Virtual Reality for Palmtop Computers

GEORGE W. FITZMAURICE, SHUMIN ZHAI, and MARK H. CHIGNELL
University of Toronto

We are exploring how virtual reality theories can be applied toward palmtop computers. In our prototype, called the Chameleon, a small 4-inch hand-held monitor acts as a palmtop computer with the capabilities of a Silicon graphics workstation. A 6D input device and a response button are attached to the small monitor to detect user gestures and input selections for issuing commands. An experiment was conducted to evaluate our design and to see how well depth could be perceived in the small screen compared to a large 21-inch screen, and the extent to which movement of the small display (in a palmtop virtual reality condition) could improve depth perception. Results show that with very little training, perception of depth in the palmtop virtual reality condition is about as good as corresponding depth perception in a large (but static) display. Variations to the initial design are also discussed, along with issues to be explored in future research. Our research suggests that palmtop virtual reality may support effective navigation and search and retrieval, in rich and portable information spaces.

Categories and Subject Descriptors: H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*retrieval models*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*input devices and strategies*; *interaction styles*; I.3.1 [**Computer Graphics**]: Hardware Architecture—*three-dimensional displays*; I.3.6 [**Computer Graphics**]: Methodology and Techniques—*interaction techniques*; I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*virtual reality*

General Terms: Design, Human Factors

Additional Key Words and Phrases: 3D control and display, palmtop computers, virtual reality

1. MOTIVATION

Portable hand-held devices (palmtops) are emerging in the computer industry but suffer from a variety of input and output problems due to their compact size. One primary problem is their relatively small, claustrophobic screen size. Palmtop computers are apparently limited in the functionality they can offer based on their small screen size; however, even larger monitors may have "real estate" problems as large amounts of information are being browsed or manipulated. Many users are finding that their 19-inch or larger computer monitors do not provide enough viewing space for some tasks.

Authors' addresses: G. Fitzmaurice, Department of Computer Science, University of Toronto, CSRI, 6 King's College Road, Toronto, Ontario, Canada, M5S 1A4; email: gf@dgp.toronto.edu; S. Zhai and M. H. Chignell, Department of Industrial Engineering, University of Toronto, 4 Taddle Creek Road, Toronto, Ontario, Canada, M5S 1A4; email: {zhai; chignell}@ie.utoronto.ca

The computer screen generally has a preeminent role in displaying feedback to the user. Thus, improving the communication bandwidth of the computer display is an important issue. One strategy for increasing bandwidth is to use 3D displays. Providing the user with a means of navigating through a 3D display creates a type of virtual reality (VR) [Bishop and Fuchs 1992; Ellis 1991; Kruger 1991].

A minimal 3D display consists of a flat 2D display panel with an image rendered in 3D. Objects are drawn with a sense of depth using techniques such as occlusion, perspective, and shading, given one or more light sources. More sophisticated 3D displays adjust their scene depending on a user's physical position and viewpoint. For example, a fishtank VR system [Arthur et al. 1993, Ware and Slipp 1991] tracks a user's head position and dynamically renders the appropriate scene given the current user viewpoint on a large, stationary computer monitor. Alternatively, there are a set of head-mounted display (HMD) systems [Sutherland 1968] which position a small monitor directly in front of one or both eyes, similar to a pair of goggles. As the user moves his or her head, the displays are updated to reflect the new viewpoint. The specific variation of this design developed by Sutherland allows users to see both the computer-generated data and their physical surroundings simultaneously.

Given the trends in portable computing, we are exploring how palmtop computers can be equipped with a 3D display. Our palmtop-based solution has the following four design goals. First, users interact with one or more hand-held mobile displays which are constantly aware of their spatial position and orientation within a 3D workspace. Second, the displays are high-fidelity color monitors capable of presenting text, graphics, video, and audio data. Third, the palmtop units have the ability to generate 3D models and 3D scenes. Finally, our systems are designed with portability in mind (the display is coupled with the processing unit). These design goals will enable users to bring their 3D information spaces with them or to access and interact with 3D spaces situated in the physical environment [Fitzmaurice 1993].

The following section briefly reviews and characterizes some of the virtual-reality research. Next, we define our interaction model and describe the Chameleon prototype. A performance experiment with the Chameleon is then presented along with lessons learned. Variations to the current prototype design are then explored in some detail.

2. BACKGROUND AND RELATED RESEARCH

Human beings are constantly reacting to a 3D world and have acquired a remarkable capability to comprehend 3D spatial relationships. Researchers and engineers are learning to take advantage of the natural 3D human information-processing capabilities for man-machine system interfaces.

Although a generally accepted definition of virtual reality has not been achieved, some key components in constructing VR systems have been identified. Zeltzer [1992] states that *autonomy*, *interaction*, and *presence* are three independent components in VR systems. Wickens [1992] identifies the use of

3D displays, dynamic animation, closed-loop interaction, ego-referenced displays, and enhanced sensory experience as five characteristics of VR technology. Interestingly, the interaction and presence components are usually considered as two independent elements, but this is not the case in the palmtop VR system described below (where input control and output display are integrated as a single unit).

Two seemingly disparate approaches have been pursued in the VR community: *outside-in* systems—simulating 3D environments within the 2D computer screen (e.g., rooms, information visualizer, and cone trees [Card et al. 1991]—and *inside-out* systems—virtual-reality environments in which users are completely immersed into an artificial 3D space by wearing a head-mounted display and, potentially, a DataGlove.™

The outside-in approach assumes that the user is external to the virtual world, looking inward via a viewing window. This approach, sometimes referred to as "fish tank environments" [Arthur et al. 1993; Ware and Slipp 1991], is practical and unintrusive to the user. Some systems include head-tracking devices and/or a stereo display. Designers of these systems must pay particular attention to providing the proper amount of depth cues and intuitive user interface designs to prevent the user from becoming disoriented. This approach offers a less engaging experience and sense of presence than the alternative approach. However, many researchers have successfully pursued the use of this outside-in style of 3D displays to help human operators to better perceive the natural task environment such as those in teleoperation (e.g., [Zhai and Milgram 1991]). These kinds of displays can also facilitate a user's comprehension of abstract concepts such as those in abstract hierarchical structures (e.g., [Chignell et al. 1992; Robertson et al. 1991; Xiao and Milgram 1992]).

The inside-out approach requires that the user be the focal point of the virtual world (ego-centric). The user is then in the midst of the environment, and surrounded by it. This approach requires a great deal of extra hardware but may be rewarded by a more engaging 3D experience. Designers of these systems must constantly grapple with the tradeoff between system response time and the quality of the rendered virtual environment. Providing high levels of detail of local objects along with coarse approximations of objects situated within a panoramic view will quickly tax many current hardware configurations. Nevertheless, this approach allows users to exploit their spatial and physical skills to navigate through the space, and it should become more viable as processing power and computer graphics technology improve.

A great deal of attention has been drawn to this type of immersive virtual reality which implies head-mounted displays [Bryson and Levit 1992; Sutherland 1968]. Furness [1986] introduced the idea of a "super cockpit" by providing the pilot a combination of physical controls with a synthetic scene to improve a pilot's spatial awareness. Aiming at providing a transparent man-machine interface to control an anthropomorphically designed multiple-degrees-of-freedom telerobot, Fisher et al. [1986] described a "telepresence" system which employs a head-mounted display and Dataglove.™ The ultimate goal of this system is to make the human operator feel "present" at the remote site. Indeed, one of the most advanced telepresence systems [Tachi et al. 1990] demonstrated significant advantages over traditional teleoperation in a number of tasks.

Large screens have also been used to enhance the 3D effect. With a large stereoscopic screen and a force-feedback robot arm as an input device, Brooks et al. [1990] demonstrated a new tool for scientists to better comprehend chemical reaction processes.

In this paper, we attempt to achieve 3D comprehension with a very small screen. Other than frequently used depth cues in 3D display, such as stereopsis, interposition, and perspective, the act of movement also provides a great deal of 3D sensation, as suggested in the motor theory of space perception [Gibson 1950; Wallach and O'Connell 1953]. As Gibson argued, "seeing" and "acting" are not separable. He also states that 3D comprehension in humans develops as a result of the intimate interaction of vision and muscular movement during ceaseless exploration of one's environment. The palmtop virtual-reality concept we are proposing utilizes kinesthetic movement to overcome the disadvantages of a small screen size while integrating visual and motor exploration of 3D displays.

3. INTERACTION MODEL

The interaction model for our prototype is an "ego-centric" spatially based "virtual cubic spreadsheet" in which a palmtop screen shows one cell of the spreadsheet at any given time (Figure 1). Objects within each cell can be rendered in 2D or in 3D perspective. Moving the palmtop left or right (x axis) will change the viewing of the spreadsheet on a column basis while raising or lowering the device (y axis) will allow for row movement within the spreadsheet. Moving the palmtop closer or further away from the user (z axis) allows a sequence of spreadsheets to be viewed. This z axis may best be used to alter the temporal view of data; older levels are visible as the device is pushed away, and future levels are made visible as the device is pulled closer to the user.

A simple yet powerful modification of this model is to introduce the notion of tilting the palmtop for localized viewing of adjacent cells or for issuing commands. We can pretend that the device is resting on a pin located in the center of the device. Tilting the device will cause the display to jump to one of the 8 adjacent cells depending on which edge(s) are moved (Figure 2). For example, the adjacent left cell would be displayed if the right edge of the device is tilted upward.

Alternatively, this tilting action could be a natural way of issuing commands. This is similar to the concept of pie menus [Hopkins 1991; Kurtenbach et al. 1993] except that menu choices are selected by a physical tilting action in 3D instead of a mouse or pen-directional marking in 2D. For example, to browse textual data, a palmtop application could be designed which uses the tilting mechanism as a means of issuing commands. Tilting the top part of the palmtop toward the user would indicate a "Scroll-Up by

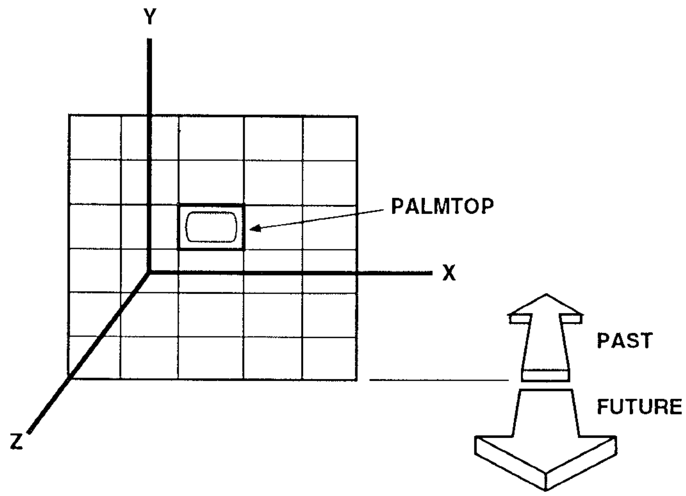

Fig. 1. Virtual cubic spreadsheet. The hand unit displays one cell of the spreadsheet at any given time. As a user moves the unit to different spatial locations, the appropriate cell of the spreadsheet is shown in the display of the hand unit. In this instance the z axis has been associated with time; past events are further away, and future events are closer to the person holding the palmtop.

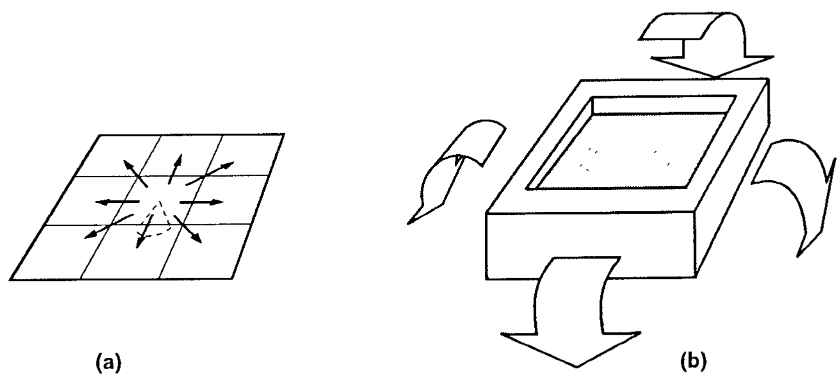

Fig. 2. Localized viewing of adjacent cells (a) can be achieved by making use of a tilting mechanism. The user pretends that the hand unit is resting on a pin. Tilting the unit (b) will allow for the appropriate adjacent cell to be viewed.

line" command while tilting the bottom part upward indicates a "Scroll-Down by line" command. Similarly, a "Next Page" and "Previous Page" command could be associated with a right and left tilting action, respectively (Figure 3a). A second example provides a means of browsing video data on a palmtop using the tilting mechanism. Tilting the hand unit toward the right would indicate a "forwardVideo" command while tilting toward the left would issue

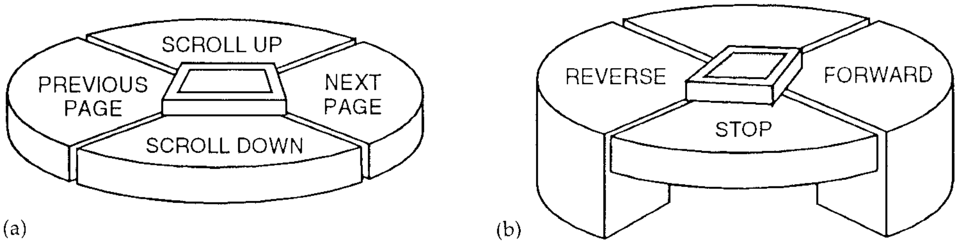

(a) (b)

Fig. 3. Pie menus. (a) text-browsing commands; (b) video-browsing commands.

a "reverseVideo" command. Finally, the degree to which the tilt occurs could indicate parametric levels of the command: slight tilts would be small units (1 frame per second) while greater tilts would indicate larger units (100 frames per second) (Figure 3b).

This "ego-centric" spatially based "virtual cubic spreadsheet" interaction technique relies on having the user maintain a simple mental model (i.e., the spreadsheet or cube) and exploits both spatial memory and muscle memory.

4. THE CHAMELEON PROTOTYPE

The goal of our design is to provide a framework for exploring how palmtop computers with high-performance graphics could be used with gestural inputs to provide portable VR environments for visualizing, exploring, and interacting with 3D information spaces. The prototype framework, known as Chameleon, supports this type of investigation which is described below.

Many of today's palmtop computers do not have the capabilities of high-performance graphics workstations such as the Silicon Graphics workstations. However, we have been able to simulate a palmtop of the future which would have these features as well as additional input capabilities.

The Chameleon prototype configuration (see Figure 4) uses a small 4-inch color LCD monitor, which fits comfortably in one's hand, as the basis for the palmtop of the future. A Silicon Graphics 4D/310GTX Iris workstation is used as the main processing unit and graphics engine. Currently, the images on the workstation are fed into the small hand unit by using a video camera. Technologically more robust solutions will improve the video quality (e.g., using an NTSC output channel on the SGI workstation). This configuration gives the illusion that the 4-inch monitor (i.e., palmtop) has the capabilities of an SGI workstation, albeit with an inferior image size and resolution.

To facilitate input controls, the palmtop has been given the ability to monitor its position and orientation information within a 3D workspace. The Ascension Bird™ (by Ascension Technology) is a 6-degree-of-freedom input device and is attached to the small monitor. This allows the system to detect user gestures in terms of positional data (x, y, and z dimensions) as well as orientation data of the palmtop (pitch, yaw, and roll). In addition, a mechanical push button has been attached to the top of the device to obtain input selections from the user (see Figure 5). Both the bird and the response button

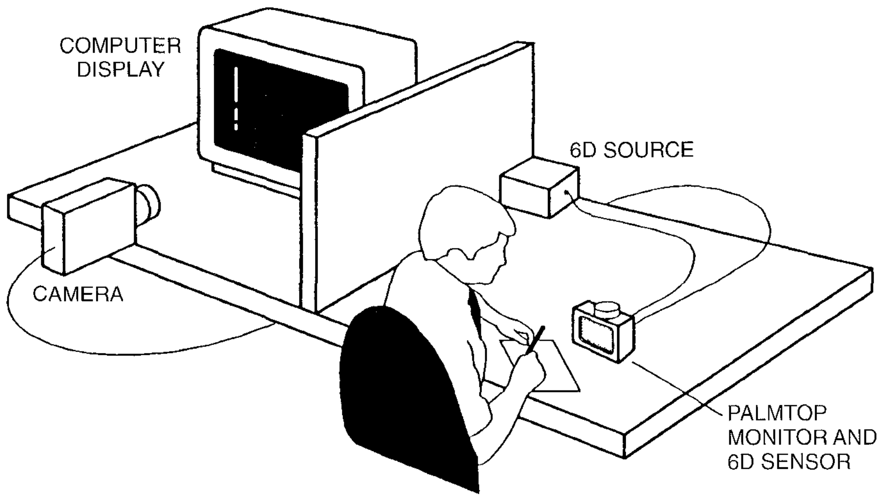

Fig. 4. Configuration of the Chameleon prototype.

Fig. 5. Palmtop unit consisting of a video display, response bottom (top) and a 6D input sensor for providing position and orientation information (x, y, z, pitch, yaw, and roll).

allow the user to make simple gestures to navigate within a 3D workspace and to issue commands.

It is difficult to precisely measure the dynamic response of the Bird™ device as well as other 6D devices. Adelstein et al. [1992] have developed a controlled testbed and methodology for measuring these devices. Informal data measurements for our optimized bird device driver have an approximate response delay of 50 ms providing translational accuracy on the order of 0.1 inch and angular accuracy on the order of 0.5 degrees within a range of a 3-foot cube. All of these measurements are sufficiently fast and reliable enough to support novel interaction techniques.

To navigate through a 3D workspace, the translation (x and y axes) and zoom (z axis) data of the palmtop are constantly fed back to the SGI workstation which dynamically updates the 3D scene being relayed back to the palmtop unit. Objects appearing in the 3D scene are drawn in perspective. The net effect is that the palmtop unit acts as a porthole window into a 3D workspace. For example, as a user translates the palmtop unit to the left, he or she moves toward the left wall of the 3D workspace.

In addition, we have opted to make the cell boundaries in the virtual cubic spreadsheet transparent instead of opaque (Figure 6a). Users, therefore, not only see the contents of the cell they are occupying in the spreadsheet but also all the deeper cells (i.e., those in the z axis; see Figure 6b). Finally, the prototype allows for perspective viewing instead of an orthogonal view from the current cell (Figure 6c). This means that the user views cells adjacent to the current cell in a slight fisheye manner [Furnass 1986; Robertson et al. 1991]. Both the transparent cells and perspective viewing provide users with additional orientation information. However, these characteristics can be changed depending on the application or information space being browsed. For example, if a user is browsing a spreadsheet of numbers, it may be too confusing to have transparent cells since the numbers would overlap along the z axis.

The granularity size of the cells in the virtual cubic spreadsheet may also be set. By default, the cells are at a very fine granularity. Coarser cell granularities offer a gridding effect which is best suited for discrete instead of continuous translations.

To illustrate the navigation effect, Figure 7a shows a user moving deeper into a 3D workspace. Figures 7b and 7c show a before-and-after view of the palmtop unit and the contents of the palmtop screen as the user zooms into the 3D space along the z-axis. The 3D space consists of a room containing cubes randomly positioned in the space. The cubes could represent a spatial database or a user's filing system in which each cube symbolizes a file. Note that as the palmtop is moved from the start to the end position, the user sees a smooth zooming animation during the traversal of the path.

The system can also switch into "clutch" mode in which the translation and zoom controls are registered only while the response button is depressed (i.e., the clutch is engaged). When the button is released, the position is frozen, and the user is able to move the palmtop to a closer, more comfortable position or to place the device on the desktop. If the clutch is engaged again,

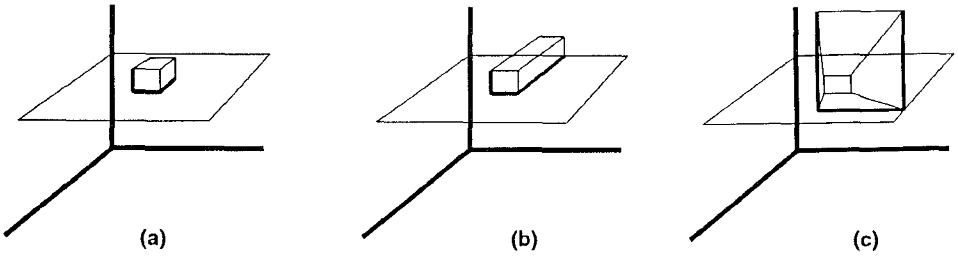

Fig. 6. Cell options for the virtual cubic spreadsheet: (a) opaque, (b) transparent orthogonal, and (c) transparent perspective.

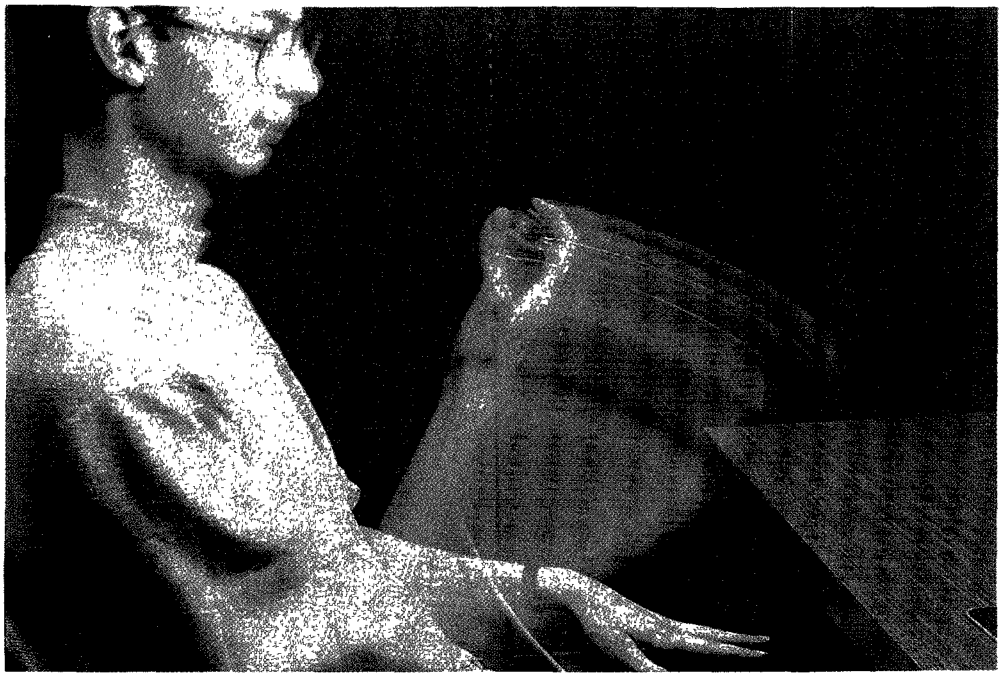

Fig. 7a. Palmtop unit in clutch mode. A user holds a button to register movement within the 3D workspace.

the model adjusts the 3D workspace to the current position of the palmtop to resume movements from the last registered position.

A cross-hair cursor is fixed in the center of the screen to allow the user to select objects within the virtual world. Users line up the target object over the cross-hair and click on (or in some situations double-click) the response button. This causes an imaginary ray centered at the cross-hair to emanate from the palmtop unit toward objects in the virtual world. The first object

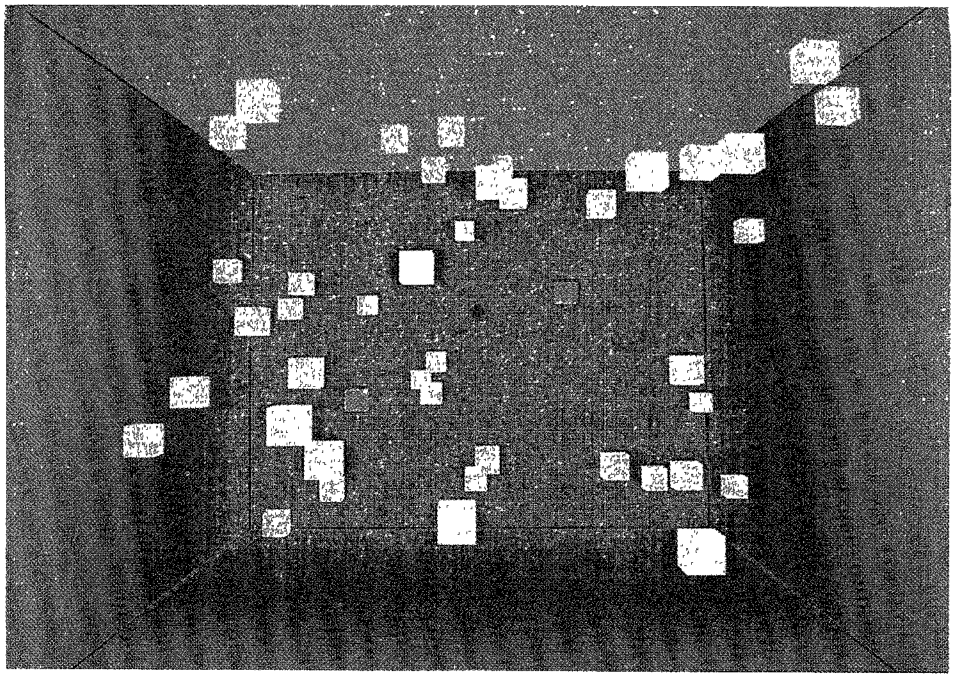

Fig 7b. What the user sees on the palmtop before gesturing deeper into the 3D workspace.

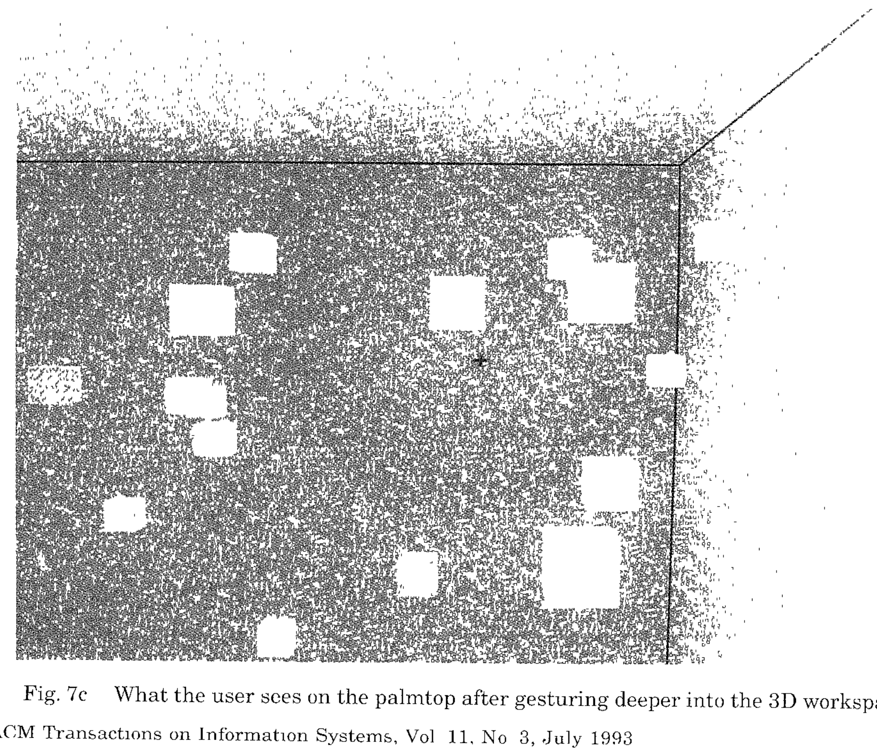

Fig. 7c What the user sees on the palmtop after gesturing deeper into the 3D workspace.

encountered is selected. Visual and auditory feedback are important and can improve a user's accuracy in selecting objects. For example, as the cross-hair cursor is in range over a target, the target can change color. In addition, a nonspeech audio sound (e.g., a low click tone) can be made to indicate that the user is over a target. This type of reinforcement feedback is important since the display resolution can be low and since targets can be small.

Simple gestures and application-specific controls can be built into the system. For example, if the workspace contains a 3D tree hierarchy of a file system, then users would be able to select a node and gesture downward in a tugging motion to cause the tree to rotate along the x axis. Conversely, selecting a node and tugging upward would cause the tree to rotate in the opposite direction.

The current prototype does not make use of the orientation information of the 6D input device; users always look straight ahead into the 3D workspace. Future versions of the prototype may allow for true one-to-one mappings of the palmtop orientation to the 3D model. Note that this interface may not be optimal since the user would have to physically move into a new viewing perspective to get a corresponding view into the 3D workspace. For example, rotating the device by 90 degrees along the y-axis to get a side view means that the user would also have to move his or her head to the side in order to see the display on the palmtop. While the scenes generated for the palmtop display are drawn in 3D perspective, we are looking at providing additional interactive viewpoint control [McKenna 1992]. Moreover, tracking the user's head position along with the palmtop position may provide greater depth sensation.

The Chameleon prototype configuration allows us to perform rapid prototyping of new applications since our development efforts are not tied to a highly customized programming environment often found on palmtop computers. That is, depending on the requirements of the applications we are exploring, it is possible to use any computer as our base machine (e.g., Silicon Graphics or Sun workstation or a Macintosh computer) and any development environment best suited for the task.

Although our prototype palmtop device is currently tethered by cords (due to the video feed and 6D input device), it provides a rich environment for testing new situations, applications, and user interactions in a technology configuration which we anticipate will be available in a few years in a highly portable form.

5. EXPERIMENT

The performance of the Chameleon system was tested in an experiment. The purpose was to see how well depth could be perceived in the small screen compared to the large screen and the extent to which movement of the small display (in a palmtop virtual-reality condition) could improve depth perception.

The subjects in the experiment were 12 volunteers from the University of Toronto community. Each subject participated in 30 trials in each of 3 conditions. For each trial the subject was shown a field of cubes within a 3D space and was asked to select which of two highlighted cubes was further away on the z-axis (i.e., into the screen). One condition consisted of a 21-inch monitor with a static display. The second condition consisted of the 4-inch television monitor, also with a static display. In the third condition, the 4-inch television monitor was again used, but now the subject was able to move through the display using the palmtop virtual-reality gestures for zoom and translation that were described earlier in this paper. Figures 8a through 8c show the three conditions. In the first two conditions, the mouse was used to position the cursor within the display and make the selection. In the third condition, the selection was made by moving the monitor until a cursor was positioned over the box, and then a response button attached to the monitor (as shown in Figure 5) was clicked. A limited amount of training was provided in each of the three conditions, with subjects practicing the task until they said they were comfortable with it. In no case did the initial training period last more than a few minutes.

The number of errors were tabulated for each subject in each condition, making a total of 36 (3×12) data points. These data were then subjected to a one-way analysis of variance. There was a significant difference between the three conditions ($F[2,33] = 34.9$, $p < 0.001$), with the large display (2.08 errors per subject) and the palmtop condition (3 errors per subject) leading to fewer errors than the small (stationary and static) display (6.08 errors per subject). Post hoc analysis showed no significant difference in the number of errors produced in the large-display condition versus the palmtop virtual-reality condition.

Response times were also collected across the three conditions. The median response times for each condition were calculated for each subject and were then subjected to analysis of variance. Response times varied significantly across the three conditions ($F[2,33] = 112.9$, $p < 0.0001$), with the large- and small-display condition having a significantly shorter mean response time (2.75 and 3.75 seconds, respectfully) than the palmtop virtual-reality condition (10.67 seconds). This result is not too surprising since the third condition allowed the subjects to explore the environment. If two cubes had similar depth placement, then the subject could investigate or visit each cube. In the other two static scene conditions, the user must make a decision, perhaps an arbitrary one, without investigation.

The results of this experiment show that, with very little training, perception of depth in the palmtop virtual-reality condition is about as good as corresponding depth perception in a large (but static) display. However, performance was much poorer (an average of 3–4 more errors per 30 trials) when the static small display was used.

The ability to move and explore the display seems to be critical in developing adequate depth perception on a small screen. While depth perception was relatively good in the palmtop virtual-reality condition, the response time data showed that selecting the box that was further away took much longer in the palmtop condition. After the experiment, the subjects indicated that this longer time was partially due to the difficulty in physically making the

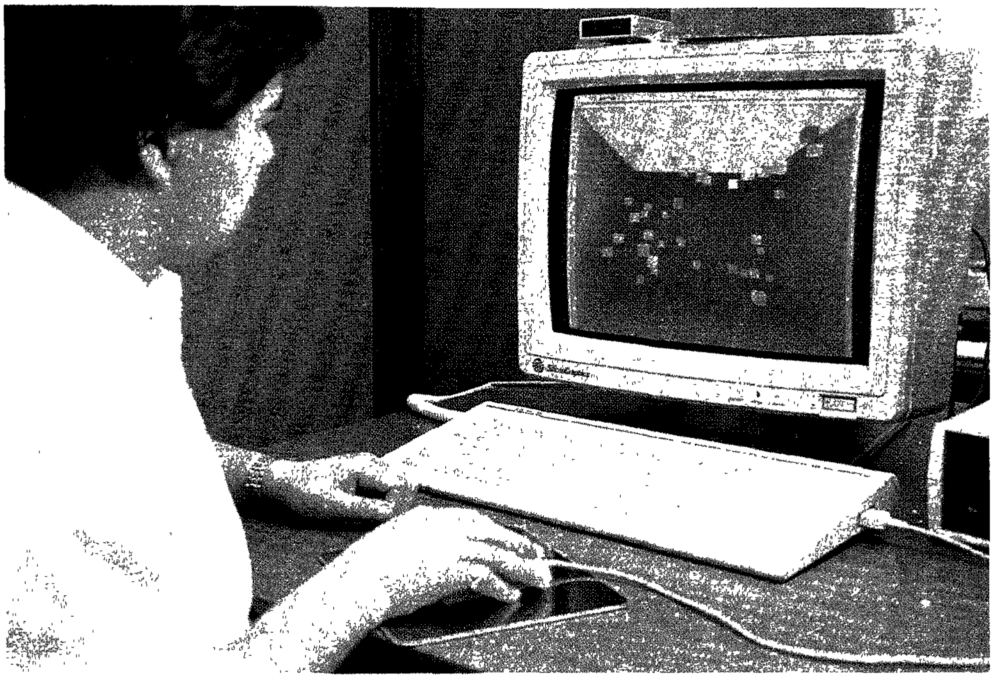

Fig. 8a. First experiment condition consisting of a large 21-inch monitor, static display, and mouse input.

selection, rather than in making the perceptual decision about which of the two boxes was further away on the z-axis. We anticipate that with further experience subjects would be able to make selections with the palmtop virtual-reality system much more quickly.

While running the experiments, a few subjects initially stated that the palmtop virtual-reality controls were not intuitive and that they perceived them as being completely reversed. That is, moving the device to the left should move the entire 3D world or a target object to the left with the user always in a fixed location. These comments revealed that they had an object view instead of an ego-centric view for the controls. After further usage, they all were able to switch to the ego-centered model. We found that users having difficulty with understanding the ego-centered model could make the transition more quickly if they physically rocked in their seat while keeping the palmtop at a constant distance from their bodies. For example, to zoom in and out of the 3D workspace the subject would rock forward and then backward. The physical act of moving the body corresponded directly with the controls for the palmtop and complemented the visual display to reinforce the notion of moving around in a virtual world.

6. DESIGN VARIATIONS

After conducting the experiment, we have begun to explore deeper into the design space for the palmtop unit in terms of modifying the movement

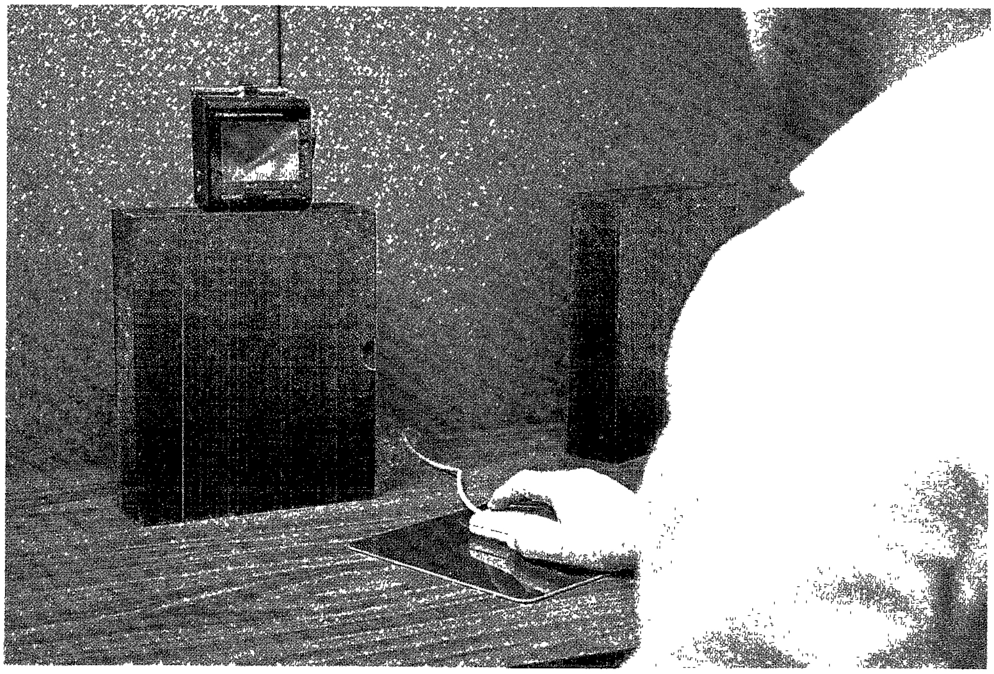

Fig 8b. Second experiment condition consisting of a small 4-inch screen, static display, and mouse input

controls, varying the size and shape of the 3D workspace, considering the interaction of the palmtop with other computational devices, and exploring applications. Each of these are described in some detail below.

Movement Controls

We have designed but not yet implemented in our prototype two mechanisms: glancing and mirror mode for alleviating some of the physical orientation problems. To take a temporary glance either to the left or right of the current palmtop position, the user would rotate the device between 15 and 45 degrees. The angles received by the 6D input device would be exaggerated such that a 45 degree rotation corresponds to a temporary full 90 degree smooth viewing change in the desired direction. This allows for the display of the palmtop to be seen by the user without much, if any, adjustment by the user's head. To compensate for small unintentional rotations, no change in viewing perspective will be registered between +15 and −15 degrees.

To view the back side of objects within the 3D workspace, the user could switch into mirror mode in which the palmtop acts as a mirror. Instead of users seeing deeper into the z-axis, they are now viewing objects between their eyes and the palmtop unit. That is, instead of the "eye" being placed on the back of their hand, it is temporarily placed on the front of their hand. To

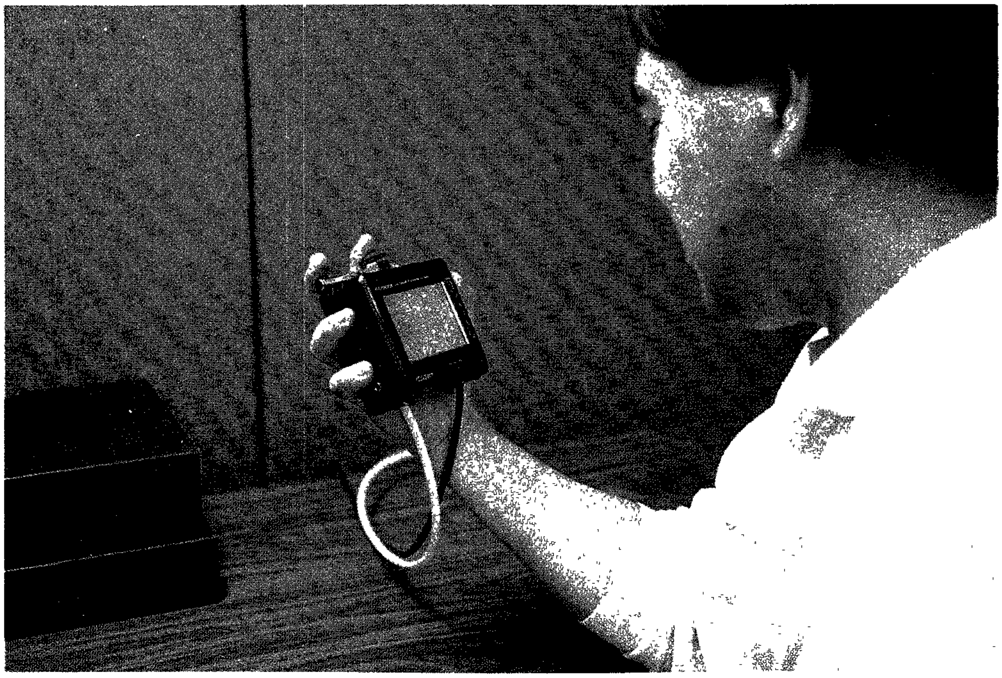

Fig. 8c. Third experiment condition consisting of a small 4-inch monitor, dynamic display using the palmtop virtual-reality system for input.

see the back side of an object, the user would switch into mirror mode once the object is located between the user and the palmtop unit.

We have also explored ways of providing physical support of the device, in terms of both position and orientation, to reduce the anticipated muscle fatigue after prolonged use of the palmtop. This can be thought of along the same lines as the "boom" displays [Fisher et al. 1986] which suspend a goggle-like display in the air, counterbalanced on a mechanical boom apparatus. A simple way to provide physical support is to use a mechanical arm similar to that of a desk lamp with a holster attached to the end of the arm (Figure 9). This would allow users who work at a desk to use the palmtop in a more comfortable manner.

Considerable research has been undertaken to systematically map out the design space for 6D input devices. For example, Ware and Osborne [1990] have begun to evaluate metaphors for exploration and virtual-camera control in virtual environments using 6D input devices. Mackinlay et al. [1990] have performed a semantic analysis of input devices and have applied the theories toward building controls for 3D egocentric motion. Finally, Zhai et al. [1993] has defined a taxonomy of 6D input devices which we are using to define and investigate alternative movement controls for the palmtop. Three variations are described below which are currently being investigated.

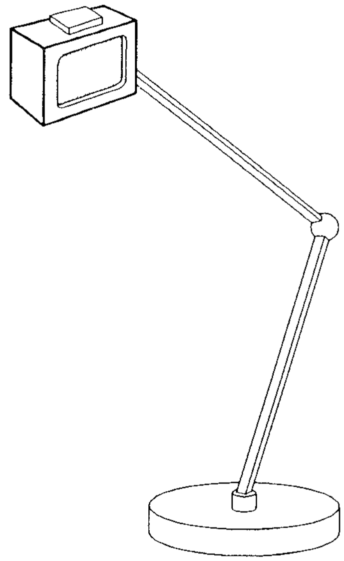

Fig 9 A holster for the palmtop to reduce muscle fatigue during prolonged use of the device

Visual Joystick. The visual-joystick design reduces the need to make global or extended hand gestures by adopting a rate control mechanism. Users make small hand motions with the palmtop unit in the direction they wish to travel. The traveling velocity is determined by the distance of the gesture. For example, a one-inch gesture to the right would cause the user to slowly travel to the right-hand side of the virtual space. A five-inch gesture would indicate a much faster velocity towards the right. To slow down, the user moves the palmtop back toward the original starting position. To stop moving, the user brings the palmtop back to the actual original starting position. Note that it will be convenient to define a small dead region at this position to prevent unintentional movements. In addition, the clutch can be used to engage the movement controls. The visual-joystick strategy allows for a rich level of movement controls in a relatively small gesturing space.

Tilt and Go. Tilt and go minimizes spatial gesturing by using the orientation of the palmtop as an indication of direction. That is, tilting the palmtop to the right will cause the user to travel to the right-hand side of the virtual space. Movement velocity can be mapped into the depth of the tilt. Slight tilts cause slow movement while deeper tilts cause faster movements. Alternatively, small push buttons can be attached to the palmtop unit to serve as accelerator and brake pedals.

Independent Directional Buttons. The third design variation separates each of the dimensional movements into individual physical buttons. Three buttons can be used to control the x, y, and z movements. We are currently considering using three-way, spring-loaded buttons in which the rest state occurs while no force is being applied to the button. The button can detect if the user is pushing on the left or right edge of the button; when released, the button returns to the rest or off state. Using separated button controls may be ideal for tasks which require the user to line up objects along an axis or for precise movement control along one dimension in general. The design, however, may suffer from requiring extra cognitive overhead from the user.

Alternatively, we may build software controls similar to the snap-dragging concepts [Bier 1990] for precision control and placement of objects in three dimensions. This design, as well as the others, will be investigated further to determine their feasibility and performance characteristics.

Levels of Fidelity

Another issue concerns the level of fidelity of the system. There are currently three envisioned interaction environments which have been formulated. Each level offers more input/output fidelity for the user. Note that a variety of environments could be derived from those defined here. Our prototype is considered the first environment in which the palmtop unit operates independently on a user's mental model of a virtual cubic spreadsheet.

Physical Static Frames of Reference. Instead of relying on the user to maintain a mental model of the virtual cube, the second environment offers a physical grid as a static frame of reference. Labels can be applied to the grid to reinforce regions of the space or to indicate coarse-level contextual information. The grid could be designed to rest the palmtop in a given cell without requiring the user's support. Foldup grids could be designed for increased portability.

Physical Dynamic Frames of Reference. The static frame of reference can be significantly improved if we replace it with display screens. A large two-page monitor could act as a dynamic backdrop for displaying information. The monitor could be used in a variety of ways: (1) to dynamically partition, highlight, or indicate regions of the space, (2) as a means of outlining the silhouette of a virtual object being viewed in the space, or (3) to act as a mirror and thus display a rendered image of the back side of a virtual object. A second monitor embedded in a desk and positioned flush against the first monitor could further promote the virtual cube into a more realistic 3D imaging chamber (Figure 10).

The bottom or "floor" monitor could display contextual information or simply supply a means of displaying shadows of virtual objects within the chamber. Finally, a larger virtual cube could be realized if a large, rear-projection screen is used as a backdrop monitor in conjunction with the palmtop hand unit. This is similar to the CAVE project [Cruz-Neira et al. 1993] which has a rear-projection design for room size virtual cubes.

Workspaces

Wrapping the 3D Workspace around the User. Instead of restricting the 3D workspace to a relatively small cubic volume, we can propose alternative 3D workspaces. For example, in the office setting, a donut shape surrounding the user's chair may be more appropriate (Figure 11). The user could pivot in

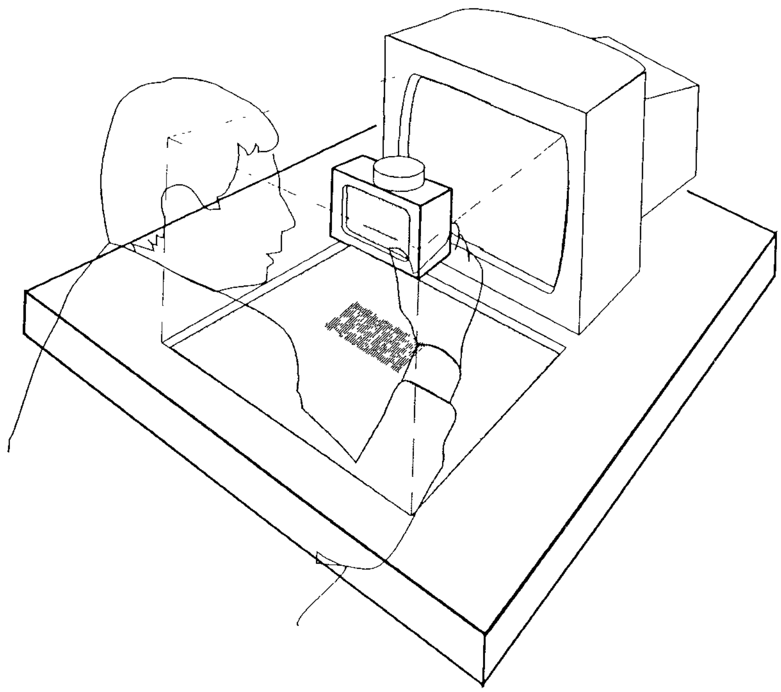

Fig. 10. A crude "imaging chamber" consisting of a backdrop and floor monitor to be used in conjunction with the hand unit.

his or her chair to access information with the advantage that the donut offers more surface area and spatial displacement.

The Filmstrip Model. This final variation suggests that the palmtop unit could be used in conjunction with a conventional computer monitor in which information is placed on an imaginary roll of film. Two virtual reels are defined to the left and right of the user's workspace. The conventional computer monitor acts as a view screen which enlarges and shows more detailed information compared to the monitor on the palmtop unit (Figure 12). The palmtop acts as a previewer mechanism and a means of gesturing for the film to be rolled forward or backward. Note that the film does not necessarily have video data on it but could contain document icons, spreadsheets, graphical figures, and chunks of text.

Applications

A variety of applications can be developed that make use of the core palmtop virtual-reality ideas.

Personal Calendars. Consider having the cubic spreadsheet represent a user's personal calendar. The days of the week (i.e., Monday, Tuesday, etc.) could be defined as the columns, and each week would be contained in a row

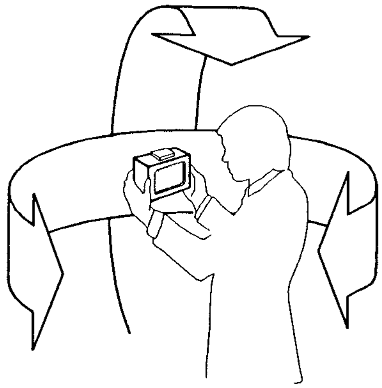

Fig. 11. Extending the virtual-reality workspace around the user in a donut shape.

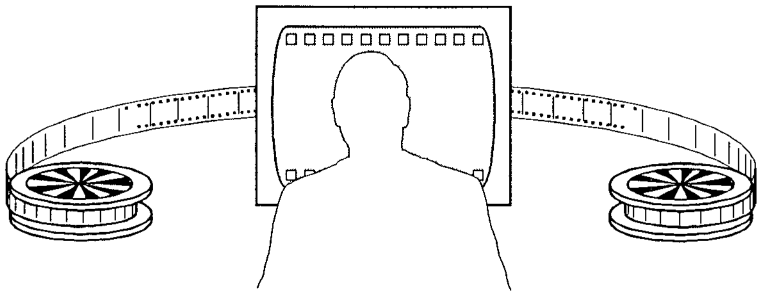

Fig 12. Film reel metaphor in which data placed on the surface of the film can be viewed in detail by the conventional computer monitor. The palmtop unit serves as a means for advancing or rewinding the film and to preview portions of the film not currently shown on the main monitor.

(or perhaps hours in a given day—depending if the user was observing a weekly or monthly view). Past months or weeks would be visible as the user pushes the device away while future dates would be visible as the user pulls the device closer to himself. If the technology could be sufficiently miniaturized, this application could be programmed into a wrist watch for improved accessibility.

There are also interesting improvements or enhancements that may be reached by using nonspeech audio for augmented output. For example, consider the calendar application in which the user wishes to determine which day has the fewest appointments. By assigning sounds to appointments and having users quickly wave the palmtop device from left to right (i.e., Monday to Friday) they can get a sense of what appointments have been made that week and which days look or sound open.

Information Landscapes. Another type of application could provide users with the ability to store their files in a 3D portable information landscape.

Files could be represented as icons, 3D icons, a sample of the data (e.g., a video frame for a movie), or as common physical objects. These representations would be placed in a 3D scene and accessed by the user gesturing to navigate to the target objects. The design attempts to exploit a user's spatial and muscle memory within a spatial database representation [Herot 1980].

7. CONCLUSIONS

The palmtop virtual-reality concept we are proposing utilizes kinesthetic movement to overcome the disadvantages of a small screen size. We believe that the Chameleon prototype configuration offers a fertile environment for exploring integrated control and output in a hand-held environment. The prototype provides a proof of concept for palmtop virtual reality, and our experimental results demonstrate that depth perception can be about as good in a palmtop environment as it is in a large screen display, providing that exploration of the display is permitted.

A number of design variations were described and will be explored as future research. More specifically, we wish to determine the feasibility and performance characteristics of some alternative movement controls operating in complex 3D information spaces.

Applying some of the virtual-reality concepts to palmtop computers seems to be beneficial as user communities move to more mobile and nomadic computing. Within a palmtop environment, we want users to be able to visualize and interact with large portable information structures. This palmtop virtual-reality strategy will serve to increase the communication bandwidth between the user and computer for highly mobile and portable computers.

ACKNOWLEDGMENTS

The authors greatly appreciate the contributions of Ferdie Poblete, Robert Gilbertson, Sarah Zuberec, Ron Baecker, Bill Buxton, Paul Milgram, and the reviewers for their suggestions and support.

We gratefully acknowledge the support to our laboratory from the Natural Sciences and Engineering Research Council of Canada, the Information Technology Research Centre of Ontario, Apple Computer, and Xerox PARC.

REFERENCES

ADELSTEIN, B. D., JOHNSTON, E R., AND ELLIS, S R 1992. A testbed for characterizing dynamic response of virtual environment spatial sensors. In *Proceedings of the ACM 5th Annual Symposium on User Interface Software and Technology* ACM, New York, 15–22

ARTHUR, K W., BOOTH, K. S., AND WARE, C 1993. Evaluating 3D task performance for fish tank virtual worlds *ACM Trans. Inf Syst 11*, 3. This issue

ASCENSION TECHNOLOGY CORP 1992. *The Ascension Bird*. Burlington, Vt

BIER, E A 1990 Snap-dragging in three-dimensions, 1990 Symposium on Interactive 3D Graphics *Comput Graph. 24*, 2 (Mar), 193–204.

BISHOP, G., AND FUCHS, H. 1992. Research directions in virtual environments: Report of an NSF Invitational Workshop. *Comput. Graph. 26*, 3, 153–177.

BROOKS, F. P., JR., OUH-YONG, M., BATTER, J. J., AND KILPATRICK, P. J. 1990. Project GROPE—haptic display for scientific visualization. *Comput. Graph. 24*, 4 (Aug.), 177–185.

BRYSON, S., AND LEVIT, C. 1992. The virtual wind tunnel. *Comput. Graph. Appl. 12*, 4 (July).

CARD, S. K., ROBERTSON, G. G., AND MACKINLAY, J. D. 1991. The information visualizer, an information workspace. In *Proceedings of CHI'91*. ACM, New York, 181–188.

CHIGNELL, M., POBLETE, F., AND ZUBEREC, S. 1993. An exploration in the design space of three dimensional hierarchies. In *Proceedings of the Annual Meeting of the Human Factors and Ergonomics Society*. To be published.

CRUZ-NEIRA, C., SANDIN, D. J., AND DEFANTI, T. A. 1993. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proceedings of SIGGRAPH '93*. ACM, New York.

ELLIS, S. R., III. 1991. Nature and origins of virtual environments: A bibliographical essay. *Comput. Syst. Eng. 2*, 4, 321–347.

FISHER, S., MCGREEVY, M., HUMPHRIES, J., AND ROBINETT, W. 1986. Virtual environment display system. In *Proceedings of the 1986 Workshop on Interactive 3D Graphics*. ACM, New York, 77–87.

FITZMAURICE, G. W. 1993. Situated information spaces and spatially aware palmtop computers. *Commun. ACM 36*, 7 (July), 38–49.

FURNASS, G. G. 1986. Generalized fisheye views. In *Proceedings of CHI '86*. ACM, New York, 16–23.

FURNESS, T. A., III. 1986 The super cockpit and its human factors challenges. In *Proceedings of the Human Factors Society—30th Annual Meeting*. Human Factors Society, 48–52.

GIBSON, J. J. 1950. *The Perception of the Visual World*. Houghton Mifflin Company, Boston, Mass.

HEROT, C. 1980. Spatial management of data. *ACM Trans Database Syst. 5*, 4 (Dec.), 493–514.

HOPKINS, D. 1991. The design and implementation of pie menus. *Dr. Dobb's J. 16*, 12 (Dec.), 16–26.

KRUGER, M. 1991. *Artificial Reality*. Vol II. Addison-Wesley, Reading, Mass.

KURTENBACH, G., SELLEN, A., AND BUXTON, W. 1993. An empirical evaluation of some articulatory and cognitive aspects of "marking menus." *J. Human-Comput. Interaction 8*, 2 (Jan.).

MACKINLAY, J., CARD, S. K., AND ROBERTSON, G. G. 1990. A semantic analysis of the design space of input devices. *J. Human-Comput. Interaction 5*, 2–3, 145–190.

MCKENNA, M. 1992. Interactive viewpoint control and three-dimensional operations. In *Proceedings of the 1992 ACM Symposium on Interactive 3D Graphics*. (Special issue of *Comput. Graph.*, 53–56.)

ROBERTSON, G. G., MACKINLAY, J. D., AND CARD, S. K. 1991. Cone trees: Animated 3D visualizations of hierarchical information. In *Proceedings of CHI'91*. ACM, New York, 189–202.

SUTHERLAND, I. 1968. A head-mounted three dimensional display. In *Proceedings of the Fall Joint Computer Conference*. Thompson Books, Washington, D.C., 757–764.

TACHI, S., ARAI, H., AND MAEDA, T. 1990. Tele-existence master slave system for remote manipulation (II). In *Proceedings of the 29th Conference on Decision and Control*. IEEE, New York.

WALLACH, H., AND O'CONNELL, D. H. 1953. The kinetic depth effect. *J. Exp Psychol. 45*, 205–217.

WARE, C., AND OSBORNE, S. 1990. Exploration and virtual camera control in virtual three dimensional environments. In the *1990 Symposium on Interactive 3D Graphics. Comput. Graph. 24*, 2, 175–183.

WARE, C., AND SLIPP, L. 1991. Using velocity control to navigate 3-D graphical environments: A comparison of three interfaces. In *Proceedings of the Human Factors Society, 35th Annual Meeting*. Human Factors Society, 300–304.

WICKENS, C. D. 1992. Virtual reality and education. In *Proceedings of IEEE SMC Conference*. IEEE, New York, 842–847.

XIAO, Y., AND MILGRAM, P. 1992. Visualisation of large networks in 3-D space: Issues in implementation and experimental evaluation. In *Proceedings of the 1992 CAS Conference* Vol. 1. IBM Canada Ltd. Laboratory, Center for Advanced Studies, Toronto, 247–258

ZELTZER, D. 1992. Autonomy, interaction, and presence. *Presence—Teleop. Virtual Envir 1*, 1 (Winter), 127–132.

ZHAI, S., AND MILGRAM, P 1991. A telerobotic virtual control system In *Proceedings of Cooperative Intelligent Robotics in Space II* SPIE, vol. 1612. Boston, Mass. 311–320.

ZHAI, S , MILGRAM, P., AND DRASCIC, D. 1993. An evaluation of four 6 degree-of-freedom input techniques. In *InterCHI'93, ACM Conference on Human Factors in Computer Systems* (Amsterdam, The Netherlands, Apr). ACM, New York.

Received November 1992; revised April 1993; accepted April 1993