SpikeNav: Using Stylus Tilt in Three-Dimensional Navigation

Robert Bridson
University of British Columbia & Exotic Matter AB
201-2366 Main Mall, Dept. Computer Science
Vancouver, V6T 1Z4, Canada
rbridson@cs.ubc.ca

ABSTRACT

We propose a new method for 3D navigation using a commodity stylus and tablet already in common use among 3D artists. SpikeNav exploits tilt sensitivity to provide a more direct interaction style. Conceptually the stylus is turned into a spike which the user can jab into the model by contact with the tablet; moving the stylus drags the model around and tilting it rotates it naturally, making the stylus a temporary handle for integral direct manipulation. Pressure triggers a mode switch between rotating and dollying along the view direction with tilt, conceptually using the stylus as a lever instead.

Keywords: pen input, interaction techniques.

INTRODUCTION: THE PROBLEM

Navigation in 3D is a ubiquitous task in visual effects and related computer graphics work. Commodity tablet/stylus input devices are similarly common in film studios at artist workstations. However, apart from painting or sculting applications, the additional channels of the stylus (tilt, pressure) are unused: often the stylus is used simply as a mouse. This poster argues this is a great opportunity, in particular for 3D navigation. The proposed method, SpikeNav, exploits tilt sensitivity to provide direct physical 3D interaction with standard hardware.

The accompanying video illustrates this navigation in the context of visualizing complex simulation data generated by physics-based animation software, e.g. the shape of a cloud of smoke. Developing an effective simulation for a shot often involves detailed study of this output to determine parameters or debug the underlying code: effective navigation and interactive visualization is a critical tool. At least two styles of navigation must be supported in this application: global viewing to get a sense of general motion, alternated with zooming in on specific details in interesting or problematic regions, viewed from all angles in both cases.

UIST'09, October 4–7, 2009, Victoria, BC, Canada

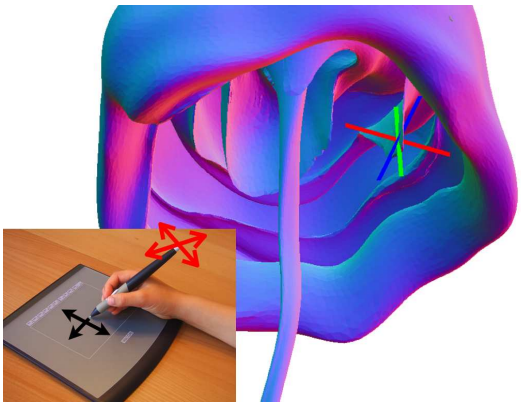

Figure 1: SpikeNav uses both translation and tilt (4D) to control viewpoint in a 3d scene, treating the stylus as a handle stuck into a selected point on the surface. Here a complex smoke surface is being explored for problems, with the stylus's current anchor point highlighted by 3D cross-hairs.

RELATED WORK

The de facto standard navigation for the target audience is the Autodesk Maya approach, mapping 2D mouse motion to a 5D cinematic camera model (3D translation, orbiting rotation around the scene's vertical y-axis and around the camera's x-axis; this preserves vertical orientation, and makes rotations order-independent, but at the cost of less intuitive control when looking straight up or down). Different mouse buttons select whether translation in the film plane, orbiting rotation, or dollying along the view axis are controlled. This work uses the same camera model, but a significantly different metaphor for controlling it. Many other approaches to mapping 2D mouse input to 3D rotations have been explored, such as ARCBALL [6].

Several authors have looked at uses of tilt information outside of the traditional brush metaphor: e.g. 2D cursor appearance [7] or menu selection [8]. Xin et al. compared speed and precision of 2D widget control between tilt and pressure pen input [10]. Bi et al. explored using the longitudinal roll of a pen as an additional input channel for novel tasks [2], though we did not consider roll as it is not as commonly available as tilt. Evans et al. [3] provided an early example of mapping pen input to 3D control, but focused on rotating around a single axis without tilt sensitivity.

The Rockin'Mouse [1] is most similar to this poster: Balakrishnan et al. augmented a mouse with tilt control, enabling notably faster performance for some 3D tasks—and further provided arguments for working with devices such as mice (and, we argue, pens) that also perform well for the common 2D tasks such as menu selection that 3D artists constantly need to perform. However, unlike the Rockin'Mouse work, SpikeNav uses standard hardware, provides a more direct physical metaphor (the stylus, once jabbed into the model, acts as a natural handle to move and rotate it), and tackles orientation as well as translation control.

SpikeNav also is related to the Scene-in-Hand metaphor [9], as the stylus essentially becomes a handle on the scene upon contact with the tablet, and by extension can be used to directly translate and rotate the scene.

MODELING THE STYLUS AS A SPIKE

As mentioned, we use the standard 5D cinematic camera model. When the user touches the stylus to the tablet (a Wacom Intuos2 in the video) SpikeNav finds the 3D point in the scene corresponding to the cursor screen location: this is the "anchor" point. The user has conceptually inserted a spike (or fork) into the scene here. The stylus is now a handle to directly manipulate the scene relative to the camera.

In the regular mode, motion of the stylus point on the tablet is mapped to translation in the film plane of the camera, always keeping the anchor under the cursor and at the same camera depth. Tilting controls rotation of the spike and attached model: tilt along tablet-x rotates the scene around the world's vertical axis, and along tablet-y around the camera's horizontal axis.

Dollying mode is triggered when stylus pressure exceeds a threshold (say, 80%) and ends when pressure goes below a more moderate level (say, 50%), inspired by Li et al.'s work on mode switching with pens [4]. Dollying along the view direction is controlled with one tilt axis: south-west to north-east for right-handed users, consistent with a physical metaphor of applying a lever on the right-hand or top side of the scene, and would be flipped for left-handed users. Change in tilt is converted to an exponential factor multiplying the current camera depth of the anchor, similar to Mackinlay et al.'s Point-of-Interest navigation mode [5].

LIMITATIONS AND FUTURE WORK

SpikeNav shows that a more direct physical style of navigation, previously associated with specialized 3D input devices, may be achievable with a commodity device already commonly used by 3D artists. It currently is just at the prototype stage and requires more study. Some probable limitations to investigate include limits on range of tilt due to grip posture—is excessive "clutching" required, and can this be remedied if so?—questions of precision in tilting, whether the lever/dollying metaphor is natural enough, and—given that tablet use is common but not universal in the industry—a fallback mode for plain mouse input. Finally, it would be potentially very powerful to extend SpikeNav with control over time to view animation data, perhaps mapping the horizontal tilt axis to time and the vertical axis to dollying in the alternate mode.

ACKNOWLEDGMENTS

This work was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada. We also like to thank Karyn Moffatt for advice and guidance on the literature, and Tyson Brochu for providing example simulation data.

REFERENCES

1. Ravin Balakrishnan, Thomas Baudel, Gordon Kurtenbach, and George Fitzmaurice. The Rockin'Mouse: integral 3D manipulation on a plane. In *Proc. ACM CHI*, pages 311–318, 1997.

2. Xiaojun Bi, Tomer Moscovich, Gonzalo Ramos, Ravin Balakrishnan, and Ken Hinckley. An exploration of pen rolling for pen-based interaction. In *Proc. ACM UIST*, pages 191–200, 2008.

3. Kenneth B. Evans, Peter P. Tanner, and Marceli Wein. Tablet-based valuators that provide one, two, or three degrees of freedom. In *Proc. ACM SIGGRAPH*, pages 91–97, 1981.

4. Yang Li, Ken Hinckley, Zhiwei Guan, and James A. Landay. Experimental analysis of mode switching techniques in pen-based user interfaces. In *Proc. ACM CHI*, pages 461–470, 2005.

5. Jock D. Mackinlay, Stuart K. Card, and George G. Robertson. Rapid controlled movement through a virtual 3D workspace. In *Proc. ACM SIGGRAPH*, pages 171–176, 1990.

6. Ken Shoemake. ARCBALL: a user interface for specifying three-dimensional orientation using a mouse. In *Proc. Graphics Interface*, pages 151–156, 1992.

7. Feng Tian, Xiang Ao, Hongan Wang, Vidya Setlur, and Guozhong Dai. The tilt cursor: enhancing stimulus-response compatibility by providing 3D orientation cue of pen. In *Proc. ACM CHI*, pages 303–306, 2007.

8. Feng Tian, Lishuang Xu, Hongan Wang, Xiaolong Zhang, Yuanyuan Liu, Vidya Setlur, and Guozhong Dai. Tilt menu: using the 3D orientation information of pen devices to extend the selection capability of pen-based user interfaces. In *Proc. ACM CHI*, pages 1371–1380, 2008.

9. Colin Ware and Steven Osborne. Exploration and virtual camera control in virtual three dimensional environments. In *Proc. Symp. Interactive 3D Graphics*, pages 175–183, 1990.

10. Yizhong Xin, Xiangshi Ren, and Dawei Li. A comparison of pen pressure and tilt in precision parameter manipulation. *Intl. Conf. Computer Science and Software Engineering*, 2:1070–1073, 2008.