

Interactive Augmented Reality Techniques for Construction at a Distance of 3D Geometry

Wayne Piekarski and Bruce H. Thomas

Wearable Computer Laboratory, School of Computer and Information Science, University of South Australia
Mawson Lakes, South Australia, 5095, Australia
{wayne, thomas}@cs.unisa.edu.au

Abstract

This paper presents a series of new augmented reality user interaction techniques to support the capture and creation of 3D geometry of large outdoor structures, part of an overall concept we have named construction at a distance. We use information about the user's physical presence, along with hand and head gestures, to allow the user to capture and create the geometry of objects that are orders of magnitude larger than themselves, with no prior information or assistance. Using augmented reality and these new techniques, users can enter geometry and verify its accuracy in real time. This paper includes a number of examples showing objects that have been modelled in the physical world, demonstrating the usefulness of the techniques.

Keywords:

Augmented reality, construction at a distance, 3D user interfaces, wearable computers

Categories and Subject Descriptors (according to ACM CSS): I.3.6 [Computer Graphics]: Methodology and Techniques - Interaction Techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - Virtual Reality

1. Introduction

This paper presents new interactive augmented reality (AR) techniques we have developed that form part of an overall concept we have named construction at a distance, for use on outdoor wearable computers. Previously, we have been investigating how mobile AR users can capture planar shaped objects from the physical world using simple 3D primitives and constructive solid geometry (CSG), and also the placement of prefabricated objects¹⁶. Our work is novel since AR systems traditionally only render content, and a new useful feature is to be able to enter in new geometry information in situ. Fred Brooks² describes input devices, interaction techniques, and most importantly the creation and capture of 3D geometry as still being problems in the 3D computer graphics area.

We have come to the conclusion that specific user interface techniques are required to produce usable AR modelling applications outdoors. In this paper, we present our new techniques to support the capture and creation of outdoor 3D geometry; these are as follows: AR working planes, infinite carving planes, laser carving, laser colouring, texture map capture, and surface of revolution. These new techniques provide a set of tools that are capable of




Figure 1 - View of an automobile with corner marker points used to define the region for a carving operation

capturing and creating geometry not previously possible with our original system, such as trees, automobiles, and concave buildings.

Our implementation of techniques for construction at a distance uses a mobile augmented reality wearable computer that can be used outdoors. The user interacts with the computer using hand and head gestures. Using this inter-

face, an autonomous user can capture 3D geometry in Earth coordinates²³ with the aid of their physical presence. The AR rendered 3D geometry allows the user to immediately verify the results against the physical world in real time, which is not possible using any other existing object capture methods. The quality of the models generated by our techniques is limited only by the resolution and accuracy of currently available tracking devices, and as this technology improves, so will the results obtained. For example, a GPS accurate to 50 centimetres can only model outdoor objects on the order of metres (such as a car), while a 1 millimetre magnetic tracker allows small objects like drink cans or bottles. The work presented here does not attempt to address issues of registration or accuracy of tracking devices, but simply works within these limitations to provide the best solutions that can be achieved with the current technology. The main goal of our work is to design AR applications that demonstrate the technology in use, and takes advantage of new developments in technology as they occur. For example, GPS technology has improved from 5-10 metre accuracy to 2 centimetres with the introduction of Real Time Kinematics.

Currently in virtual reality (VR), a number of techniques have been developed for use in modelling applications. These applications traditionally provide tools to create objects in a virtual world, to manipulate them, and to fly around the world and perform scaling operations to work on a variety of object sizes. While techniques for action at a distance such as Spot Lights¹¹, Apertures⁸, and Image Plane Techniques²⁰ have been developed, these only perform simple manipulations on existing objects and cannot be used to create new ones due to the 2D nature of the techniques. VR systems also tend to operate within a fixed area, and so while movements of the body and head are used to control the environment, flying²¹ and scaled world techniques¹² are also needed when the tracked area is inadequate or objects are not within direct reach.

Although AR systems are similar in some ways to VR, they are unique in that they require registration of the physical and virtual worlds, and so flying around and perform scaling operations is not possible. While it is possible to reuse existing action at a distance techniques that maintain the user's position and scale, other techniques are required to perform creation of data within an AR environment. Our implementation of construction at a distance takes the concept of working planes from CAD systems¹³ and extends this to the area of augmented reality. AR working planes enable users to capture and create geometry that is out of reach, at a larger scale than the user, and co-located with the physical representation. Existing VR techniques cannot capture or create models of objects the size of skyscraper buildings without breaking the 1:1 relationship between the user and the virtual world, and do not allow the real time verification against the physical world

that AR provides. With our construction at a distance concept, the scale of the world is fixed and the user's presence controls the view and creation of primitives. The virtual geometry is created using absolute coordinates and is always registered against the physical world in real time.

Figure 1 and Figure 2 show an automobile that has been modelled using the laser carving technique discussed in this paper. The screen images in this paper were all captured using video overlay augmented reality, and show exactly what the user saw through the HMD while creating the examples.

This paper firstly presents a summary of existing interaction work in AR and VR, and also the capabilities of current physical capturing techniques, to provide comparisons to our work. Next, an overview of our previous work presented in¹⁶ is provided to give a foundation for the new techniques presented in this paper. The complete collection of our new modelling concepts and techniques AR working planes, infinite carving planes, laser carving, laser colouring, texture map capture, and surface of revolution are then named and defined, with examples and screen images. The following section then contains a discussion of the usability of the system while working outdoors. Finally, the implementation section briefly discusses how the mobile AR system is constructed.

2. Related work

This section discusses existing AR and VR modelling work, as well as techniques currently used for the capture and creation of physical models.

2.1. Outdoor AR systems

Traditionally, AR research has been inclined to focus on the presentation of information, which is created offline using desktop based systems. Systems such as Feiner's Touring Machine⁷ and the original Timmth work¹⁵ had limited capabilities to enter information, and were mostly




Figure 2 - Final resulting automobile model overlaid in AR view, with real automobile in background

focused with information presentation. The Tinmith-Metro paper in ¹⁶, along with different methods by Baillot in ¹, were the first to implement true outdoor AR object entry and editing operations. The work presented in ¹ uses the intersections of projected lines from different locations to enter in vertices, and then connecting them together to form shapes. Due to the generic nature of the technique, it is possible to enter in very arbitrary shapes, as there are fewer constraints, at the cost of increased modelling time and complexity for the user.

2.2. Indoor modelling systems

There has been a wide range of work in the area of 3D indoor modelling using monitors and head mounted displays (HMDs): Clark's surface editor ⁴ implemented direct manipulation of splines using a HMD and a wand; Sachs' 3-Draw ²² performs the creation of arbitrary models using direct manipulation of a stylus and tablet; Liang's JDCAD ¹¹ pioneered many new techniques such as lasers and spot lights for action at a distance using 3D input devices; Butterworth's 3DM ³ developed new user interfaces for immersive VR modelling; Forsberg's work with apertures ⁸ extended Liang's spot lights to use a circular cursor on the hand projected from the head into the scene; Pierce's image plane techniques ²⁰ extended Forsberg's aperture projection concept to introduce a series of selection methods based on the projection of the user's hands and fingers; Mine's CHIMP ¹² implemented within arms reach techniques based on proprioception and scaled world operations; Stoakley's Worlds-In-Miniature ²⁴ demonstrated remote manipulation using small copies of the world held within the hands; MultiGen's SmartScene ¹⁴ is a commercially available application which implements many of the previously mentioned techniques. A common trait of many of these systems is that they use direct manipulation to create new vertices and objects within arms reach, and action at a distance for manipulation when the object is not reachable. By using scaling techniques in combination with flying, the user may manipulate small buttons in an elevator, or place a large building on a street. Due to their non-exact input methods, these systems are also limited to conceptual modelling tasks and not precision modelling, since it is not possible to easily enter accurate data and compare the model against that of a physical object.

Computer Aided Design (CAD) systems have existed since Sutherland introduced the interactive Sketchpad system in ²⁵. These systems are used to create accurate 2D and 3D diagrams of real world objects, and use desktop input devices such as mice and keyboards. Since a mouse is 2D, when the cursor moves over the display, points cannot be created because there is no depth information. CAD systems introduced the concept of working planes ¹³ that can be defined using numeric input, from a different camera

angle, or relative to the surface of another object. The 2D cursor is then projected onto the working plane to produce full 3D coordinates.

2.3. Physical capture techniques

In the construction and surveying domain, a number of techniques are used to capture the geometry of outdoor structures. The most basic method (although very tedious) is to physically measure the structure, and then record the dimensions and geometry on paper. This information can then be used to recreate the object as a 3D graphical model using a desktop CAD system. Brooks ² also discusses many current capture methods, and points out that this is still an area with a number of problems.

Debevec's Façade system ⁶ allows the capture of objects using photographs of a building taken from various angles, and then processes these offline to produce 3D models. Façade is not automated and requires the manual specification of similar points between the images by the user to reconstruct geometry.

A common capture method used by surveyors is employing portable laser scanners, such as the commercially available I-SiTE ⁹ or Cyrax ⁵ scanners. These devices sweep a laser over a large outdoor area and capture distance measurements, constructing a detailed mesh of the object. The specifications for the I-SiTE scanner claim to measure a point every 30 cm at a distance of 100 metres, and to be able to reconstruct a large building from four different angles within one hour.

With all these techniques, to produce an absolute position model, a GPS must be used to offset some reference point in the model, and the equipment must be very carefully setup before use, which can be time consuming. The setup and capture time is also fixed no matter what the desired model mesh complexity is, and so for simple tasks it is not possible to capture an entire building with a basic mesh in a few minutes. Most importantly, the Façade and laser scanner methods are both ineffective in environments where it is not possible to obtain perfectly clear views of the building from all angles, such as when obscured by trees, clouds, or other objects. Given these limitations, we believe that our work is a useful tool to complement these techniques.

2.4. Previous modelling work

Tinmith-Metro is our original AR outdoor 3D modelling application ¹⁶, which allowed the placement of outdoor street furniture and the capture of simple building shapes. As previously mentioned, Tinmith-Metro extends image plane techniques ²⁰ to support object manipulation (translate, rotate, scale) and object selection in mobile augmented reality. Figure 3 is an example showing a virtual table that has been placed down on the ground and is being

manipulated into the correct position using these techniques.

One often overlooked aspect of AR systems is the user interface, and in a system with a large number of commands, making these all accessible in an intuitive fashion is important. Since traditional desktop interaction devices such as keyboards and mice are not suitable for outdoor mobile computing, new user interface methods are required. The primary input device is a set of pinch gloves which are tracked using a video camera and fiducial markers, and is the only input device required by the system. Figure 1 and Figure 3 show examples of the system being controlled by these gloves in an outdoor setting. At the bottom left and right are the menu strips, indicating the options that will be executed if the matching finger is pressed. The user navigates through the menu to perform modelling operations using a series of pinch and palm gestures. The menu interface is discussed in more detail in ¹⁸, and will not be discussed further since this paper focuses on the construction at a distance techniques.

With the ability to place down prefabricated objects and also very primitive shapes such as planes, a CSG engine can be used to combine these objects together to form more complex objects as required (see Figure 6). Given a set of simple input primitives and enough time, it is possible to iteratively apply CSG operations to produce an object of any complexity. As the simplest primitive capable of modelling arbitrary shapes, large planes are created to define a 3D convex volume, as shown in Figure 5. To create these planes, the user positions their head so that they can (at least partially) sight down a wall, and then on command the system fires a plane in this direction from the current body position. Sighting along a very long straight object is a very accurate means of placing the eye so that it is in line with the object, and has been used by sailors for centuries - by aligning landmarks while offshore, sailors can locate previous waypoints easily and accurately, even at large distances.




Figure 3 - Previously created table being moved into its correct position using direct manipulation hand gestures

The user walks around the object and creates the planes representing an object, and when this process is completed, the CSG engine is used to produce a solid shape that can then be used for further CSG operations to produce the final shape. Some problems with the original method was that it was not possible to preview the final shape until all the planes were in place, and if any were incorrectly placed the resulting object would be non-existent.

The main advantage of this technique is that it allows users to capture a building without having to actually stand next to or on top of it, and the user can model it from a distance while being partially occluded by other objects. This is very useful as many existing techniques such as photo and laser based scanning require a full view of the object, GPS trackers do not work well near large buildings, and in many cases standing on top of the building is not possible or too dangerous.

3. Construction at a distance techniques

In this section, we present the current collection of our new modelling techniques: augmented reality working planes, infinite carving planes, laser carving, laser colouring, texture map capture, and surface of revolution. Many objects that were previously not possible to model with our original Tinmith-Metro (or required an excessive number of nested operations) are now relatively easy, and some objects which were previously not possible at all with other methods are now possible as well. An important feature of these techniques is that they can be combined in various orders to build up more complex shapes, and iteratively applied so the user can start with simple objects and then refine them based on comparisons made in real time against the physical object. Brooks ² describes this iterative refinement process guided by the eye as an efficient way of focusing time onto the areas of the model that require it.

3.1. Augmented reality working planes

While image plane techniques and our AR extensions are useful for selecting and manipulating objects at a distance, they are not useable for the placement of new objects as there is no distance information. As a result, we define a new technique known as augmented reality working planes that is based on previous interfaces used in desktop CAD systems. AR working planes can perform a wider range of tasks than is possible with previous image plane techniques, such as construction at a distance. AR working planes are infinite in size and can be fired from the user's current location in a specified orientation, or created at an object with the surface normal matching the object or pointing toward the user. AR working planes can be created either in absolute world coordinates, or relative to the head orientation of the user. Relative AR working planes can be used to perform all operations previously performed




Figure 4 - Working plane at object distance, used for 3D manipulation operations with 2D cursors




Figure 5 - Infinite carving planes created by user to carve a solid convex shape from an infinite solid




Figure 6 - Concave objects can be interactively created by carving away other shapes with CSG

using image planes (see Figure 3 and Figure 4), except now the approach for both action and construction at a distance is now unified.

As the glove is moved by the user, the cursor is projected onto the current AR working plane to calculate its full 3D coordinates. Since the plane is infinite in size, the user can project the cursor onto the plane at almost any angle, although accuracy decreases toward the singularity that occurs when the plane is inline with the user's view. As in desktop CAD systems, the cursor can create vertexes in world coordinates for other operations, allowing the specification of objects at a distance.

3.2. Infinite carving planes

This section describes the infinite carving planes construction at a distance technique. It is based on our original CSG with primitive objects construction technique, but the generation of more complicated 3D models is made simpler when used in conjunction with the other new techniques presented in this paper. Previously, each of the primitive objects was specified and then a CSG intersection operation was performed to produce the final volume, sometimes giving undesired results if the user made a mistake at any point, and requiring the user to delete everything and start over again. With infinite carving planes, the object is now carved away iteratively and the user receives real time feedback of the infinite volume being bounded, allowing immediate undo in case of a mistake, and greatly improving ease of use.

Figure 5 shows a building (in thick black lines) and the location of the mobile AR user as they are sighting down each of the walls. When infinite carving planes mode is first started, the AR system creates an (approximately) infinite solid volume which will be used for carving. When the user is aligned with a wall, a menu option is selected, and an infinitely long vertical plane is fired from the user's head into the world. This plane divides the previous infinite solid into two parts, and the left or right portion (decided by the user from the menu) is carved away from the solid and removed. As the user sights down each wall, the solid will be reduced down to an object which is no longer infi-

nite in the X and Y axes. At completion, a floor is automatically created at ground level, and the roof is left unbounded for carving using other techniques, since it is not possible to sight down the roof of a very tall building. The final 3D shape is stored using absolute world coordinates and reflects the geometry of the physical building.

While our previous work¹⁶ demonstrated the repeated use of CSG to form concave objects, (see Figure 6) this was very time consuming, tedious, and sometimes impossible. By using the other techniques described in this paper, it is now possible to easily create objects such as those with zig-zag roofs, holes, and indentations, using source objects generated using infinite carving planes.

3.3. Orthogonal laser carving

Our new laser carving technique allows a user to modify an existing shape and cut away portions to produce new highly concave shapes in an intuitive way. Figure 7 shows a detailed diagram of how the carving operation is performed. When the user creates the first point, the system projects a virtual laser from the head through the hand and finds the closest polygon under the cursor. This polygon is then used to create an AR working plane that is coplanar to the polygon, and all points created will be projected from the user's hand against this plane. Each point on the AR working plane is then connected together to form a 2D concave outline, and orthogonally extruded at an infinite distance along the surface normal. This extruded shape is then passed to the CSG engine and used as a carving tool on the original object.

This technique is first demonstrated with a simple example of modelling a building with a sloped roof, as shown in Figure 8. The user must first model a box that is the same ground area as the building, and with a roof higher than any parts of the building (using infinite carving planes). To carve the roof, the user positions themselves in front of the building so that the entire slope is easily visible. The user then indicates with the hand cursor the first point defining the roof on the left, then the peak of the roof, and then the right side of the roof. To complete the selection, the user must enter points around the overall building

to indicate which part of the graphical object to keep. The parts of the object outside the selection region are then carved away to produce the final shape in Figure 8 and Figure 9.

One important aspect of the projection is that since the carving is performed using orthogonal extrusion from the AR working plane, the technique is position invariant and points can be entered from any location in front of the polygon. It is therefore possible to cut a flat roof on a 100 metre high building while standing at ground level and looking up. If the virtual laser projected from the hands was used to cut the object directly, the system would produce pyramid shaped extrusions, which in many cases is not desirable. Also, although the technique can create objects that are concave in cross section, some arbitrarily concave objects must be created by combining with other techniques.

For some buildings, the user may only desire to create a flat roof or a single slope, and so by creating only one point the system will create a horizontal cutting plane, and with two points a diagonal cutting plane is created. More than two points implies the user wishes to cut with an outline and so it must be fully specified. Other CSG operations may be used in conjunction with laser carving - the region selected may be used for an intersect or a difference operation, allowing holes to be cut into the middle of objects, or to split an object into two separate parts.

A second example showing the powerful capabilities of orthogonal laser carving is a small automobile modelled outdoors in Figure 2. Firstly, a box is created approximating the overall dimensions of the car. Next, the user views the car from the side, and intersects points against the box surface to define the silhouette. Figure 1 shows the markers being placed, and Figure 2 shows the final solid shape of the car. The object can then be carved from the front (or any other angle) to further refine the model until it suits the user's requirements.

3.4. Laser colouring

Once a building has been created, it may be desirable to place windows, doors, and other annotations onto the model. While it may be possible to draw these using a texture map (which cannot be zoomed arbitrarily), or to place extra polygons outside the building to represent these (covering the original building), the building model itself is still a solid shape. If these new polygons are removed or manipulated, the original solid object remains as the changes are only superficial. A more desirable scenario is that polygons of a different colour are actually cut into the subdivided surface of an object, so that if they are deleted it is possible to see features inside the object that were previously concealed by the walls. We have named this technique laser colouring. Using the same steps as laser carving,

points are created against a coplanar AR working plane to form an outline. Instead of carving away the object, it is instead subdivided and the colour of the outlined polygons is changed.

3.5. Texture map capture

Since the mobile AR computer already includes a video camera for the AR live video overlay, it is possible to pro-




Figure 7 - Virtual lasers perspectively projected onto an AR working plane coplanar with object, then carved orthogonally along the surface normal. This method is invariant to the user's height, position, and plane angle.



Figure 8 - AR view showing building model with sloped roof registered against the physical object



Figure 9 - VR view of building with sloped roof, showing the overall geometry from an overhead viewpoint



Figure 10 - VR view of cube model showing 3D geometry along with captured textures to improve detail

ject images from this camera directly on to polygon surfaces. We call this technique texture map capture, and it allows us to easily map textures to polygons simply by looking at them.

Once the object is modelled and in the correct location, the user stands at a point where the object textures to be captured are easily visible. The capture menu option is selected and the system projects the polygon vertices onto the video overlay and maps the still image to the surface of the polygon. The user repeats this operation for each face until the object is completely textured. An example is shown in Figure 10, where a 1 metre box was modelled and then texture maps placed onto the polygon faces.

The best results for this capture are gained when the object fills as much of the HMD as possible (while still being completely visible), and is perpendicular to the user's viewing direction, such as in the laser carving example in Figure 7. Since OpenGL only supports linear texture mapping, if the angle to the object is large (like the roof of the box in Figure 10), the image will be highly distorted as a non-linear mapping is required.



Figure 11 - AR view of surface of revolution tree with marker points projected onto an AR working plane

Texture maps may be combined with the other techniques to produce complex shapes with textured surfaces. When these models are viewed indoors, this type of approach provides the user with a better understanding of the final models compared to the use of simple coloured polygons.

3.6. Surface of revolution

When working outdoors, many objects do not fit against approximations like boxes or polygons, and contain many curved features. In an attempt to capture objects such as trees, we have used surface of revolution techniques (as used in many desktop CAD systems) to capture their shapes more accurately than is possible using infinite carving planes. This method is completely self-contained and is not dependent on models being supplied from other techniques. This technique also demonstrates another scenario where AR working planes can be used to create objects that are not possible with existing action at a distance techniques.

To create a tree (with a reasonably symmetrical shape), the user sights toward the central trunk or axis of the tree, and fires an AR working plane from the body. The user then walks to a point approximately 90 degrees around the tree so that the normal of the AR working plane is now pointing towards them. Using techniques similar to those in laser carving, the user projects points onto the AR working plane to outline the tree, as shown in Figure 11. The user next creates either one or two points to indicate the axis of rotation for the tree, and the system then generates a preview of the final object. Once the user is satisfied with the model, it is solidified and placed into the virtual world.

For modelling regular objects such as pine trees, which are highly symmetrical about the trunk, this technique works well. For trees that grow with deformities and other abnormal features, the models will not match very well. To better approximate these forms of physical objects, laser carving may be applied to refine the surface of revolution graphical object until the user is satisfied with the shape.

4. Integrated example

While each technique allows us to produce useful shapes, the true power of the techniques is expressed when they are used in combinations. This integrated example is designed to highlight all the features of the modelling techniques in a short simple description, and not to create a truly usable building. A user walks outside to an empty piece of land, and creates a building that they would like to preview and perhaps construct in the future if it fits the environment to their liking. The user rapidly constructs a model to what they would like, and then visualises it to see the final result. Alternately, this model could be viewed in real time or later on an indoor workstation. As well as the



Figure 12 - AR view of constructed final model, including prefabricated street furniture items



Figure 13 - VR view of the final model, with external viewpoint to show detail not visible on an immersive display

discussed construction example, this application could also be used for creative purposes such as an abstract art visualisation tool, or a landscape garden editor and viewer. It should be noted that the physical techniques such as photo based reconstruction or laser scanning are not usable at all as the object does not exist yet, and so our techniques can be used for creation and not just capturing.

Figure 12 and Figure 13 show different views of the object at the end of the construction process. The first step is to create the perimeter of the building shape using our existing bread crumbs technique described in 17. The user walks around the building site and places down markers at the desired ground locations, forming an outline. Next, the outline is extruded upwards so that it can be carved. Using laser carving, the user firstly cuts a main roof to make the object finite, and then carves a slope using various control points. After the overall roof structure is created, the object is lifted into the air using the gloves. At this point, the supporting columns, trees, tables, and avatar people are created using street furniture placement of prefabricated mod-

els, and put into position. The building is then visually lowered onto the supporting columns. At this point, the user performs further carving and a large hole is created through the centre of the building. Next, laser carving is used to cut out two large sections of the building, causing it to exist as three unattached solid shapes, most visible in Figure 12. After around 10 minutes for this example, the desired model is complete and the user can now walk around the object to preview it from different angles.

5. Discussion

This section contains a discussion of the usability of the system while working outdoors. This understanding was gained by operating the system in the field with real life examples.

5.1. External views

AR is designed to overlay virtual objects onto the user's view of the world in real time. With our modelling techniques, AR allows the user to verify the resulting 3D models. However, in some cases it is more useful to work in an external VR style view (such as orbital view 10) with the live video replaced by a sky and ground texture. In examples like infinite carving planes and bread crumbs, it is more intuitive to use the technique while working in an external VR view (see Figure 13) than in the immersive AR view with a limited field of view (see Figure 12), as the user has a higher situational awareness of the objects that have been placed around them.

These external views were also used for the captured videos and photographs presented with this paper as it shows an easy to understand overhead preview of the final object shape, which might not be otherwise obvious at ground level with immersive AR. By interfacing our outdoor AR system with an indoor VR system (using the same software) it is possible for users indoors to observe what objects are being created outdoors 15.

5.2. Error analysis

The techniques in this paper rely on the position and orientation sensors for all tracking, and so increasing the accuracy of these devices will produce improved results and alter the minimum sized model that can be properly captured. However, errors from each sensor have different effects on the captured models since one is measured as a distance and the other as an angle. Also, when rendering the AR display, results are affected not only by the errors in the current tracker data, but also those from the capture process.

The position sensor used in this paper is a Trimble Ag132 GPS, with an accuracy of better than 50 centimetres at 10 updates per second, and working reliably amongst small buildings and light tree cover. To ensure the most

accurate positioning possible, an indicator on the HMD shows the quality of the GPS signal. For orientation, an Intersense IS-300 hybrid magnetic and gyroscopic sensor with very fast update rate is used, although the accuracy is unreliable when there are magnetic distortions present in the environment or when the user is moving quickly and disturbing the sensor. Since the accuracy of the IS-300 can vary unpredictably, this is the most critical component in terms of reliability and accuracy.

When modelling a new object, projections are calculated based on the user's current location and the direction they are looking. For the highest accuracy, it is desirable to be as close to the object as possible, minimising the distance the projection can stray from the desired direction caused by angular error of the orientation sensor. When viewing an existing virtual object, the registration with the real world caused by the GPS will be the most accurate when viewed from a distance due to perspective, while standing very close to an object will cause these errors to be very noticeable. For registration errors caused by the IS-300, these remain constant (at all distances) on the display due to their angular nature.

While GPS is fast and reliable (with accuracies of Real Time Kinematic units at 2 cm in ideal conditions), improvements in orientation sensing technology are required to make the application easier to work with. While the IS-300 is one of the best mid-range priced trackers on the market, it is still not completely adequate for modelling outdoors. Due to its use of magnetic sensors, distortion caused by sources of magnetic interference such as electrical cables, heavy machinery, street lamps, and the backpack itself affect the accuracy of the tracking. To correct this, the backpack uses a touch pad mouse mounted on the user's chest to apply a correction offset. The touchpad allows the user to fine tune the calibration while moving around outdoors through varying magnetic fields. We are currently investigating methods of automating this calibration but at the present, automatic methods such as optical natural feature tracking are still an ongoing research problem.

In order to take advantage of the stability of position tracking, we propose to modify our infinite carving planes techniques so that the planes are projected based on two or more position points and ignoring the orientation sensor. The user sights down the wall and marks the position, and then walks closer and marks a second point while still sighting down the wall. These points can then be used to specify an infinite carving plane, and the accuracy can be improved by increasing the spacing of the marker points or using a line of best fit on multiple points. This technique is invariant to errors in orientation since it is only based on position values.

6. Implementation

The techniques discussed in this paper have all been implemented and tested in an outdoor environment. We have developed a complete and useable modelling application known as Tinmith-Metro, which allows users to capture 3D outdoor geometry in the field. This application supports all the features discussed in this paper, plus others such as the loading and saving of models, colouring, texture maps, AR and VR style viewing, real time mobile AR demo video recording, and a user interface fully controlled using pinch gloves and hand tracking. To implement this application, we use our flexible Tinmith-evo5 software architecture¹⁹, which is a complete toolkit for the development of high performance 3D virtual environment applications.

A custom backpack computer was designed to support our research, enabling a single user to carry all the necessary equipment. Our current backpack is known as Tinmith-Endeavour, and shown in Figure 14. It is based around a Pentium-III 1.2 GHz laptop with Nvidia GeForce2, mounted onto a special polycarbonate backplane. An Intersense IS-300 hybrid magnetic and gyroscopic tracker is used for orientation sensing, and a Trimble Ag132 GPS with an accuracy of 50 cm is used for position sensing. The helmet is based on a Sony Glasstron head mounted display with a FireFly 1394 camera for live video input. The backpack operates for 2 hours with a 12V battery rated at 85 Wh, and weighs approximately 16 kg, although no efforts have been made to miniaturise or lighten the design since it is currently a research prototype with multiple uses.

7. Conclusion

This paper has presented our construction at a distance concept and the AR working planes, infinite carving planes, laser carving, laser colouring, texture capture, and surface of revolution techniques. When combined with a



Figure 14 - The Tinmith Endeavour backpack, jointly developed by the authors at the Wearable Computer Lab of the University of South Australia with the Defence Science Technology Organisation - Salisbury

user interface using hand and head gestures, users can capture the geometry of 3D outdoor structures quickly and intuitively. We have demonstrated a number of examples showing each technique used for outdoor object construction with an augmented reality user interface for real time verification. The examples demonstrated real scenarios where existing methods such as image and laser based capturing would not be possible and hence show the usefulness of these techniques for outdoor modelling work. Finally, we discussed insights gained from these field tested examples.

8. References

- Baillot, Y., Brown, D., and Julier, S. Authoring of Physical Models Using Mobile Computers. In *5th Int'l Symposium on Wearable Computers*, pp 39-46, Zurich, Switzerland, Oct 2001.
- Brooks, F. P. B. J. *What's Real About Virtual Reality?* IEEE Computer Graphics and Applications, Vol. 19, No. 6, pp 16-27, Nov 1999.
- Butterworth, J., Davidson, A., Hench, S., and Olano, T. M. 3DM: A Three Dimensional Modeler Using a Head Mounted Display. In *Symposium on Interactive 3D Graphics*, pp 135-138, Cambridge, Ma, Mar 1992.
- Clark, J. H. *Designing Surfaces in 3-D*. Communications of the ACM, Vol. 19, No. 8, pp 454-460, Aug 1976.
- Cyra Technologies Inc. *Cyrax 2500 Laser Scanner*. <http://www.cyra.com>
- Debevec, P. E., Taylor, C. J., and Malik, J. Modeling and Rendering Architecture from Photographs: A hybrid geometry- and image-based approach. In *ACM SIGGRAPH 1996*, pp 11-20, New Orleans, La, Aug 1996.
- Feiner, S., MacIntyre, B., and Hollerer, T. A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment. In *1st Int'l Symposium on Wearable Computers*, pp 74-81, Cambridge, Ma, Oct 1997.
- Forsberg, A., Herndon, K. P., and Zeleznik, R. Aperture Based Selection for Immersive Virtual Environments. In *9th Int'l Symposium on User Interface Software and Technology*, pp 95-96, Seattle, Wa, Nov 1996.
- I-SiTE. *I-SiTE 3D Laser Scanner*. <http://www.isite3d.com>
- Koller, D. R., Mine, M. R., and Hudson, S. E. Head-Tracked Orbital Viewing: An Interaction Technique for Immersive Virtual Environments. In *9th Int'l Symposium on User Interface Software Technology*, pp 81-82, Seattle, Wa, Nov 1996.
- Liang, J. and Green, M. Geometric Modelling Using Six Degrees of Freedom Input Devices. In *3rd Int'l Conference on CAD and Computer Graphics*, pp 217-222, Beijing, China, Aug 1993.
- Mine, M., Brooks, F. P., and Sequin, C. H. Moving Objects In Space: Exploiting Proprioception In Virtual-Environment Interaction. In *ACM SIGGRAPH 1997*, pp 19-26, Los Angeles, Ca, Aug 1997.
- Mine, M. R. *Exploiting Proprioception in Virtual-Environment Interaction*. PhD Thesis. Chapel Hill, NC, University of North Carolina, 1997.
- Multigen. *SmartScene*. <http://www.multigen.com>
- Piekarski, W., Gunther, B., and Thomas, B. Integrating Virtual and Augmented Realities in an Outdoor Application. In *2nd Int'l Workshop on Augmented Reality*, pp 45-54, San Francisco, Ca, Oct 1999.
- Piekarski, W. and Thomas, B. Tinmith-Metro: New Outdoor Techniques for Creating City Models with an Augmented Reality Wearable Computer. In *5th Int'l Symposium on Wearable Computers*, pp 31-38, Zurich, Switzerland, Oct 2001.
- Piekarski, W. and Thomas, B. H. Bread Crumbs: A Technique for Modelling Large Outdoor Ground Features. In *Int'l Symposium on Mixed and Augmented Reality*, Darmstadt, Germany, Oct 2002.
- Piekarski, W. and Thomas, B. H. The Tinmith System - Demonstrating New Techniques for Mobile Augmented Reality Modelling. In *3rd Australasian User Interfaces Conference*, Melbourne, Vic, Jan 2002.
- Piekarski, W. and Thomas, B. H. *Tinmith-evo5 - A Software Architecture for Supporting Research Into Outdoor Augmented Reality Environments*. Technical Report, University of South Australia, Adelaide, SA, Report No. CIS-02-001, Jan 2002.
- Pierce, J. S., Forsberg, A., Conway, M. J., Hong, S., Zeleznik, R., and Mine, M. R. Image Plane Interaction Techniques in 3D Immersive Environments. In *1997 Symposium on Interactive 3D Graphics*, pp 39-43, Providence, RI, Apr 1997.
- Robinett, W. and Holloway, R. Implementation of Flying, Scaling, and Grabbing in Virtual Worlds. In *1992 ACM Symposium on Interactive 3D Graphics*, pp 189-192, Cambridge, Ma, Mar 1992.
- Sachs, E., Roberts, A., and Stoops, D. *3-Draw: A Tool For Designing 3D Shapes*. IEEE Computer Graphics and Applications, Vol. 11, No. 6, pp 18-24, Nov 1991.
- Spohrer, J. C. *Information in Places*. IBM Systems Journal, Vol. 38, No. 4, pp 602-628, 1999.
- Stoakley, R., Conway, M. J., and Pausch, R. Virtual Reality on a WIM: Interactive Worlds in Miniature. In *CHI 1995 - Conference on Human Factors in Computing Systems*, pp 265-272, Denver, Co, May 1995.
- Sutherland, I. Sketchpad: A man-machine graphical communication system. In *AFIPS Spring Joint Computer Conference*, pp 329-346, 1963.

For more information, as well as videos of the system in use outdoors, readers are invited to visit our project's web site at <http://www.tinmith.net>