Bringing 2D Interfaces into 3D Worlds

Joseph J LaViola Jr

Ph.D Candidate

Brown University,

Department of Computer Science

Providence, Rhode Island

Lead Consultant and Founder

JJL Interface Consultants, Inc.

http: http://www.cs.brown.edu/people/jjl

email: jjl@cs.brown.edu

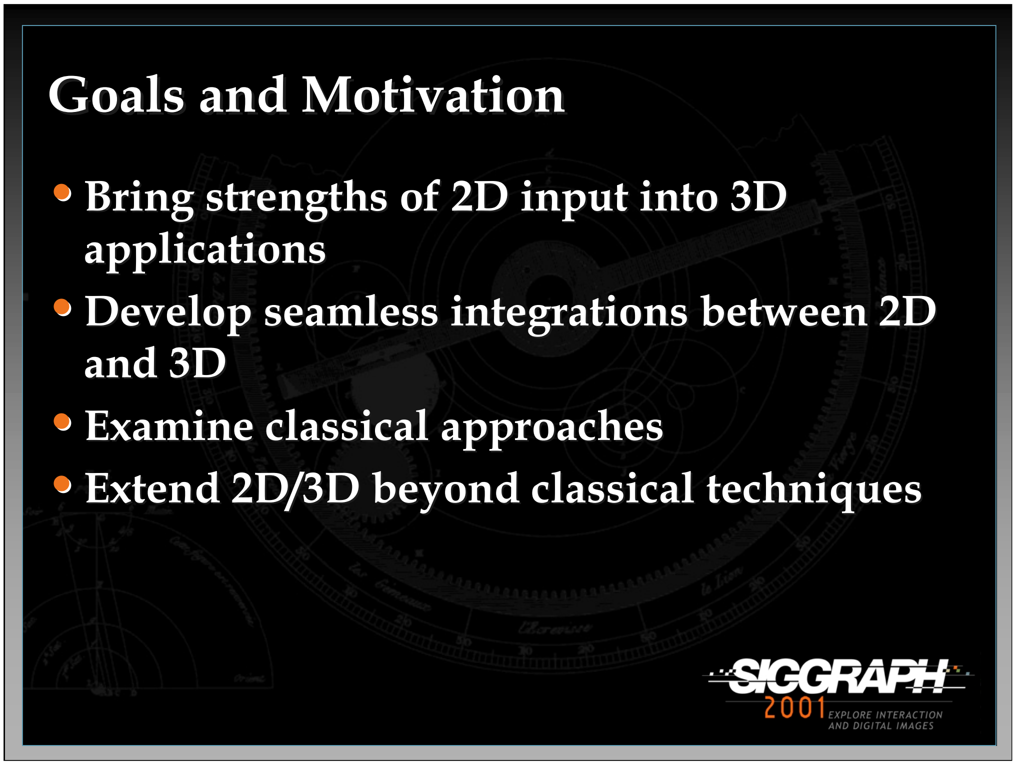

In this lecture, we will discuss the advantages of 2D interaction and how we can use various 2D techniques in 3D applications. We will examine both classical approaches and state of the art research results.

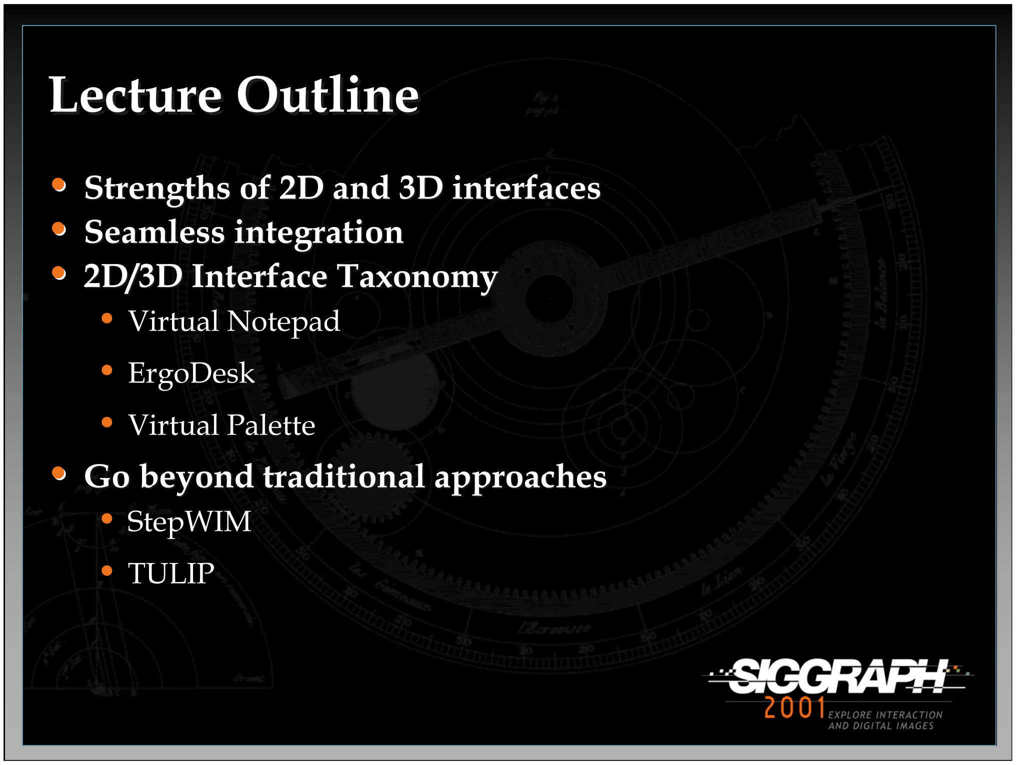

In the first part of the lecture, we will examine the strengths and weaknesses of 2D and 3D interfaces and how they can be seamlessly integrated as one interface to 3D graphics applications and virtual environments. Second, we will briefly develop a 2D/3D interface taxonomy based on previous work in the area and examine some interfaces that fall into this categorization. Finally, we will look at a couple of interfaces that go beyond our taxonomy in the sense that they bring 2D interaction concepts into 3D applications in unconventional ways.

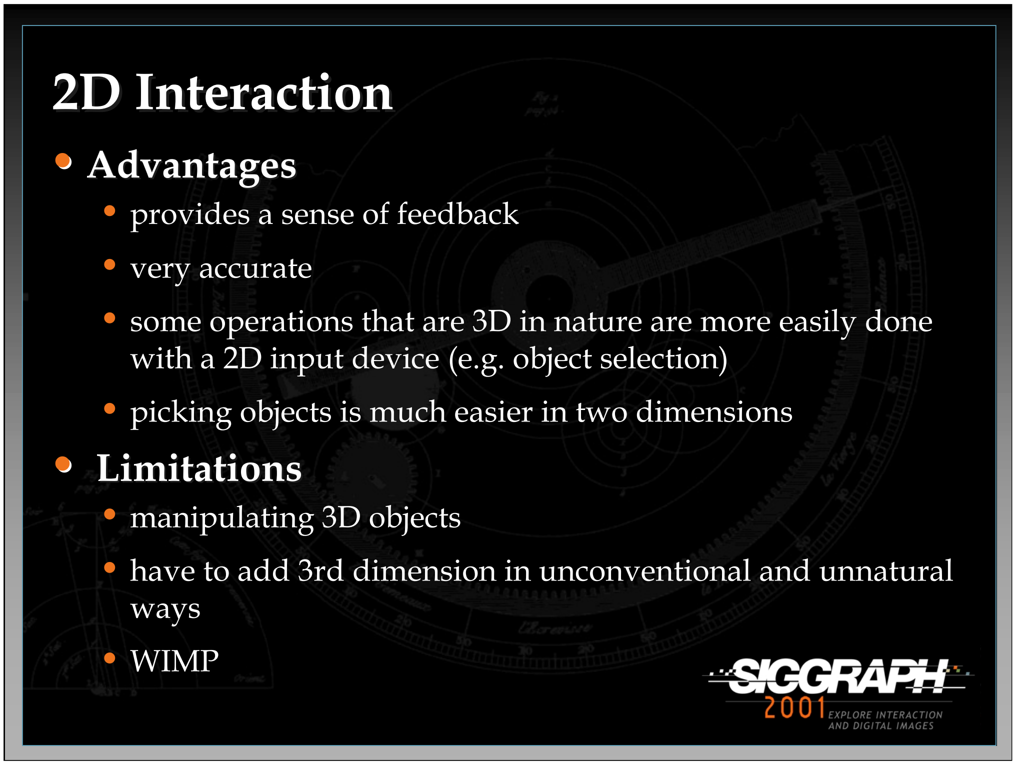

2D interaction techniques have both advantages and disadvantages as shown in the slide.

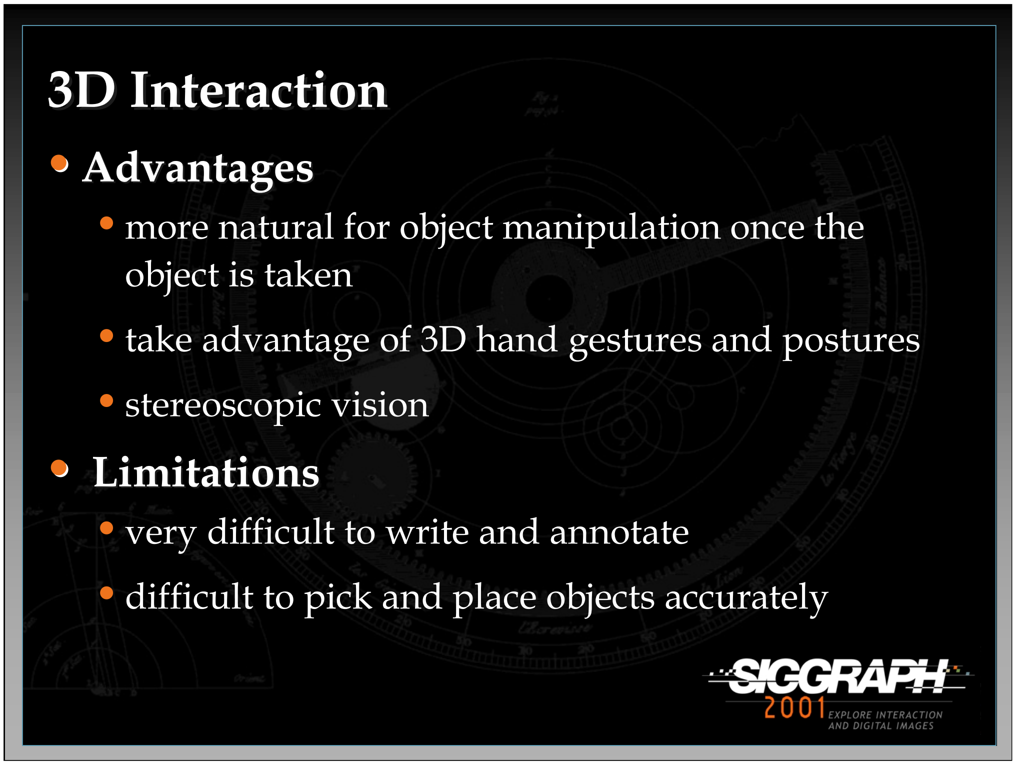

3D interaction techniques have both advantages and disadvantages as shown in the slide.

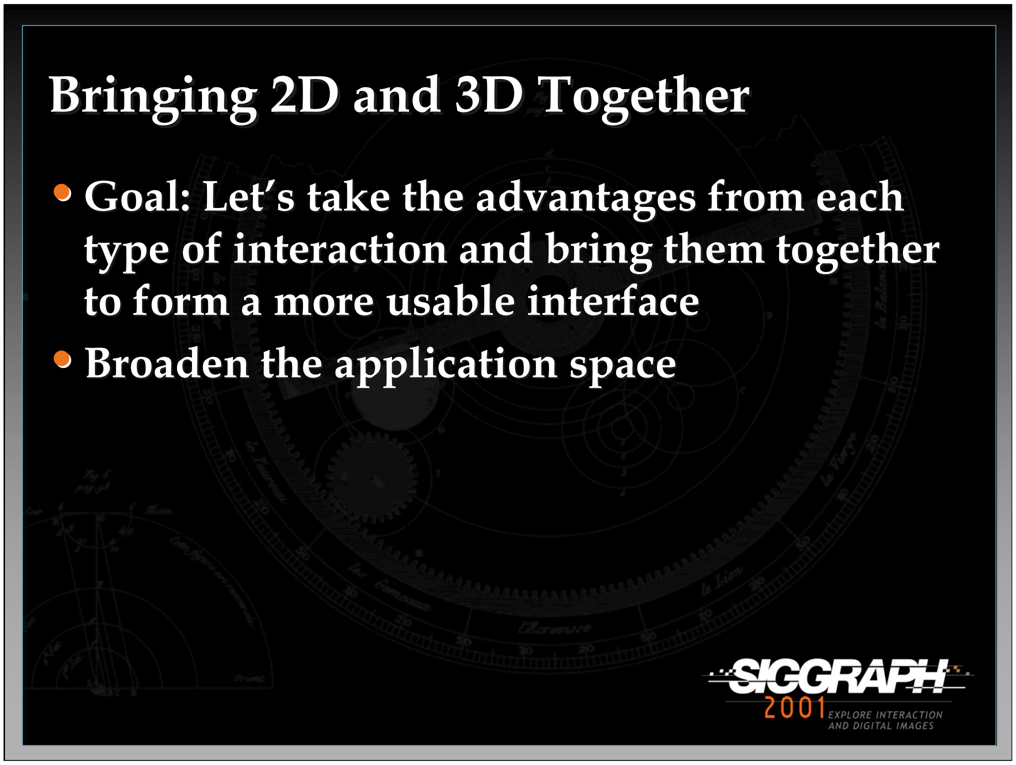

By taking advantage of the benefits of both 2D and 3D interaction techniques and metaphors, we can create interfaces for 3D applications that are easier to use and more intuitive for the user. The key research issue is how to combine these two input styles in a seamless manner and to determine whether a particular task is better suited for either 2D or 3D interaction so we can maximize user performance.

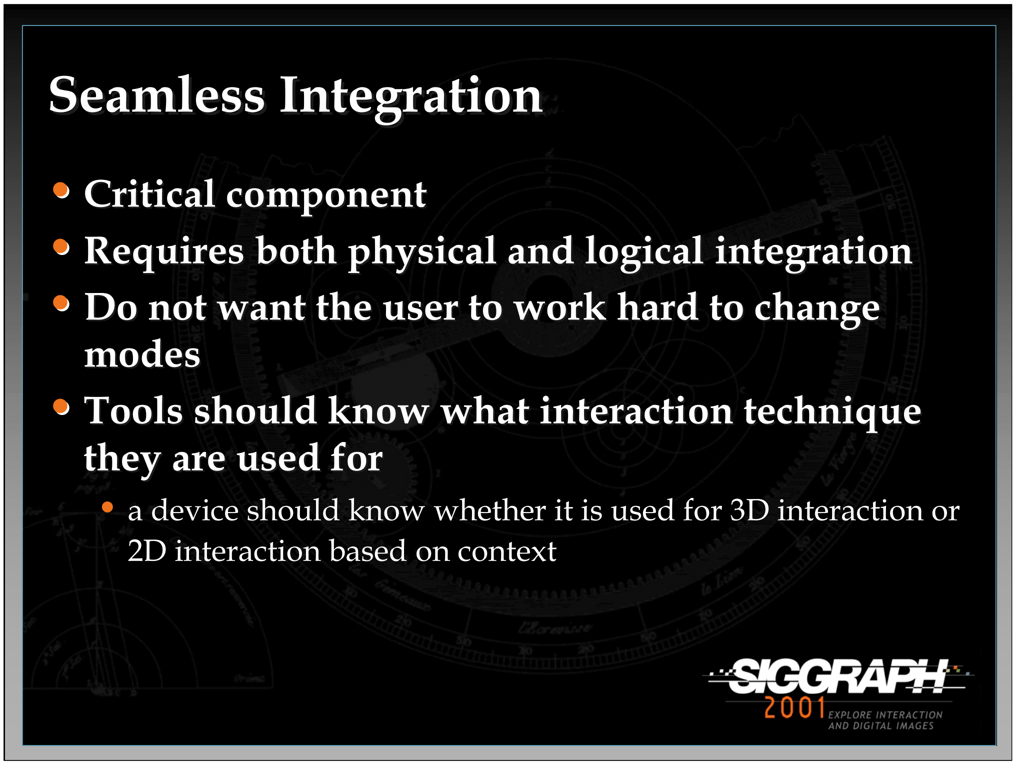

The seamless integration of the 2D and 3D interface techniques in a 3D application is a critical design consideration from both a logical and a physical perspective. Physical integration is important because we do not want to make it difficult for the user to switch between 2D and 3D devices. Logical integration is also important because we want the devices used in the application to know whether they are used for either 2D or 3D interaction. This knowledge helps to reduce the user's cognitive load.

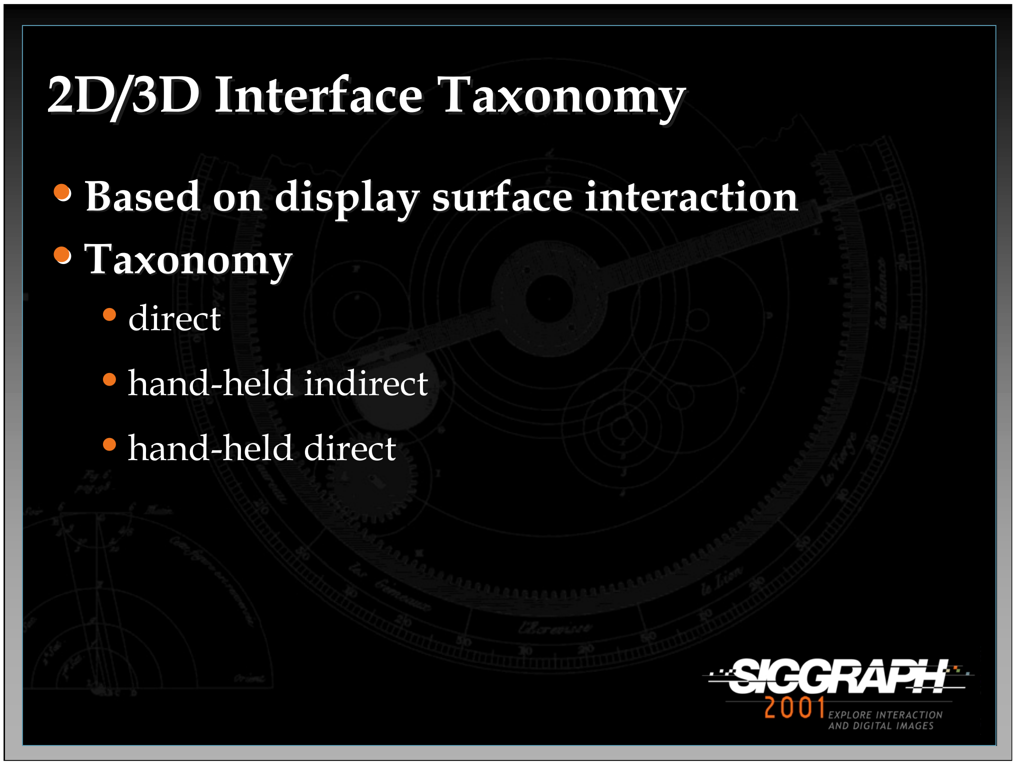

There is a pattern in the 2D interface in 3D application literature which leads to a simple taxonomy based on how the user performs the 2D operations. In general, there has to be a some type of surface with which the user can interact on. The first category in the taxonomy is *direct display surface interaction*. Any display surface with allows the user to perform 2D operations falls into this category. Desk-based displays which allow the user to draw on the display surface are a good example. The second category is *hand-held indirect display surface interaction*. This category includes applications which require the user to hold a pad, whether transparent or opaque, in order to perform 2D operations. A classic example is in virtual environments where users must wear HMDs and cannot see the physical world. A virtual display is presented to the user in the virtual world and is correlated with the physical pad. Finally, the third category is *hand-held direct display surface interaction*. This category includes applications which use hand-held displays or computers that allow the user to interact in 2D on their display surfaces.

In the next few slides we will examine a few examples of 2D/3D interfaces that fall into these categories.

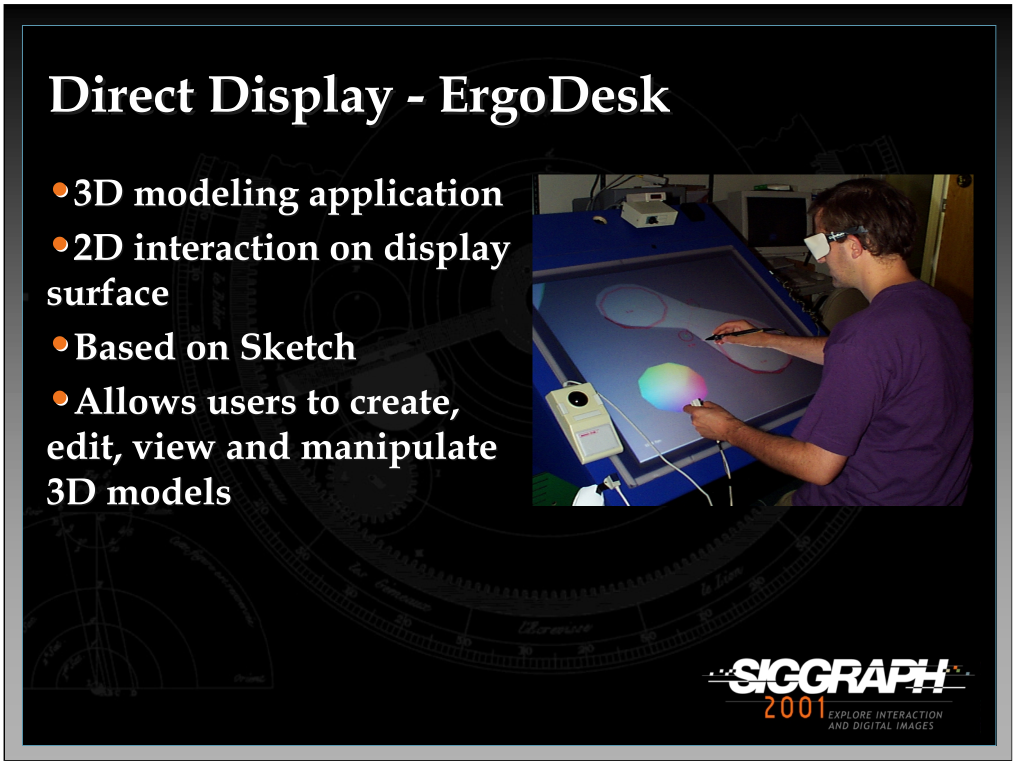

ErgoDesk in a example of a 2D/3D interface where the user interacts in 2D directly on the display surface. The 2D component of the ErgoDesk application is based on the Sketch conceptual modeling system which uses only a three button stylus (no menus or 2D interface widgets are used). Sketch interprets lines drawn by the user on the image plane of a 3D view as operations and parameters. These operations include primitive creation, primitive manipulation, and camera manipulation. Gestures that create primitives provide enough information to select which primitive to create, its dimensions and its place in 3D. Creating a cube, for example, requires the user to draw 3 gesture lines one for each of the principle axes, each line meeting at a single point. The cube is generated with it's length, width, and height corresponding to the three gesture lines and its place in 3D based on the intersection point. Primitives such as cylinders, cones, pyramids, and extrusions can also be instantiated. The primitive manipulation interface allows for automatic object constraint by gesturally drawing a motion constraint over the object before manipulating it. For example, to constrain an object's movement to a given axis, a straight line is drawn indicating what axis to constrain the object to, and when the user moves the object it will only move along that axis. Other gestures constrain objects to move along surfaces, rotate around a given principle axis, or scale and deform to fit a new gesture contour.

-continued on the next page

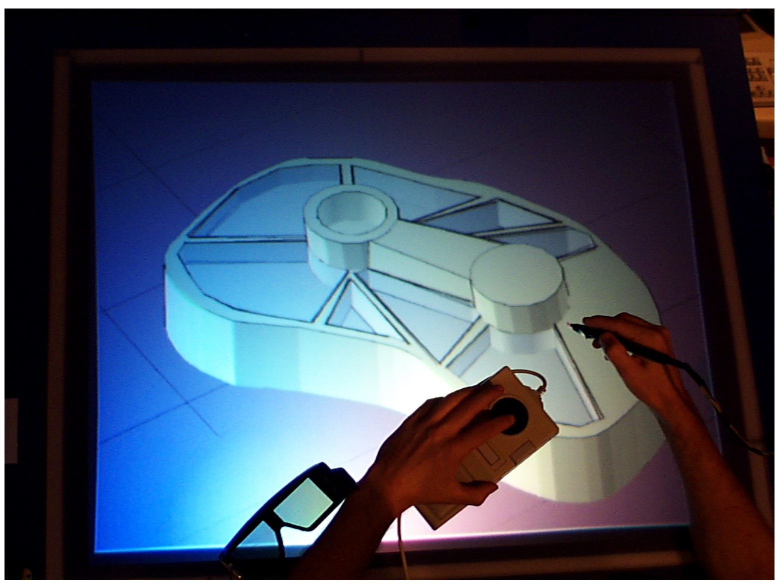

References:

Forsberg, A., LaViola J., and Zeleznik, R. “ErgoDesk: A Framework For Two and Three Dimensional Interaction at the ActiveDesk." In the Proceedings of the Second International Immersive Projection Technology Workshop, Ames, Iowa, May 11-12, 1998.

Zeleznik, R.C., Herndon, K., Hughes, J. (1996) "Sketch: An Interface for Sketching 3D Scenes." Proceedings of SIGGRAPH'96, 163-170.

The Virtual Notepad is an example of a 2D/3D interface where users cannot physically see the 2D device since they are wearing an HMD. The 2D device is tracked so a graphical representation of it is present in the virtual environment.

References:

Poupyrev, I., Tomokazu, N., Weghorst, S., “Virtual Notepad: Handwriting in Immersive VR”. IEEE VRAIS'98, 126-132, 1998.

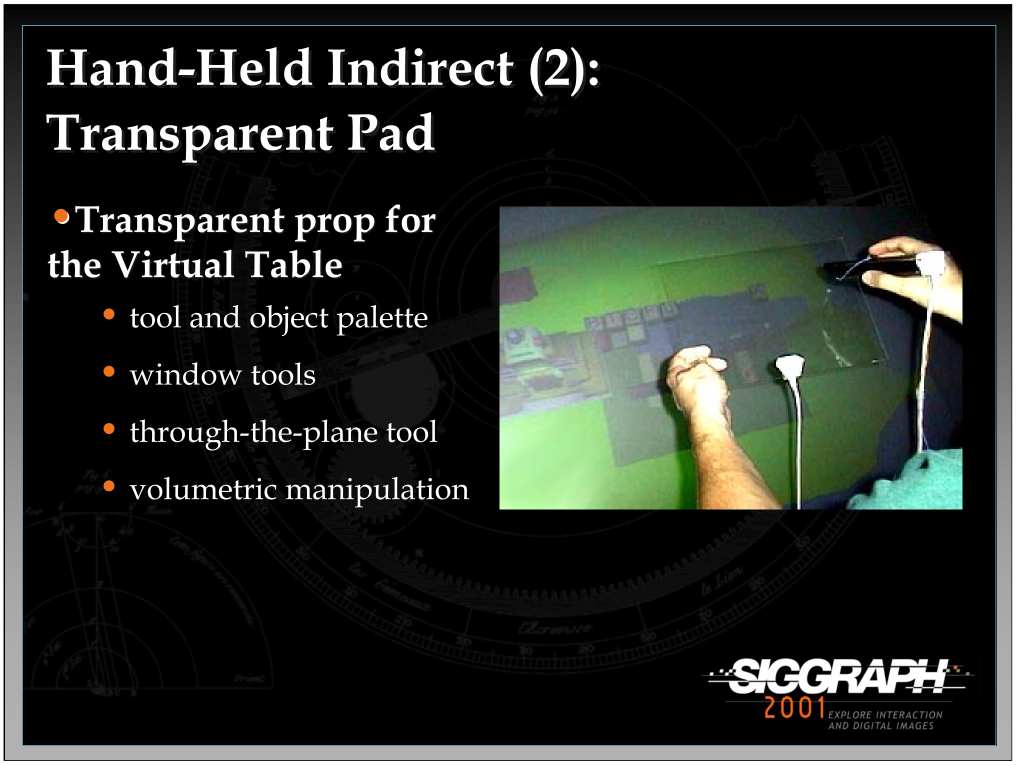

The Transparent Pad is another example of a 2D/3D interface which utilizes a hand-held pad to perform 2D operations. In this case, the pad is transparent. The pad is tracked and graphics are projected on the primary display but appear as if they are on the surface and even above the pad.

References:

Schmalsteig, Dieter, L. Miguel Encarcacao, Zsolt Szalavari. “Using Transparent Props For Interaction with The Virtual Table.”, In Proceedings of the 1999 ACM Symposium on Interactive 3D Graphics, 147-154, 1999.

Coquillart, S. and G. Wesche. “The Virtual Palette and the Virtual Remote Control Panel: A Device and Interaction Paradigm for the Responsive Workbench.” IEEE VR’99, 213-217, 1999.

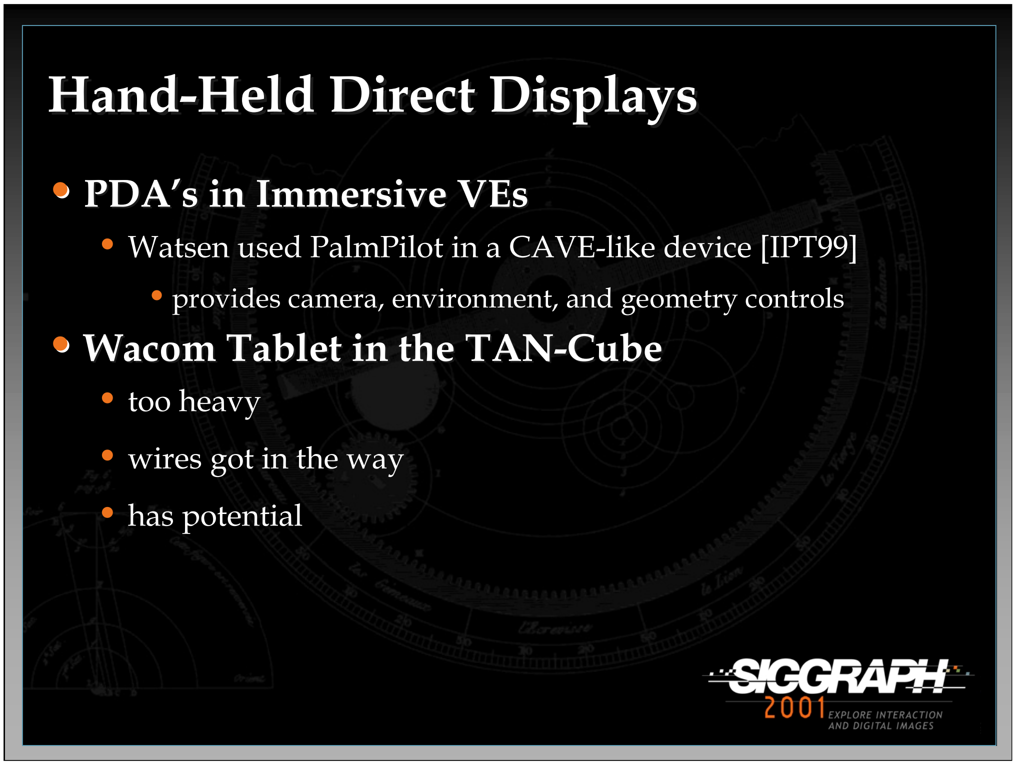

There are only a few reported cases of using hand-held direct displays for performing 2D operations in 3D applications. One of the first used a PalmPilot to control camera, environmental, and geometrical parameters in a virtual environment. With better wireless technology and more light weight, portable display devices hand-held direct display interaction should gain more prominence in 2D/3D interface research.

References:

Watsen, Kent, Rudy Darken, and Michael Capps. "A Handheld Computer as an Interaction Device to a Virtual Environment." Proceedings of the Immersive Projection Technology Workshop, Stuttgart, Germany, May 1999.

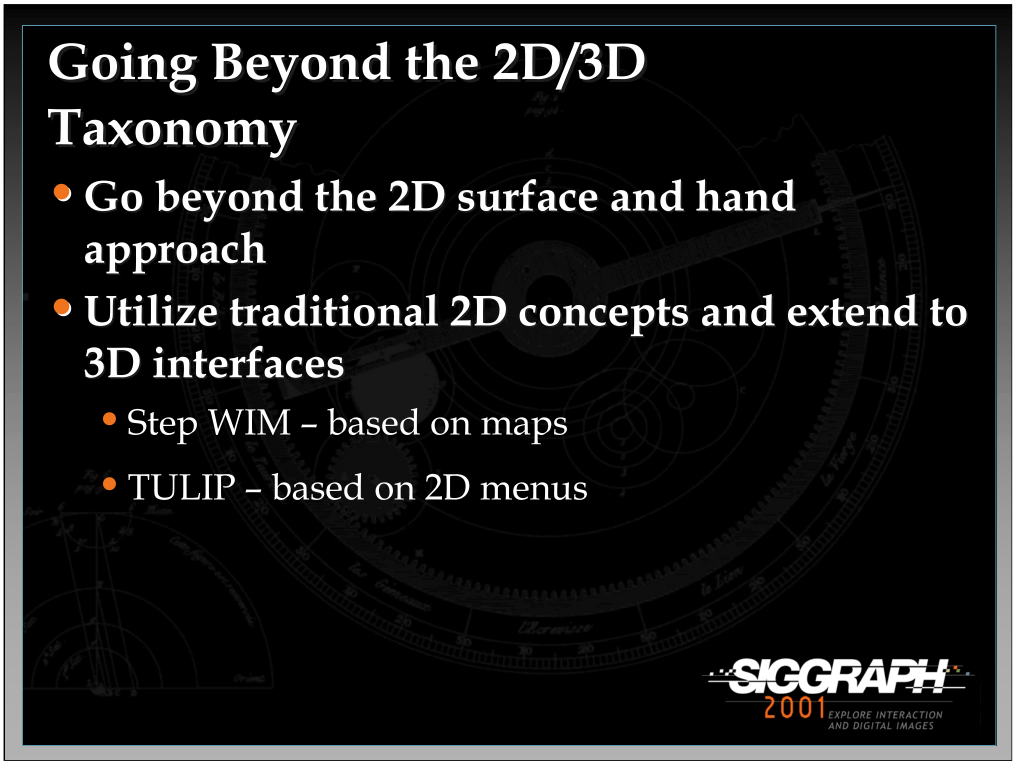

Our 2D/3D interface taxonomy has a common theme in that the 2D interface component is derived from a surface and hand metaphor. In each category, the user interacts with a display surface of some kind with his/her hands to perform various 2D and even 3D operations. An interesting research question is whether we can go beyond our taxonomy and utilize traditional 2D interface concepts without the constraint of the 2D surface and hand approach. Of course the answer is yes; otherwise the lecture would be over. In the next few slides, we will look at two examples which go beyond our 2D/3D classification. The first is the Step WIM, which keeps the 2D surface concept but is a hands-free interface which utilizes body gestures and the user's feet. The second is TULIP, which allows the user to interact with the hands but removes the 2D surface constraint allowing for interaction in 3D space.

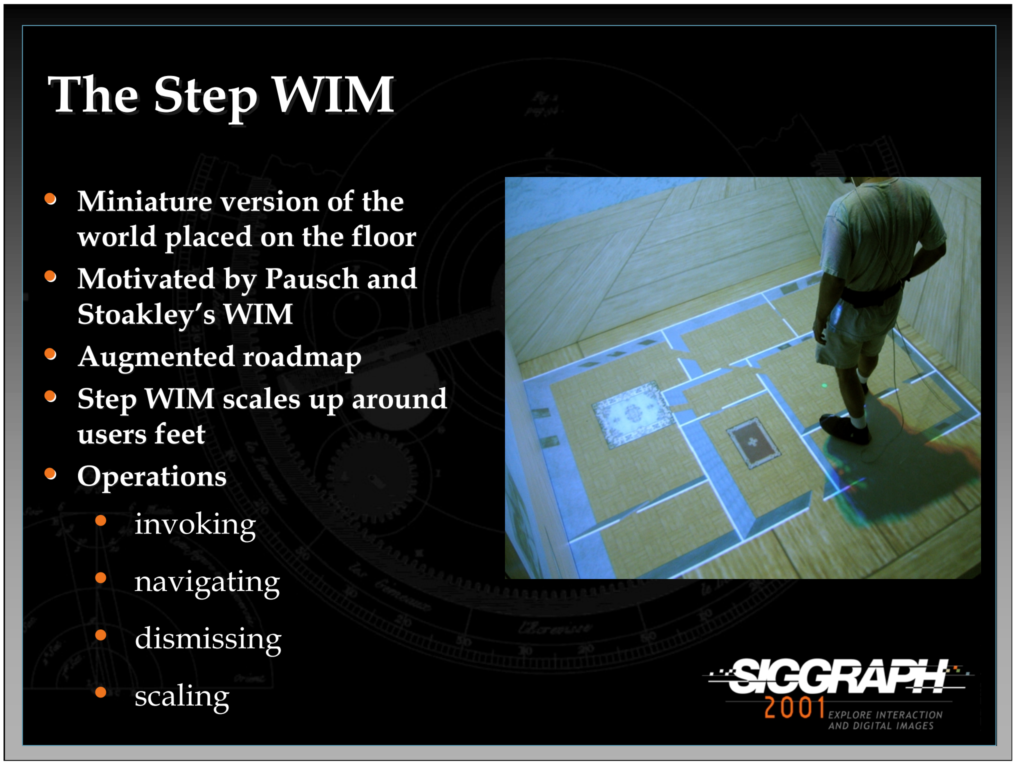

The Step WIM is a interaction widget for quickly navigating through a virtual environment. It is a miniature version of the world placed underneath the user's feet and acts as an augmented roadmap. The user can either walk around the Step WIM to get a better understanding of the virtual world or navigate to a specific place by simply walking to a desired location in the WIM and invoking a scaling command, causing the Step WIM to animate, scaling up around the user's feet, thereby seamlessly transporting the user to the specified location.

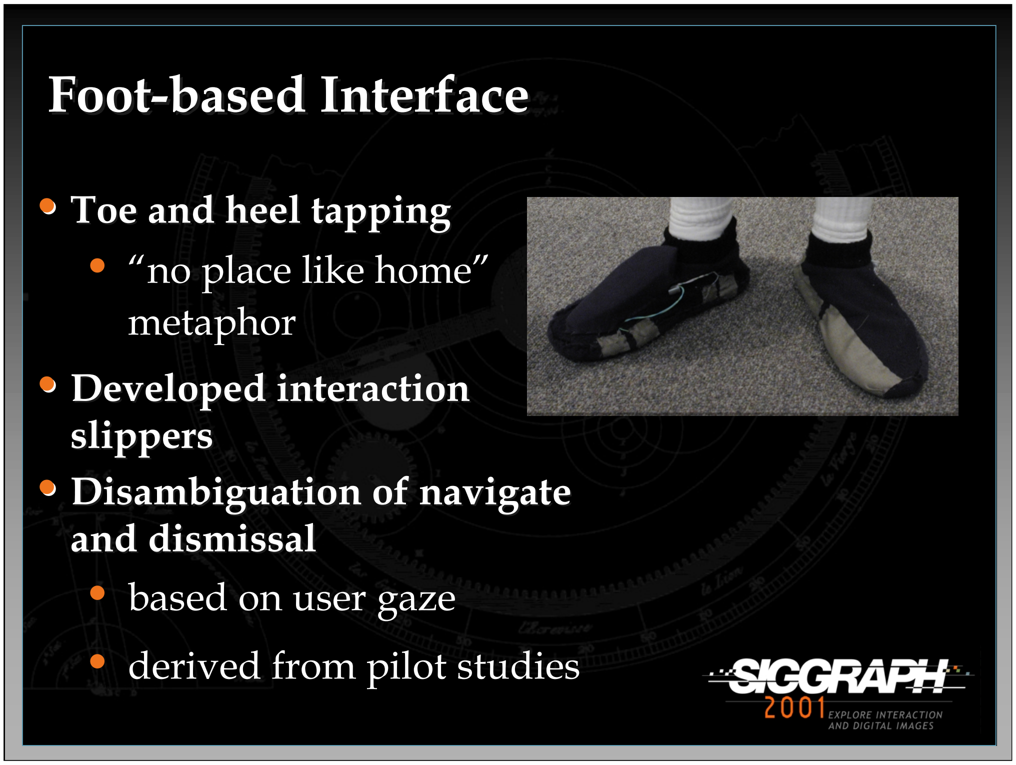

In order to invoke, navigate and dismiss the Step WIM, users can wear a pair of interaction slippers (slippers with an imbedded wireless mouse) which gives them the ability to perform toe and heel tapping. To invoke the Step WIM, users simply tap their toes together. Once the Step WIM is active, another toe tap will transport the user to a new location or dismiss the widget without navigation. Based on pilot studies, users tended to look down at the Step WIM when they wanted to navigate so disambiguation of the navigation and dismissal tasks are based on user gaze. This approach allows for two distinct operations to be mapped to one button press.

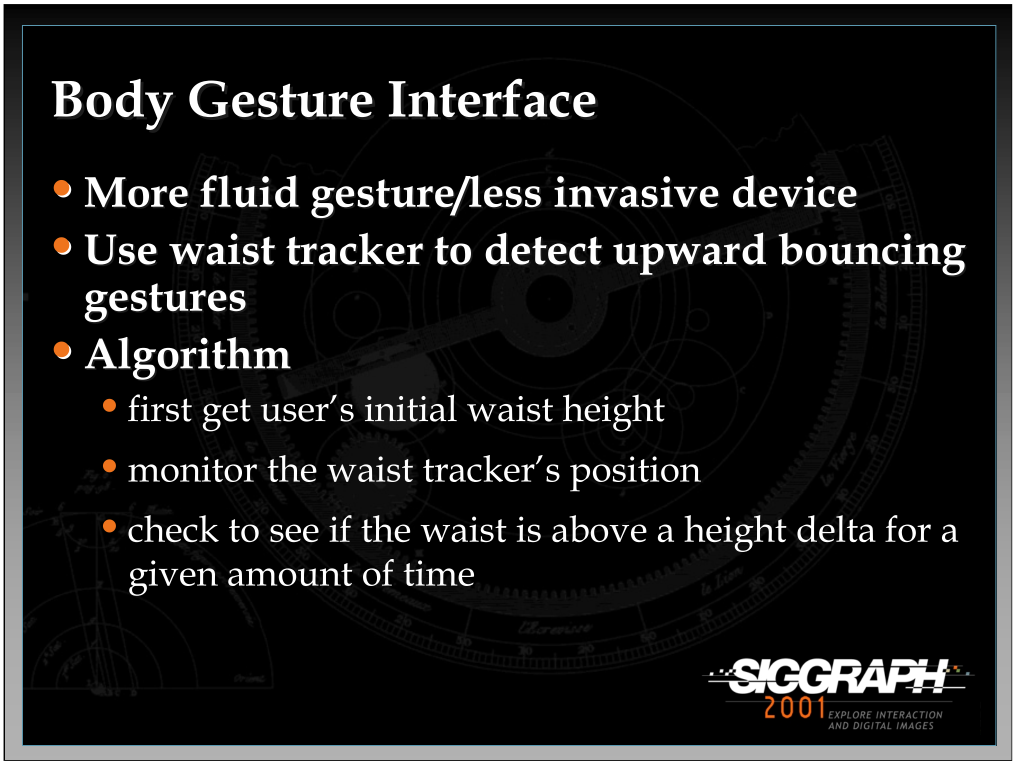

Another way to use the Step WIM is based on body gestures. Using a waist tracker, a simple gesture recognition algorithm detects upward bouncing movements. When the user performs a bouncing gesture, the Step WIM is activated. Another bouncing gesture dismisses the Step WIM or transports the user to a new location, once again depending on user gaze.

Since the virtual environment may be too large to fit within the user's walking area, the Step WIM can be scaled to varying sizes. This scaling implicitly provides different levels of detail.

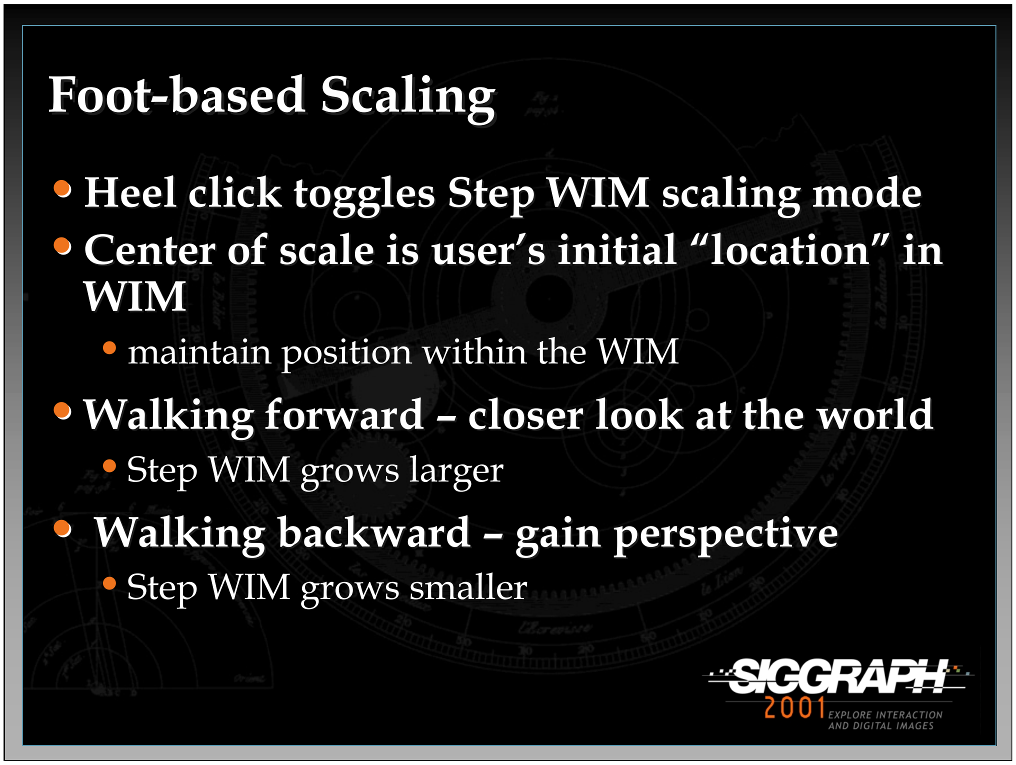

When using the interaction slippers, a heel click toggles in and out of Step WIM scaling mode. When the user walks forward from his or her initial position (the position defined with the heel click), as if to take a closer look at the world, the Step WIM grows larger. When the user walks backward from his or her initial position, as if to gain perspective, the Step WIM grows smaller.

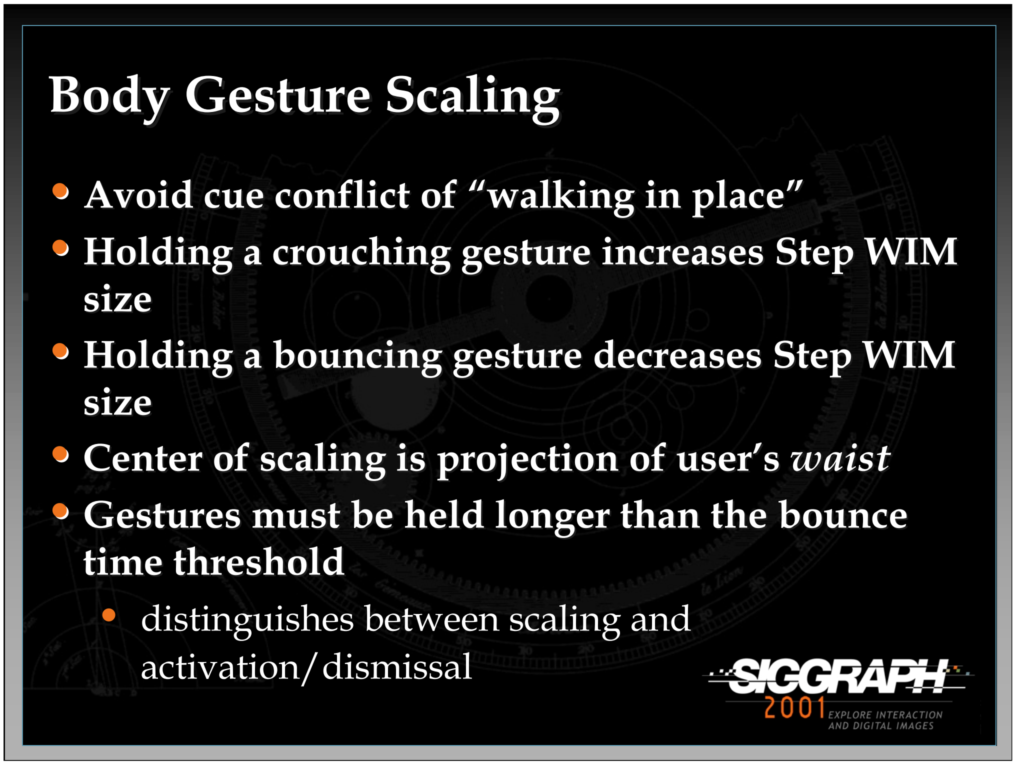

With the body gesture-based interface, holding a crouching gesture increases the Step WIM size, while holding a bouncing gesture decreases the size of the Step WIM.

In general, the Step WIM represents a 2D concept (the concept of a map) that has been incorporated into a 3D interface for navigation. This interface technique goes beyond our 2D/3D taxonomy by removing the hand component of the 2D surface and hand metaphor. More information on the Step WIM can be found the in paper entitled, "Hands-Free Multi-Scale Navigation in Virtual Environments" included in the papers section of the course notes.

References:

LaViola, J., Acevedo, D., Keefe, D., and Zeleznik R. "Hands-Free Multi-Scale Navigation in Virtual Environments", In the Proceedings of the 2001 Symposium on Interactive 3D Graphics, 9-15, March 2001.

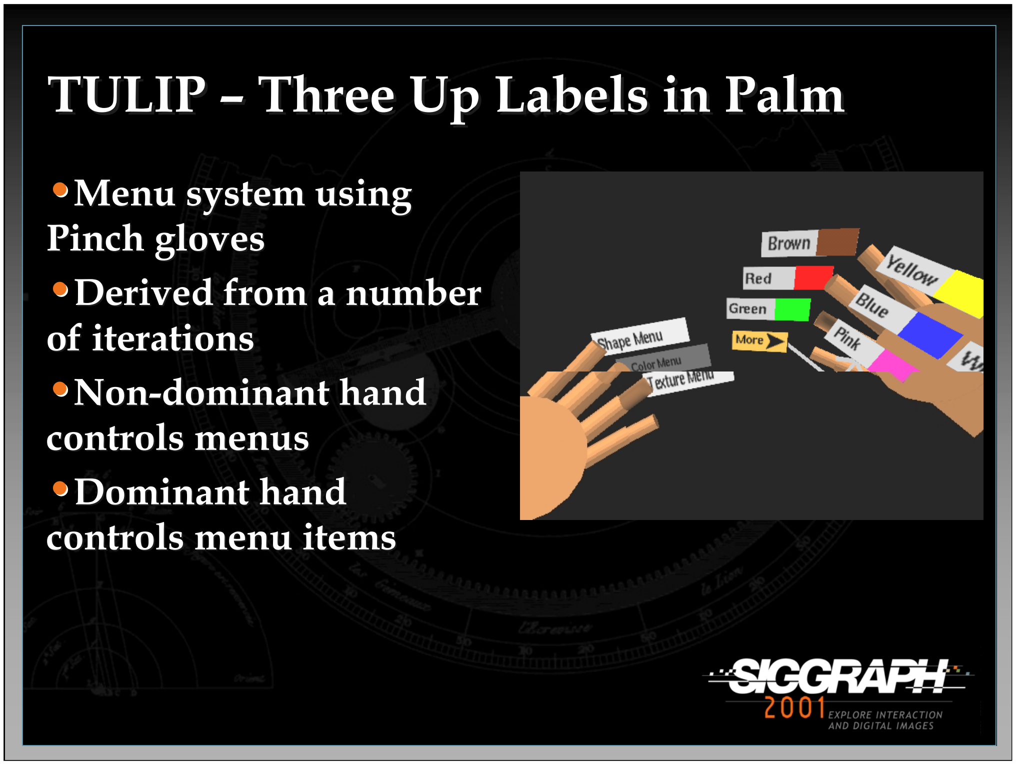

TULIP is an interaction tool which takes the concept of the 2D menus and brings it into a 3D interface. By utilizing Pinch gloves, the user has a menu system attached to the hands which is activated by pinching postures. The non-dominant hand holds menu choices and the dominant hand holds three menu items at a time which correspond to thumb to index, thumb to middle, and thumb to ring contacts. A "more" option is displayed on the pinkie which points to another set of three menu items shown in the palm of the hand. These three menu items will become available if the users makes a thumb to pinkie contact.

Although other 2D menu systems have been developed for 3D applications such as pull-down and body-centered menus, the TULIP system keeps the menu in the user's hands rather than in the virtual space or by some other part of the body.

References:

Jacoby, R. and S. Ellis. "Using Virtual Menus in a Virtual Environment", In SPIE: Virtual Data Interpretation, 1992.

Mine, M., F. Brooks, and C. Sequin. "Moving Objects in Space: Exploiting Proprioception in Virtual Environment Interaction", ACM SIGGRAPH'97, 19-26, 1997.

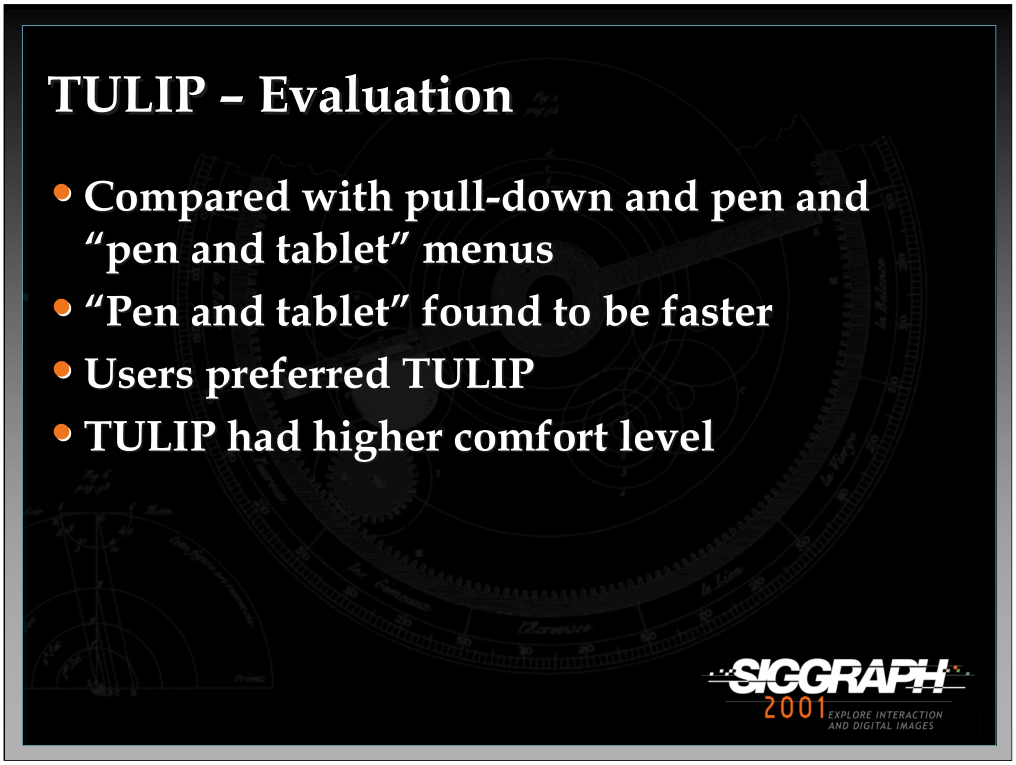

A user evaluation which compared TULIP with pull-down and "pen and tablet" menus was conducted to test TULIP's ease of learning, efficiency and comfort. The results of the study indicate that although "pen and tablet" interaction was faster than TULIP, more users preferred the TULIP system and found it to be more comfortable to use than the other two menu techniques. More details on the user evaluation and the design of TULIP can be found in the paper, "Design and Evaluation of Menu Systems for Immersive Virtual Environments", found in the papers section of the course notes.

References:

Bowman, D. and C. Wingrave, "Design and Evaluation of Menu Systems for Immersive Virtual Environments", Proceedings of IEEE Virtual Reality 2001, 149-156, 2001.

Lindeman, R., J. Silbert, and J. Hahn, "Hand-Held Windows: Towards Effective 2D Interaction in Immersive Virtual Environments", Proceedings of IEEE Virtual Reality'99, 205-212, 1999.

Conclusions

- 2D interface metaphors can be critical in 3D applications
- Seamless integration of 2D and 3D components is essential
- Make the tools that the user needs intelligent
- Important to find lightweight solutions when using hand-held devices
- Field is still in its infancy

We must continue to explore how 2D interface concepts and components can fit into 3D interfaces and virtual environments. They are powerful tools when used properly and can greatly increase productivity for users.