Editors: Lawrence Rosenblum and
Simon Julier

3D Modeling with the Tinmith Mobile Outdoor Augmented Reality System

Wayne Piekarski

University of South Australia

At the Wearable Computer Lab at the University of South Australia, we have been performing research into outdoor augmented reality (AR) systems for the last seven years. During this time the technology has vastly improved, resulting in more accurate systems that have better quality output. While tracking and registration are important issues in the area of AR, it's also important that we have suitable user interfaces that let people effectively view information and control the computer to get the desired output. Therefore, the Tinmith project explores the problem of interacting with a mobile AR system outdoors and the types of possible applications.

Working in an outdoor environment also imposes limitations not normally experienced indoors. Apart from having to be light enough for a person to carry and be self-powered, the equipment must work with interference from the harsh and uncontrollable outdoor environment. While we initially developed navigation systems with simple user interfaces, our main focus has been developing systems that support real-time modeling of 3D objects. This type of modeling is a difficult user-interface problem, so we chose this as a suitable goal to drive our research. To support this work, we have developed our own custom backpack systems that contain all the necessary electronics to run the software and provide the user with tracking and input devices. While some of the components can be sourced commercially, we must custom develop other parts to support our research. These challenges make working in an outdoor environment difficult, but it poses many interesting research problems.

1 Current version of our custom Tinmith backpack system.

Implementation

Figure 1 shows our current custom backpack system. While it's possible to buy off-the-shelf, wearable computers, these systems are generally power efficient and therefore don't contain high-end 3D graphics acceleration hardware. Also, these systems don't provide all the tracking devices and other components we require, and are therefore unsuitable for our research.

We have therefore constructed several backpack platforms over the years, which allow us to carry all the components required in a protective housing. To provide accurate 3D overlays, we use a Trimble Ag132 GPS unit that provides an accuracy of approximately 50 cm using only a satellite differential signal. For orientation sensing, we use an InterSense InertiaCube2 sensor that combines accelerometers, gyroscopes, and magnetometers. The backpack contains a large 100-watt-hour battery to power the system outdoors for over two hours. Rather than using legacy interfaces such as RS-232 for our components, we attempted to USB enable most of our peripherals so that power and data can be sent over one standard cable type, cutting down on the weight and space used by wiring. We use a standard Dell Inspiron 8100 laptop, which contains its own batteries, and an Nvidia graphics accelerator to support our 3D rendering. The head-mounted display worn by the user is from IO-Glasses, and a firewire camera captures the real world to support video overlay AR, and tracking of the user's hands. We use the Tinmith-evo5 software architecture to implement our applications to run on the backpack.[1] Our software architecture

Published by the IEEE Computer Society

- implements an object-oriented data flow design,
- is written in C++ for efficiency on mobile platforms,
- uses an object store based on the Unix file system model, and
- implements best practices from other VR and AR research systems.

Tinmith-evo5 allows us to rapidly adapt our applications to new ideas with minimal changes due to its flexibility.

```
13.0fps @ 80.8ms    18 Feb, 2004 19:51:45
S34° 48.635258' E138° 37.234064' LLH-WGS84
A63.5(+)-0.0=63.5
< No Clipboards >
```

```
(4) Tinmith-evo5 ready ...
(3) Camera is now immersive
(2) Tinmith display captured
(1) Tinmith display captured
```

2 Users use tracked gloves to manipulate 3D objects against working planes. The blue and green menu strips show commands that can be selected by pressing the appropriate finger against the thumb on the gloves. Fiducial markers on the thumbs are tracked by the camera to provide pointing information to the user-interface software.

User interface

Working with a computer while mobile imposes constraints over traditional interaction indoors. Since there are no fixed surfaces, devices like mice and keyboards are not practical. Moving around the environment freely prevents the use of input devices and trackers that require fixed infrastructure to operate. Another problem is that many devices are 2D only. When using AR outdoors the environment is 3D, so ideally we should try to take advantage of this.

Optical tracking is one of the only ways to perform tracking outdoors. Other tracking devices that rely on magnetic emissions, infrared, or mechanical linkages require equipment not practical on a backpack, or is interfered with by the environment. To control the system, we use gloves that contain fiducial markers on the thumbs and copper pads on all the fingers. The previously described head-worn video camera is used to detect the fiducial markers, so no extra infrastructure is required. The copper pads detect when fingers are pressed against the thumbs or the palm, and each of these maps against the menu options, shown in blue in Figure 2. Pressing the fingers into the palm picks the green menu options, which are used typically for approve and cancel operations.

An important feature of the user interface is that users perform pointing and command entry separately, so the user can perform both simultaneously without interference. Also, users don't need to hold their hands in the frame when only performing command entry, reducing fatigue because they can keep their hands in a comfortable pose.

One limitation of working outdoors is that objects are typically located at far distances. As objects become further away, users have trouble perceiving the size and distance of objects they are unfamiliar with. When performing 3D modeling outdoors, the user must accurately enter in the models that they desire, otherwise the results will be clearly wrong when viewed from a different location. If the user can't accurately estimate depth and size however, it appears as though modeling outdoors is not possible because of these conflicting problems.

In our previous work, we proposed the use of AR working planes as a way of providing accurate depth estimation.[2] We intentionally restrict the hand tracker so that it only provides a 2D cursor and project this cursor onto a clear working plane placed into the environment. The system can create the working plane relative to other objects in the environment, or project it from the body along the user's line of sight. The system can accurately create working planes because it doesn't require users to estimate depth, but instead relies on users to position their body from another viewpoint and from a perpendicular direction. The system then projects the 2D cursor against the working plane to resolve 3D depth.

AR working planes are based on similar techniques used in CAD systems, where 2D input devices are used to enter 3D geometry. Our system can create AR working planes relative to the world, relative to user location, or attached to the user's head. An important part of AR working planes is that they let users use their body and not just their hands to work with the environment, taking advantage of a human's ability to know where objects are placed around themselves. This allows a number of interesting operations, such as attaching an object to the user and moving it around by walking, and using the hands to pick up an object and manipulate it along the plane. Figure 2 shows how the hand with the overlaid cursor is used to pick up a 3D object. We support manipulation with one hand, but also scaling and rotation using two-handed input. Users can also manipulate objects using a cursor projected from the display's center, requiring no pointing with the user's hands.

Construction at a distance

The AR working planes techniques let users specify planes to manipulate objects against and to draw on, but this only provides a limited subset of the capabilities required for 3D modeling outdoors. We want to use the system to create 3D representations of objects already in the environment, such as trees and buildings.

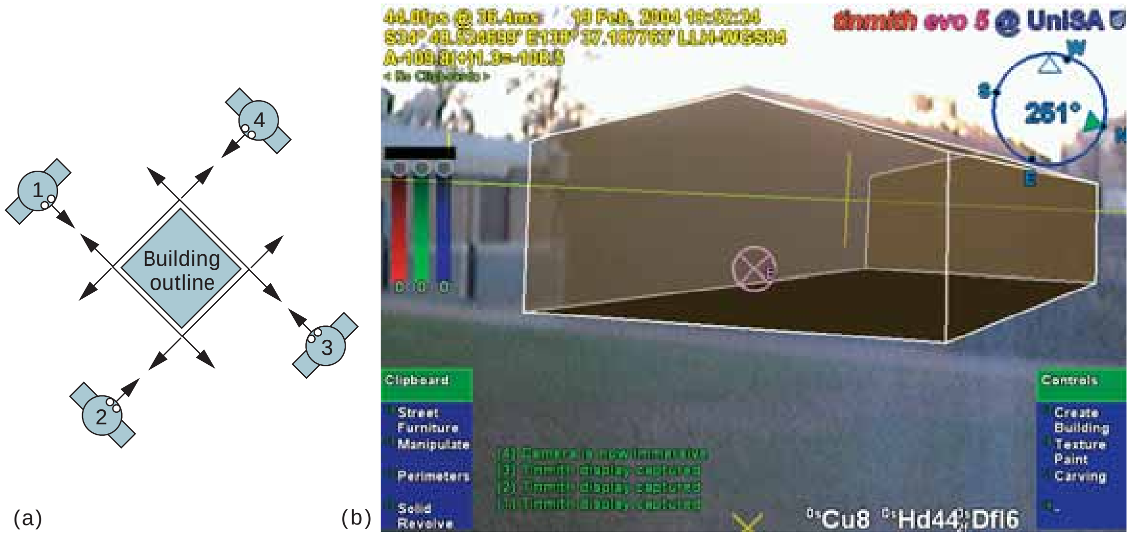

(4) Camera is now immersive
(3) Tinmith display captured
(2) Tinmith display captured
(1) Tinmith display captured

3 We used (a) the infinite planes technique to construct the building footprint and (b) the gloves to apply a sloped roof to accurately match the physical object.

We also would like to create 3D representations of objects that do not exist, so we can preview what they will look like before being constructed. While a laser scanner potentially could capture existing objects, the only way we can create a fictional object is indoors on a desktop CAD system. This is limiting because the designer can't see in real time how the object will look in the environment. We have therefore designed several techniques that we term *construction at a distance*,[3] which allow users to model the shape of large objects. Construction at a distance and AR working planes together help implement interesting modeling techniques suited for use outdoors.

The first technique we implemented is named the *infinite planes technique*. We based it on the principle used by people who pilot small boats offshore, where landmarks such as trees are aligned with hills or buildings to find locations of fishing spots. Using this concept, AR users walk around a building and align their head so they are looking along a wall. By knowing the AR user's position and orientation with various sensors, the system captures this information and draws a plane in the 3D environment. By intersecting these planes together, the system can form a solid object representing the building's shape within the tracking devices' limitations. Figure 3 illustrates this process from a top-down view and shows an example object overlaid on the real building it was based on. An important feature of this technique is that the user's perceptive capabilities don't impose any further errors on the 3D model.

We developed another technique called *orthogonal laser carving*. The user creates a shape in the environment using the previously described infinite planes technique, and then carves out more complex features along an AR working plane created against a wall surface. Figure 3 shows an example building created in this way; we carved out the pointed roof by marking out the corner points along a wall. An alternative, shown in Figure 4, is where a car is modeled by creating a box matching the approximate dimensions, and then it's carved away from multiple viewpoints until the shape matches the physical object. The user's hands mark out the corner points, and the user works around the shape to define the overall outline. An important feature is that the user can inspect the results in real time, and make any corrections immediately.

4 A user with tracked gloves carves out a solid box into the shape of an automobile.

One limitation of working from an immersive AR viewpoint is that it's difficult to see occluded objects, and distant objects will appear tiny on the display. We investigated the use of external viewpoints that let users leave their body and use head motions for orbital views, top-down perspectives, and flying to other locations. It's important that the user not become confused when in these modes, so we remove the AR video overlay and replace it with a sky box so the user understands the display is artificial. Because users can't easily see the real world, they typically don't move their body in this mode and use it only for situational awareness. This viewpoint is also useful when operating the system collaboratively; users with indoor displays can see what the mobile user is doing outdoors from an external viewpoint, such as shown in Figure 5.

Conclusions

While we've made considerable progress in the area, we still have much work ahead in improving the techniques described here. Our goal is to make our techniques useful to nontechnical users. We are currently working with people from areas such as surveying, agriculture, and architecture in developing and evaluating improved techniques and applications to support their work. ■

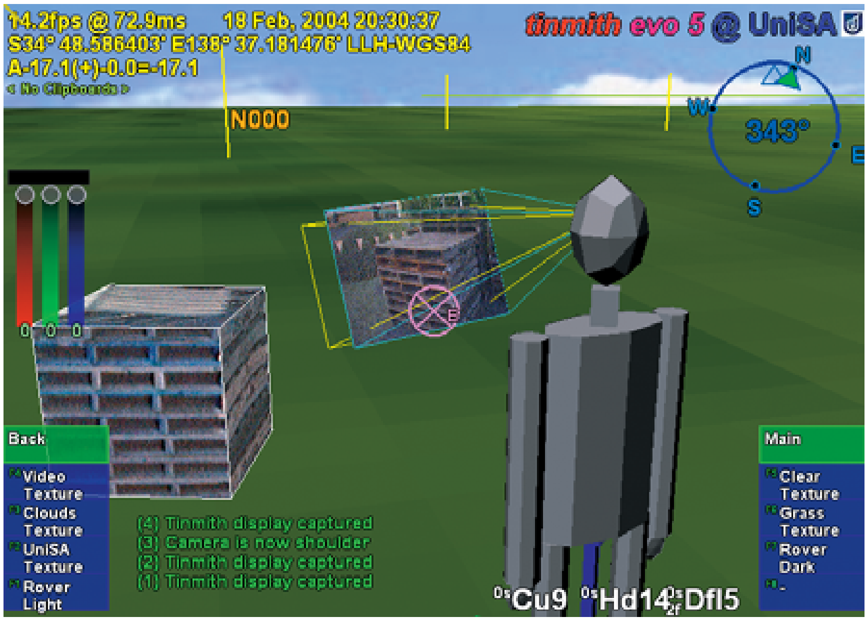

(4) Tinmith display captured
(3) Camera is now shoulder
(2) Tinmith display captured
(1) Tinmith display captured

5 **External viewpoint showing the user's body relative to other 3D objects. The live texture map shows what the user sees through the HMD, and that the system can automatically capture textures onto the surfaces of modeled objects.**

References

1. W.Piekarski and B.H. Thomas, "An Object-Oriented Software Architecture for 3D Mixed Reality Applications," *Proc. 2nd Int'l Symp. Mixed and Augmented Reality*, IEEE CS Press, 2003, pp. 247-256.
2. W. Piekarski and B.H. Thomas, "Augmented Reality Working Planes: A Foundation for Action and Construction at a Distance," *Proc. 3rd Int'l Symp. Mixed and Augmented Reality*, IEEE CS Press, 2004, pp. 162-171.
3. W. Piekarski and B.H. Thomas, "Interactive Augmented Reality Techniques for Construction at a Distance of 3D Geometry," *Proc. 7th Int'l Workshop Immersive Projection Technology/9th Eurographics Workshop Virtual Environments*, Eurographics Assoc., 2003, pp 19-28.

Readers may contact Wayne Piekarski at wayne@cs.unisa.edu.au.

Readers may contact department editors at lrosenbl@nsf.gov and julier@ait.nrl.navy.mil.