Augmenting Image Plane AR 3D Interactions for Wearable Computers

Thuong N. Hoang, Shane R. Porter and Bruce H. Thomas
Wearable Computer Lab
School of Computer and Information Science
University of South Australia

{hoatn001, porsr001}@students.unisa.edu.au – Bruce.Thomas@unisa.edu.au

Abstract

This paper presents a set of large object manipulation techniques implemented in a wearable augmented reality computer system that are optimised for the outdoor setting. These techniques supplement the current image plane approach, to provide a comprehensive solution to 3D object manipulation in an augmented reality outdoor environment. The three extended manipulation techniques, Revolve, Xscale, and Ground plane translation, are focused on using what we determined to be the best coordinate system for object rotation, scaling and translation. This paper goes on to present the generalised plane technique for the constrained translation of graphical objects on arbitrary planes to enable more complex translation operations. The paper presents the techniques from both the user interface and software development perspectives.

Keywords: 3D object manipulation, outdoor augmented reality, user interfaces, and virtual environments.

1 Introduction

Augmented Reality (AR) systems provide an interactive medium through which the real-world environment is supplemented with virtual information. This paper presents a set of action-at-a-distance techniques (Bowman & Hodges 1997) to overcome a number of limitations found in the 2D nature and inherent relative coordinate systems of image plane techniques. Image plane techniques have been shown as an intuitive and useful technology with action/construction-at-a-distance manipulation (Piekarski & Thomas 2001b) for a user operating a wearable augmented reality computing system (see Figure 1). In particular the image plane's natural coordinate systems for standard manipulation tasks (translate, rotate, and scale) are less than optimal for manipulation tasks of large virtual objects in an outdoor setting. This paper will explore a set of user interface manipulation techniques that are more appropriate for the manipulation tasks of these large virtual objects in an outdoor setting.

Our investigations into AR in an outdoor setting have covered a range of activities, including collaboration (Stafford et al. Oct. 2006), modelling (Piekarski & Thomas 2001b), and entertainment (Thomas et al. 2000). In particular we have investigated the use of image plane technique (Pierce et al. 1997) for action-at-a-distance manipulation of large scale graphical objects, such as buildings, trees, and street furniture. A common task is the placement of virtual objects co-located with physical objects with a specific location, orientation, and scale. Current user interfaces for the manipulation of graphical objects in AR focus on a unified methodology of employing the image plane techniques for all of the canonical object manipulation tasks of rotation, and translation (Bowman et al. 2005, pp. 141 - 3), with the addition of scaling. These approaches limit the manipulation to one coordinate system (one that is relative to the image plane the head mounted display), which is not suitable for all manipulation tasks.

Figure 1. Tinmith augmented reality wearable computer system

This paper presents a number of interaction techniques that have been optimised for the best relative coordinate system for the manipulation of large virtual graphical objects (large in scale to the user) in an outdoor setting. Each of the direct manipulation operations: translation, rotation, and scale, requires proper frames of reference and manipulation techniques. We found that a single technique was unsuitable, and a range of tailored techniques is required for each task at hand. These techniques leverage the following assumptions concerning the virtual objects to be manipulated:

1. the virtual objects reflect large physical objects,
2. the virtual objects would be influenced by gravity (they will settle on the ground and manipulations will naturally reflect this force),

3. the final transformation will leave the virtual objects in a physically sensible position (the virtual objects will not be protruding into the ground or house will rotated to the floor is not co-planar with ground,
4. the centres of gravity of the virtual object are placed on the ground plane,
5. the object's coordinate system is based around the centre of gravity of the object, with the Z axis normal to the ground plane.

The innovation presented in this paper is not the manipulation techniques in isolation; but by grouping them together in proper context, these techniques provide a more natural interaction for operations on large graphical objects in an outdoor augmented reality setting. We present the following three manipulation techniques: Revolve, Xscale, and Ground Plane Translation. Revolve and Xscale support rotation and scaling of objects about an axis normal to the ground plane and relative to their own coordinate axes, as opposed to the image plane technique, which rotates and scales about a normal to the image plane. Using an object's natural up direction (a normal to the ground plane) and its own coordinate system provides more natural rotation and scaling of objects. Ground plane translation allows translating objects with 2DOF (X and Z in first person perspective) constrained to the ground plane, supplementing the original image plane (X, Y) degrees of freedom. The ground plane translation supports the notion of object being moved along ground and once moved do not normally floating above or sinking below the ground.

This paper starts with an overview of previous investigations into action/construction-at-a-distance techniques. Our work focuses on the improvement of the current image plane techniques; therefore, limitations of using image plane techniques for the manipulation of large objects in an outdoor setting are examined. Direct manipulation techniques to overcome the discussed limitations are presented. We then generalise the ground plane translation technique to arbitrary planes. Finally, a set of concluding remarks is given.

2 Background

Action-at-a-distance techniques allow users to interact with objects that are not within arm's reach. Lasers and spot lights are early examples of 3D action-at-a-distance techniques (Liang & Green 1993). Apertures (Forsberg et al. 1996) are an extension of the spot lights techniques that employ a circular cursor on the hand projected from the head into the scene for selection.

Another interaction technique is that of the GoGo arm (Poupyrev et al. 1996). Movements of the user's arm provide input for an arm in the virtual world. As the user extends their arm, the virtual arm extends at a greater rate. This simple technique allows the user to interact with distant objects.

The concept of Worlds in Miniature (Stoakley et al. 1995) allows the user to interact with distant objects by performing actions within arm's reach. The user creates a miniature version of the virtual world, on which actions could be performed to interact with distant objects. Actions performed on objects in the miniature world are translated to their counterparts in the normal sized world.

Image plane techniques (Pierce et al. 1997) further extend the aperture projection concept to introduce a series of selection methods based on the projection of the user's hands and fingers onto the scene. These techniques were extended by Piekarski and Thomas (Piekarski & Thomas 2001b) for direct manipulation. Bowman's CDS (Bowman 1996) used object extrusion determined by intersecting points against a ground plane with a laser pointing technique. Leveraging the ground plane is a very useful technique, and this is especially true when operating on large virtual objects in an outdoor setting.

The Tinmith system is an outdoor AR wearable computer system in the form of a belt mounted computer system, a video see-through head-mounted display, and pinch gloves input. Tinmith supports a range of 3D modelling techniques. Pinch gloves are used as input devices, with fiducial markers placed on the thumbs for cursor manipulation. The menu system is located on both left and right lower corners of the display with 10 options available at any time, each of which is mapped directly to the fingers on the corresponding pinch glove (Piekarski & Thomas 2001a).

Tinmith implements a combined version of Pierce's image plane techniques and the AR working plane concept by Piekarski and Thomas (Piekarski & Thomas 2004). Objects selection is done by projecting distant objects and the input cursors onto an AR working plane attached to the user's head orientation, and finding which object intersects with the input cursor (Piekarski 2004). After selection, the same working plane is used for object manipulations, including translation, rotation and scaling. Moving the cursor across the working plane will translate the object to the required position, without changing the distance between the object and the user. Rotation and scaling operations are done using two input cursors, one for the origin and the other for specifying the rotation angle or scaling vector. Rotation could only be performed about the surface normal of the AR working plane; likewise, scaling is only along the surface of the working plane.

3 Limitations

Image plane techniques as implemented by Pierce et al. (1997) and Piekarski and Thomas (2001b) have a number of limitations. Image plane techniques are 2D (X_p, Y_p relative to the image plane) in nature, and have no real ability to work with depth (Z_p away from the user's image plane). The coordinate system of (X_p, Y_p, Z_p) refers to that of the image plane, as illustrated in Figure 2. The interactions are restricted to 5DOF (out of a possible 9DOF for full 3D interactions), two for translation (X_p, Y_p), two for scaling (X_p, Y_p), and one for rotation (Z_p).

The original Pierce et al. (1997) image plane techniques were developed for object selection in a virtual reality environment, incorporating a precise 6 DOF tracking system and some techniques requiring a set of data gloves. In the techniques presented by Pierce et al. (1997), the actual object manipulation was performed by secondary within-arm-reach techniques, forcing the user to employ one technique for selection and another for manipulation.

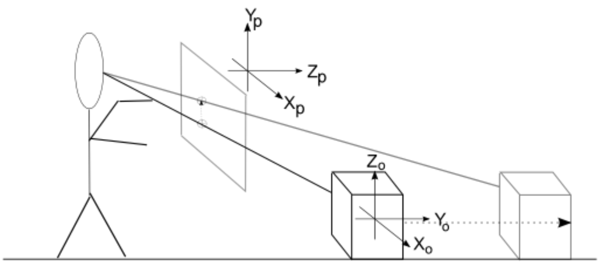

Figure 2: Translation in 3rd DOF using the original image plane (Piekarski & Thomas 2004)

The application of image plane techniques for the manipulation tasks of translation, rotation and scaling has different limitations. As previously mentioned, the image plane techniques by Pierce et al.(1997) were designed only for object selection task, and the concept was extended for object manipulation by Piekarski and Thomas (Piekarski & Thomas 2004). A number of the limitations of these techniques are examined to provide an insight into the rationale of our new techniques.

In both the Pierce et al. (1997) and Tinmith versions, the user is required to keep their head very still or errant information is passed onto the interaction task. Therefore, support for precise tracking of input devices is not as ideal as in the virtual reality environment. A translation operation employing image plane techniques may only have movements of an object that are parallel to the image plane. This was extended to be constrained to move relative to other objects or coordinate systems with Piekarski and Thomas' AR working planes (2004). In order to change the object's position in the 3rd DOF, i.e. along the axis that is perpendicular to the head-relative AR working plane, the user has to physically move backward or forward. One disadvantage of this technique is the amount of physical walking the user performs to place an object in the required position, especially when translating large virtual objects. Because the user is able to translate the object freely, within the constraint axes of the image plane, a second problem arises. This will allow a user to place an object into either an unrealistic 'floating' above or 'sinking' underground state; in other words, objects are not constrained to the ground plane. Such a constraint is too restrictive for generalised graphical operations. However when manipulating large virtual objects in an outdoor setting, this is justified in keeping with the natural movement of large objects in the physical world.

The image plane rotation technique has the user interact with the 2D projection of the objects on the AR working plane (Piekarski & Thomas 2004). The implementation of image plan rotation in Tinmith employs the intuitive use of both hand cursors to specify the origin point and the rotation point. This technique is limited in the fact that the user can only rotate objects around an axis that is perpendicular to the image plane. Rotation of large 3D graphical objects such as houses, and trees in an outdoor environment is mainly about the normal axis to the ground; however, the image plane manipulation techniques do not support this form of operation. There is another problem when specifying the axis of rotation via the normal to the image plane. The image plane technique requires the user to exactly line up a normal vector of the image plane with the desired axis. Difficulty in achieving this may cause the object to rotate about multiple axes at the same time. From a practical point of view, rotating about an axis normal to the image plane is very often required.

Image plane scaling is performed by specifying one control point and one manipulation point and altering the relative distance between cursors. Scaling operation is not bounded, which is more prone to tracking errors and undesirably scaling objects below the ground plane. Moreover, this technique is limited with support of only 2DOF and requiring both hand cursors to be visible, causing imprecise performance due to tracker errors. A major limitation is the graphical operation of scaling about an arbitrary axis is of limited use. A typical scale operation would be to make an object wider or taller, but not larger due the proportions dictated by this arbitrary viewing axis.

4 Object manipulation techniques

Currently, the Tinmith system only supports the original image plane object manipulation techniques. Due to its limitations in manipulating large 3D graphical objects, we have extended Tinmith with the following three extra object manipulation techniques: Revolve, Xscale and Ground plane translation. These operations have been constrained, especially with regards to the ground plane, to allow them to be more useful in the context of manipulating large graphical objects in an outdoor setting. The rotation and scaling manipulations employ the objects' coordinate system. We use the centre of gravity to define as the origin of the objects' coordinate system. The set of techniques are easily implemented in any wearable computer system that employs 2DOF input devices and a menu system. The existing Tinmith manipulation techniques are still available to the user for those operations not conveniently supported by these three new techniques.

4.1 Rotation techniques

The Revolve technique allows the user to rotate an object about one of its own coordinate axes, X_o, Y_o or Z_o with vertical axis Z_o normal to the ground as specified in the scene graph (see Figure 2). Once the object has been selected, the Tinmith menu displays three options for X_o, Y_o or Z_o axis. The default is Z_o that is normal to the ground plane. This design decision is from the observation that most large graphical objects we have dealt with are a representation of physical objects, and have a natural up orientation. The most common rotation is around an axis normal to the ground plane, an example of which is the rotation of a virtual house around its vertical axis to face it in different directions in an outdoor architectural modelling task. This is a rotation operation not possible with the image plane metaphor. The X_o and Y_o axes are included for completeness (see Figure 3).

The user moves the hand cursor horizontally, between vertical edges of the display, in order to rotate the selected object about the selected axis. Full horizontal movement of the cursor from left to the right vertical edges of the display is mapped to a 360 degree rotation of the object. The cursor is required to be within view of the head-mounted camera. Should it leave the camera view, Revolve operation will be paused (i.e. the virtual object is kept in the same stage as when the cursor is last seen) and only resume when the cursors are detected again. The menu options are available on both the left and the right hand menu, catering for left or right handed users. A particular drawback with this method is that the user only employs one hand and is unable to choose the control point of rotation. Although the use of two hands for object manipulation has been shown to improve performance (Buxton & Myers 1986), we experienced discomfort after long periods of activity. We believe this is due to the user having to hold both hands in front of the camera. With the Revolve technique, only the dominant hand needs to be in view of the camera for object interaction, while the non-dominant hand is used for menu selection of X_o, Y_o or Z_o axes. This will reduce the user's weariness during long periods of interaction.

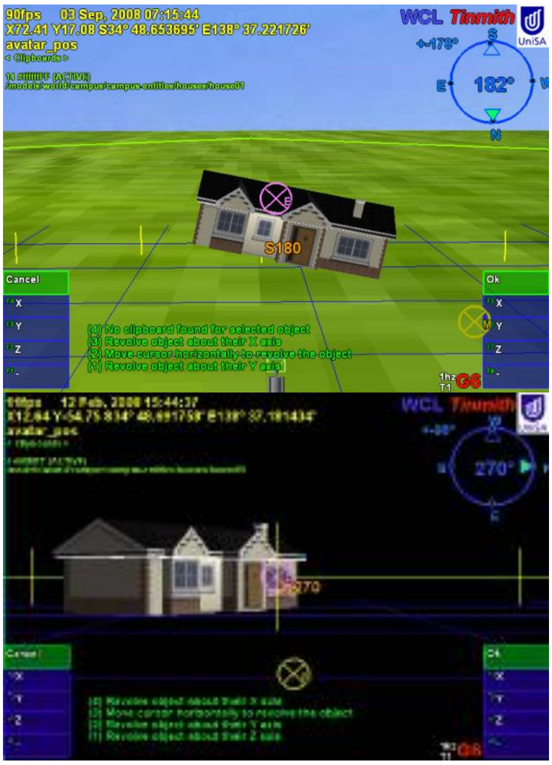

Figure 3: Revolve around the Y (top) and Z (bottom) axis

For all rotations performed by Revolve, the control point of rotation is naturally the centre of gravity of the object, but placed on the ground plane to ensure that rotation operations would not render the object in a floating or underground state. The centre of gravity is defined within the object itself, and it is not calculated by the Revolve operation. As with traditional graphical manipulation operations, Revolve enables the user to apply successive rotations to the object. This refinement allows the rotation of the object to any orientation achievable. By subsequently applying different rotation and translation operations to an object, the final orientation of the object can be achieved as if any control point was chosen. As previously mentioned, the original Tinmith rotation technique is still available to the user for those occasions where it can be used appropriately.

4.2 Scaling techniques

Scaling is performed in the object's coordinate system. Objects are constrained to being above the ground during scaling operation, i.e. they grow up and not into the ground. The Xscale technique supports scaling along the X_o, Y_o or Z_o axes of the objects, solving the similar problem of axis selection that adversely affects object rotation. The default is to scale all three axes at once, as this has been the most common mode of operation. The three scaling axes are locked together to maintain aspect ratios, as in image editors. The user operates the Xscale technique is the same fashion as the Revolve.

Upon activation of the Xscale technique, a ray is shot from the user's head position through to the input cursor; and the first object to be hit by the ray is selected for manipulation. From then, horizontal movements of the input cursor directly manipulate the scaling ratio of the object along the selected axis (see Figure 4). Selection of the scaling axis is done using the pinch glove.

The control point of scaling is the intersection of the ground plane with a line perpendicular to ground plane through object's centre of gravity. Unlike the current Tinmith scaling operation, Xscale always keeps the object above the ground plane, especially for scaling performed along the Z_o axis (see Figure 5). Objects are scaled only in the same up direction as the normal axis to the ground, whereas Xscale along the other axes scales objects uniformly in both positive and negative directions. This constrained Xscale approach helps reduce unwanted behaviours of objects being scaled below the ground.

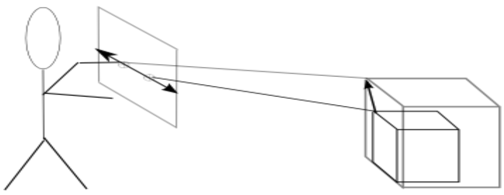

Figure 4 Xscale technique

Figure 5: Xscale for a virtual house along the Z axis

4.3 Translation techniques

The original image plane technique supports object translation in the X_p and Y_p directions in the first person perspective (or X_o and Z_o in object's coordinate system), without being able to change object distance to the user (the Z_p coordinate). To supplement for the third DOF, Ground plane technique allows user to perform translation of objects along its X_o and Y_o axes while being constrained to the ground plane.

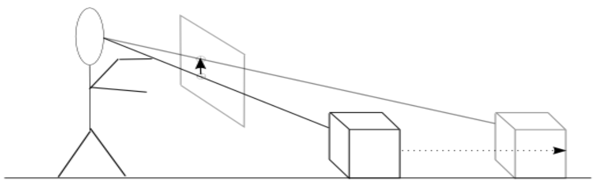

Figure 6: Ground plane translation technique

Only requiring 2DOF input devices, Ground plane technique maps the 2D movements of the input devices on the display plane to 2D translations of objects on the ground plane, which is presented in Tinmith as a virtual, plateau plane coinciding with the real world ground plane. The mapping is performed by casting a ray from the user's head position through the input device towards the object and using the intersection of the ray and the ground plane to specify object's location. This approach improves the 1 DOF object grabbing technique of ray-casting and fish-reel (Bowman & Hodges 1997) to a 2 DOF object manipulation technique. Because the distance between the intersection point of the ray to the ground plane and the user could range from zero to infinity, Ground plane technique enables non 1-to-1 mapping from bounded movements of the input devices to translation of object to infinity.

In simple words, the up and down movements of the cursor on the display move the objects further away or closer to the user. During the translation, the objects remains attached to the ground plane or a plane parallel to the ground plane if the object has a non-zero altitude.

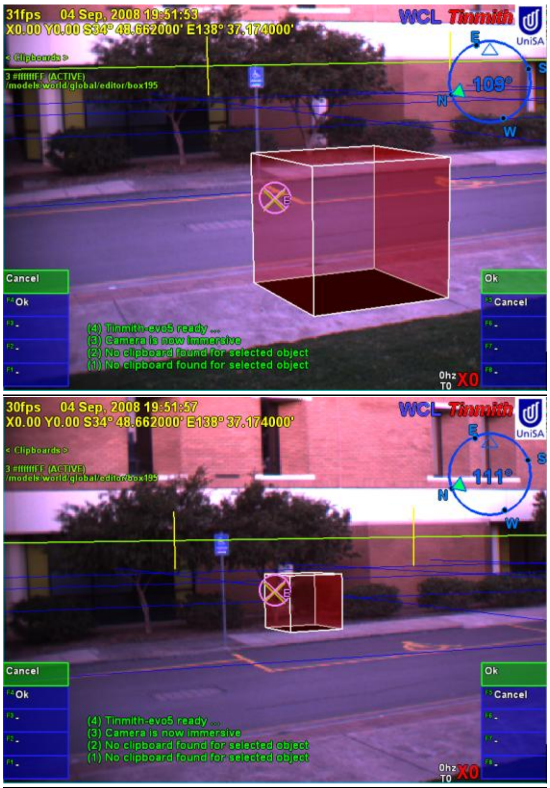

Figure 7: Ground plane translation of a virtual box. Top: before translation. Bottom: after translation.

Figure 7 is showing an example of the ground plane translation, used to move a virtual box across a real world street. After translation, the virtual box remains attached to the ground plane.

The choice of translation axes (X_o and Y_o) in the Ground plane technique is highly suitable for the augmented reality outdoor settings, considering that the manipulation of large graphical outdoor objects, such as houses or trees, often requires changing the distance between the object and the user and not altering the object's altitude. Support for those translations and constraints is lacking in the original image plane technique.

By using several values provided within Tinmith, we can calculate where to place an object on the ground plane. Firstly, to select an object to translate, we project a cursor into the virtual world that intersects with the first object in its path, as mentioned earlier. The distance between this intersection point and the ground plane is recorded (d_0) and always remains constant during the translation process.

The basic process behind the rest of the calculation is to continue to project the cursor into the virtual world until it intersects with the ground plane. At some location along this vector, there will be a point that will be the same distance away from the ground plane (d_t) as the first object intersection point, i.e. $d_t = d_0$ (see Figure 8). We can then translate the object to match up with that point along the vector.

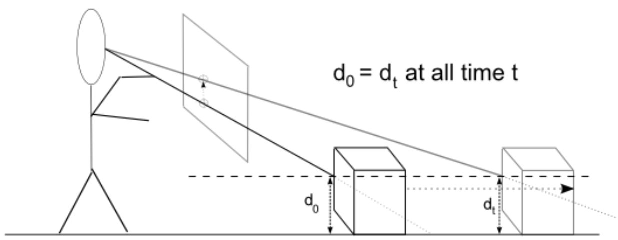

Figure 8: Calculations of ground plane translation

The resulting effect is that the object continues to remain the same distance away from the ground plane as when it was first selected. This can then be adapted for any other plane that the user wants to constrain objects to. The original image plane technique is still available to be used in combination with the Ground plane technique, to support cases of translating objects along the vertical normal axis to the ground plane.

5 Generalised plane techniques

Based on the set of interaction techniques presented above, we have developed the generalised plane technique for object manipulation. The development of the concept is inspired from the Piekarski and Thomas Working Planes (2004) and a current issue with input devices for AR systems. There is currently no known effective input mechanism to fully support 6DOF object manipulation that is suitable for an outdoor AR wearable system. Throughout our investigations into outdoor AR, however, we have discovered that in most situations, it is preferable to perform only 2DOF object manipulation with specific constraints on the third DOF. Translation of large virtual graphical objects, for example, is optimal if the objects are constrained to the ground plane, and only two-dimensional translation of objects on the ground plane is desirable.

The generalised plane technique enables 2DOF object translation on a specific constraint plane, using 2DOF input devices. The steps to perform the technique involve specifying the required plane, then selecting the object using ray casting, and completing the translation of the object to the required location by moving the input device. During the translation, the object will always remain the same distance away from the plane it is constrained to. The Ground plane technique presented in section 4.3 is an implementation of the generalised plane technique with the ground plane as the constraint plane.

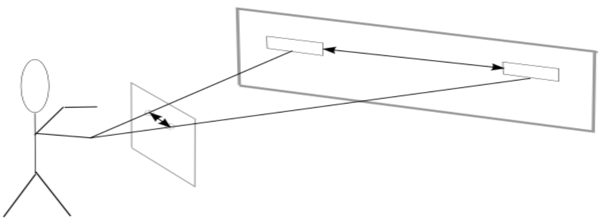

Figure 9: Generalised plane translation

Figure 9 depicts a user translating a rectangular object along a constraint plane using the generalised plane technique. It can be noticed how the movements of the cursor on the display plane are mapped to translation of the objects on the constraint plane, regardless of the angle formed by these planes, to keep the objects constrained to the constraint plane.

Depending on the implementation, methods for selecting or specifying the constraint plane may vary. An existing virtual plane could be selected if suitable; or a new plane could be created from a real world surface by using additional depth sensors, such as laser range finder, to specify at least 3 points on the surface.

Early implementation of the generalised plane translation in Tinmith includes the ability to create a virtual constraint plane on which translations of other objects are performed. The position and the normal vector of the plane are specified by the position and the head orientation of the user. In other words, the plane, represented by a bounded vertical rectangle, is created at the user's feet with the normal in the direction where the user is currently looking; a method by which any arbitrary plane can be created. The user then attaches objects by translating them towards the constraint plane. Once attached, subsequent translations of the objects are only allowed on the 2 dimensions constrained by the plane. Figure 10 shows an example of translating a virtual box on a virtual constraint plane. The plane is created not parallel to the display plane, in order to illustrate the mapping between movements of the cursor on the display plane and translations of the box on the constraint plane.

The generalised plane technique provides more benefits than the original image plane, in various aspects. The ability to choose the constraint plane allows it to perform object translation on various axes, by specifying the constraint plane that contains the required axes. The original image plane technique is restricted to the use of a head relative plane for interaction, which limits the translation to only 2 fixed axes.

Figure 10: Generalised plane translation of a virtual box on a virtual constraint plane (represented by the bounded rectangle) (top) before translation (bottom) after translation

By automatically constraining the object to a specific plane, the generalised plane technique eliminates unwanted object movements, which is often caused by input devices errors. For example, a user with the task of placing street furniture, such as benches, trees or bins, may not desire to place objects above or below the ground plane, In this case, a generalised plane technique with ground plane constraint is most suitable.

The ability to specify any arbitrary constraint plane allows the generalised plane technique to cater for situations where the required object translation is not along the standard X, Y or Z axes. For instance, the problem of moving a window object to a different position on a virtual wall situated at a 45 degree angle from the user's first person view could only be completed with a generalised plane technique, by choosing the virtual wall as the constraint plane.

The applicability of the generalised plane technique for object translation is unlimited. Examples include movements of virtual house objects alongside a real-world street pavement, or positioning a virtual architectural attachment on the side of a real-world building. Translation of small objects in regards to a larger one is also a good usage of the technique; an example includes moving a virtual rubbish bin alongside a building's wall.

6 Conclusion

We have investigated three new manipulation techniques to augment current image plane techniques to perform more natural 3D direct manipulation in an AR environment. We also have presented a new concept of generalised plane translation technique for the constrained translation of 3D objects. These operations are constrained to the ground in different ways, to support the most common manipulation tasks: rotation and scaling along the normal axis to the ground plane, and translation along the ground. Although these techniques are not new, their application in the proper coordinate system allows a more natural interaction for the user over the original image plane technique. The set of techniques presented also forms a basis for the development of the generalised plane technique, which provides a more useful and widely applicable approach to object manipulation.

7 References

Bowman, D 1996, 'Conceptual Design Space-Beyond Walk-through to Immersive Design', *Designing Digital Space*, John Wiley & Sons, p. 225—36.

Bowman, D, Kruijff, E, LaViola, J & Poupyrev, I 2005, *3D User Interfaces, Theory and Practice*, Addison-Wesley.

Bowman, DA & Hodges, LF 1997, 'An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments', *Proceedings of the 1997 Symposium on Interactive 3D graphics*, pp. 35-8.

Buxton, W & Myers, B 1986, 'A study in two-handed input', *Proceedings of the SIGCHI86 conference on Human factors in Computing Systems*, pp. 321-6.

Forsberg, A, Herndon, K & Zeleznik, R 1996, 'Aperture based selection for immersive virtual environments', *Proceedings of the 9th annual ACM symposium on User interface software and technology 1996*, pp. 95-6.

Liang, J & Green, M 1993, 'Geometric modeling using six degrees of freedom input devices', *Proceedings of the 3rd International Conference on CAD and Computer Graphics*, pp. 217-22.

Piekarski, W 2004, 'Interactive 3d Modeling in Outdoor Augmented Reality Worlds', School of Computer and Information Science, University of South Australia.

Piekarski, W & Thomas, BH 2001a, 'ThumbsUp: Integrated Command and Pointer Interactions for Mobile Outdoor Augmented Reality Systems', *Proceedings of the 9th International Conference on Human-Computer Interaction*.

Piekarski, W & Thomas, BH 2001b, 'Tinmith-Metro: New Outdoor Techniques for Creating City Models with an Augmented Reality Wearable Computer', *Proceedings of the 5th International Symposium on Wearable Computers*, pp. 31-8.

Piekarski, W & Thomas, BH 2004, 'Augmented reality working planes: a foundation for action and construction at a distance', *Proceedings of the 3rd IEEE and ACM International Symposium on Mixed and Augmented Reality.*, pp. 162-71.

Pierce, JS, Forsberg, AS, Conway, MJ, Hong, S, Zeleznik, RC & Mine, MR 1997, 'Image plane interaction techniques in 3D immersive environments', *Proceedings of the 1997 symposium on Interactive 3D graphics.*

Poupyrev, I, Billinghurst, M, Weghorst, S & Ichikawa, T 1996, 'The go-go interaction technique: non-linear mapping for direct manipulation in VR', *Proceedings of the 9th annual ACM symposium on User interface software and technology*, pp. 79-80.

Stafford, A, Piekarski, W & Thomas, BH Oct. 2006, 'Implementation of god-like interaction techniques for supporting collaboration between outdoor AR and indoor tabletop users', *Proceedings of the 5th IEEE and ACM International Symposium on Mixed and Augmented Realit.*, pp. 165-72.

Stoakley, R, Conway, MJ & Pausch, R 1995, 'Virtual reality on a WIM: interactive worlds in miniature', *Proceedings of the SIGCHI 1995 conference on Human factors in Computing Systems*, pp. 265-72.

Thomas, B, Close, B, Donoghue, J, Squires, J, De Bondi, P, Morris, M & Piekarski, W 2000, 'ARQuake: an outdoor/indoor augmented reality first personapplication', *Proceedings of the 4th International Symposium on Wearable Computers*, pp. 139-46.