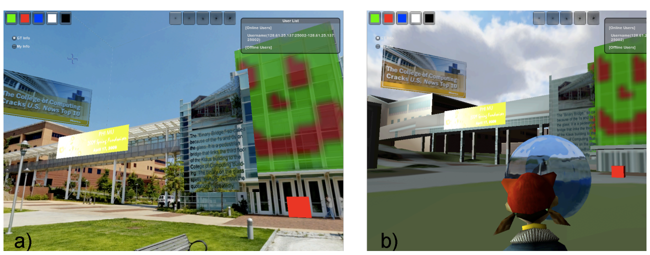

Figure 3. a) Underlying models constrain graffiti content in panorama view. b) AR bubbles facilitate transitions between AR and VR views.

To support transitioning between the VR and AR views, we placed "AR bubbles" at the designated GeoSpots in the world (Figure 3b). Using either arrow keys on the desktop and backpack-based systems or virtual gamepads on the iPhone 3GS, users can move off of their AR location and into the VR world (the VR clients are in this mode by default). Once in this VR mode, the user can travel to arbitrary locations and place annotations or graffiti on buildings. The user can also move into AR bubbles at other locations to enable the panorama backed AR interactions at those locations. This transition from AR view to VR view to panorama was easy to understand and provides a way to both move around in the local space as well as travel to more distance locations when out in the world. Although we opted to have users manually move into nearby GeoSpots through the VR interface, in recent work we have provided a map interface that allows the user to directly select and "go to" nearby GeoSpots [9].

V. Discussion and Future Work

Our efforts with Second Life, Unity 3D and later experiments we conducted with Sun Microsystems Wonderland encourage us that existing architectures can be appropriated to build these massively scalable mirror worlds [10]. And, our experiments support a belief that both fiducial markers and GeoSpots are useful tools that let AR users view and interact with mirror world content not directly tied to their current viewpoint. Combining AR access to these affordances along with VR interfaces serves the dual purpose of both expanding the audience for AR content and supporting the authoring of that content through a form of WYSIWYG interface.

Building massively scalable worlds that mirror the physical world is a considerable exercise and developing the tools to allow the facile generation of content by users is arguably beyond the scope of our work. In fact, if we believe our own predictions, the real mirror world is developing out on the web already and the systems to support it may be of less interest than the tools that organize, discover and render it. Of note is the problem that multiple authors placing overlapping content onto a single physical space presents for content discovery and visualization. As a result, we are now turning our attention to visualizing mirror world content through mashups of existing geo-referenced data and encouraging the creation of new geo-referenced AR content through an extended version of KML[4] and our iPhone client called Argon[5].

Acknowledgment

This work was supported by a grant from Alcatel Lucent of North America and equipment donated by Sun Microsystems.

References

[1] W. Piekarski and B. Thomas, ARQuake: The outdoor augmented reality gaming system, Communications of the ACM,2002 Vol 45. No 1, pp. 36-38.

[2] A. D. Cheok, K. H. Goh, F. Farbiz, W. Liu, Y. Li, S. W. Fong, X. Yang and S. L. Teo. Human pacman: a wide area socio-physical interactive entertainment system in mixed reality, Proceedings of CHI Extended Abstracts, 2004, pp. 779-780.

[3] A. Stafford, W. Piekarski and B. H. Thomas, Implementation of god-like interaction techniques for supporting collaboration between outdoor AR and indoor tabletop users, International Symposium on Mixed and Augmented Reality - ISMAR , 2006, pp. 165-172.

[4] D. J. Murphy, M. Kahari and Ville-Veikko Mattila, An augmented reality view on mirror world content, with Image Space, 2010 IEEE Virtual Reality Conference (VR).

[5] T. Kantonen, C. Woodward and N. Katz, Mixed reality in virtual world teleconferencing, Virtual Reality Conference (VR), 2010 IEEE, pp. 179-182.

[4] http://research.cc.gatech.edu/kharma

[5] http://www.argon.gatech.edu

[6] T. Lang, B. MacIntyre and I. J. Zugaza, Massively multiplayer online worlds as a platform for augmented reality experiences, IEEE Virtual Reality Conference, 2008, pp. 67-70.

[7] C. Wingrave, Y. Haciahmetoglu and D. Bowman, Overcoming world in miniature limitations by a scaled and scrolling WIM, Proceedings of IEEE Symposium on 3D User Interfaces, 2006, pp. 11-16.

[8] B. Avery, B. H. Thomas, J. Velikovsky, J. and W. Piekarski, Outdoor augmented reality gaming on five dollars a day, Proceedings of the Sixth Australasian Conference on User interface, v. 40, Newcastle, Australia, Jan. 30 – Feb. 03, 2005.

[9] Hill, A., MacIntyre, B., Gandy, M., Davidson, B. and Rouzati H. KHARMA: An open KML/HTML architecture for mobile augmented reality applications, Proceedings of 9th IEEE/ACM International Symposium on Mixed and Augmented Reality, Seoul, Korea, Oct 13-16, 2010.

[10] Hill, A., Bonner, M., MacIntyre, B., Herbert, S. and Britton, D. ClearSpace: Mixed reality virtual teamrooms, Proceedings of Human-Computer Interaction International 2011, July 9-14, Orlando, Fl, in press.