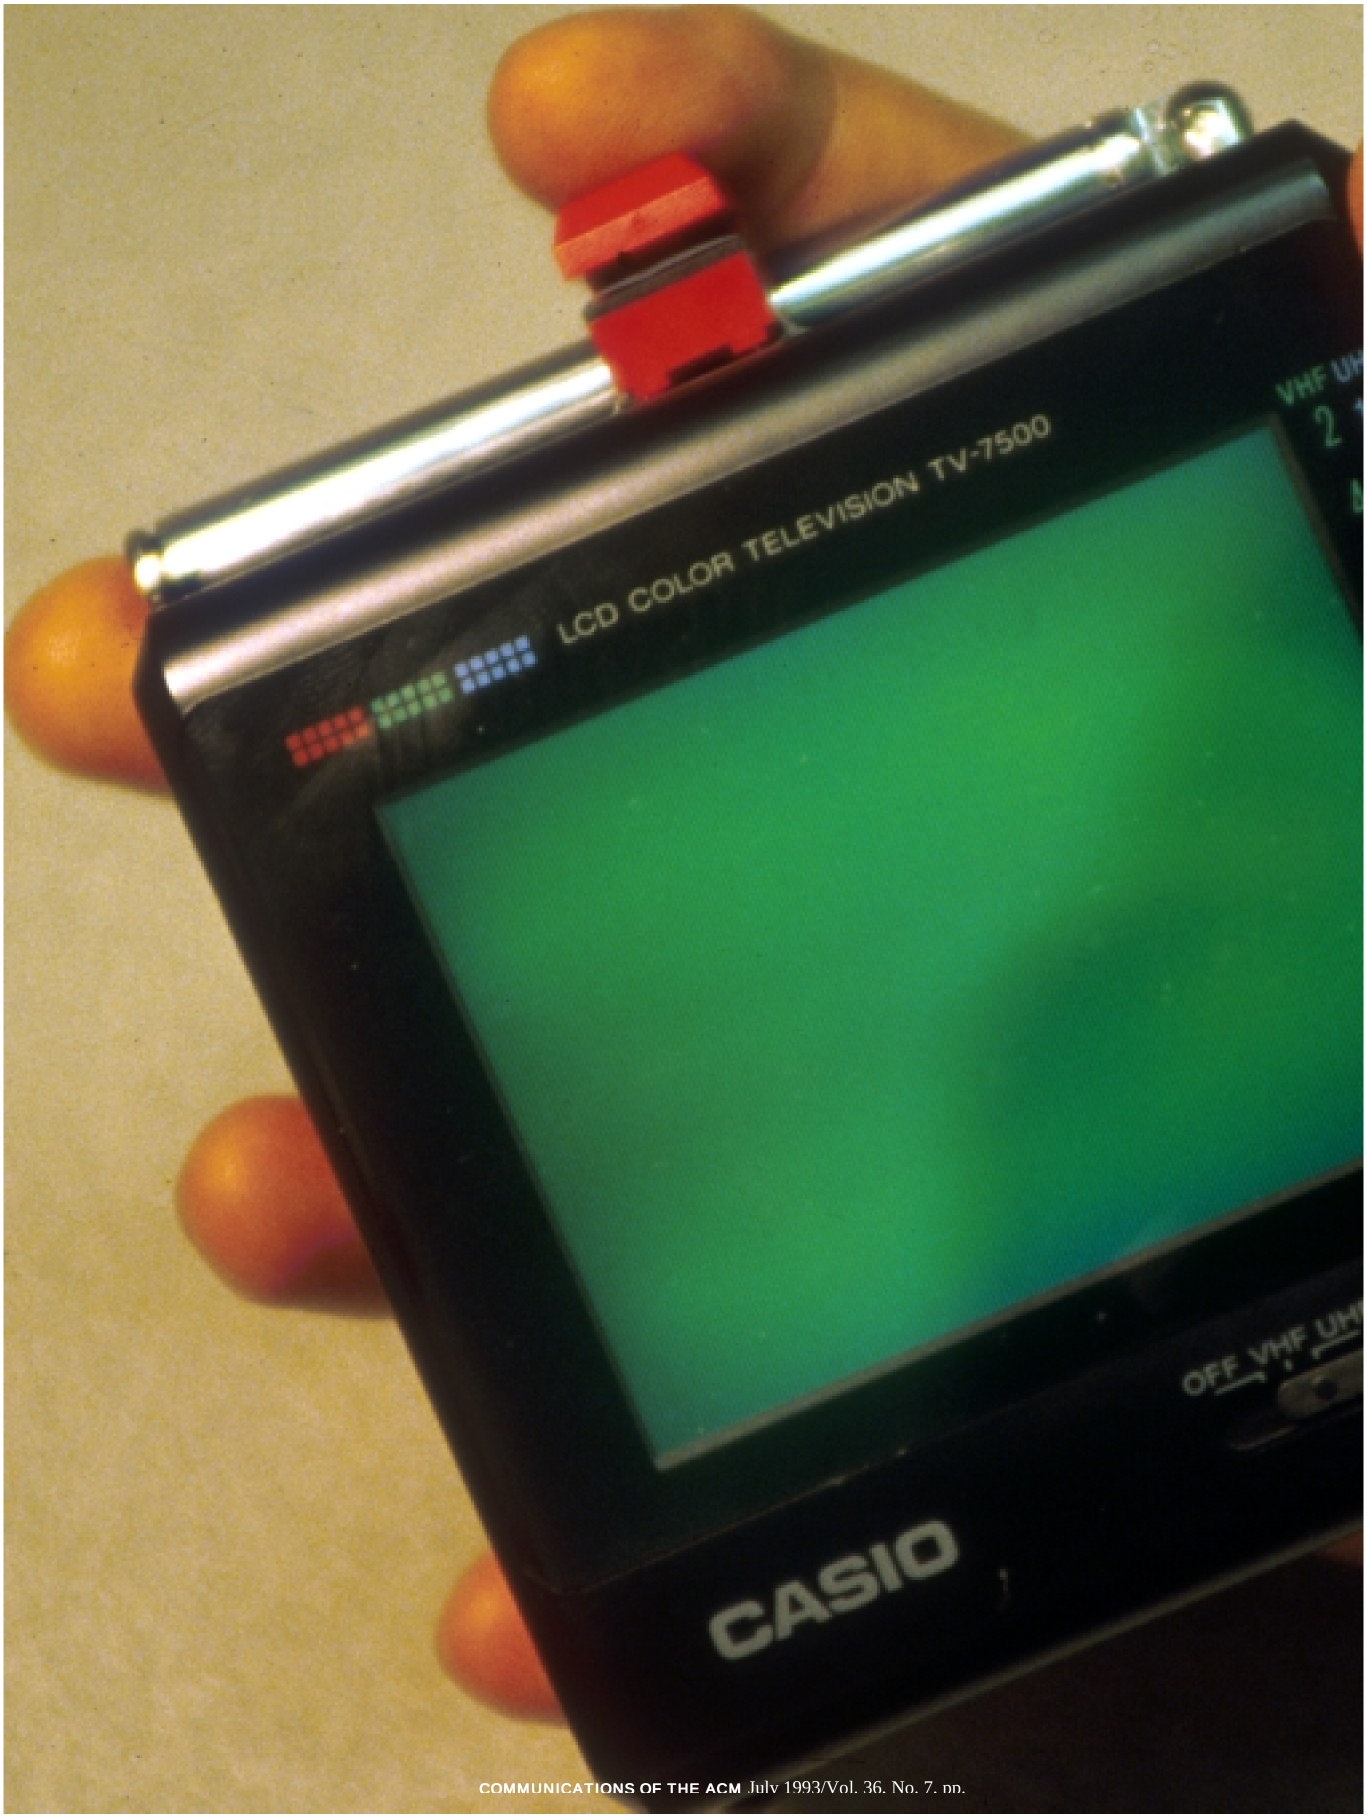
LCD COLOR TELEVISION TV-7500
VHF UH
OFF VHF UHF
CASIO

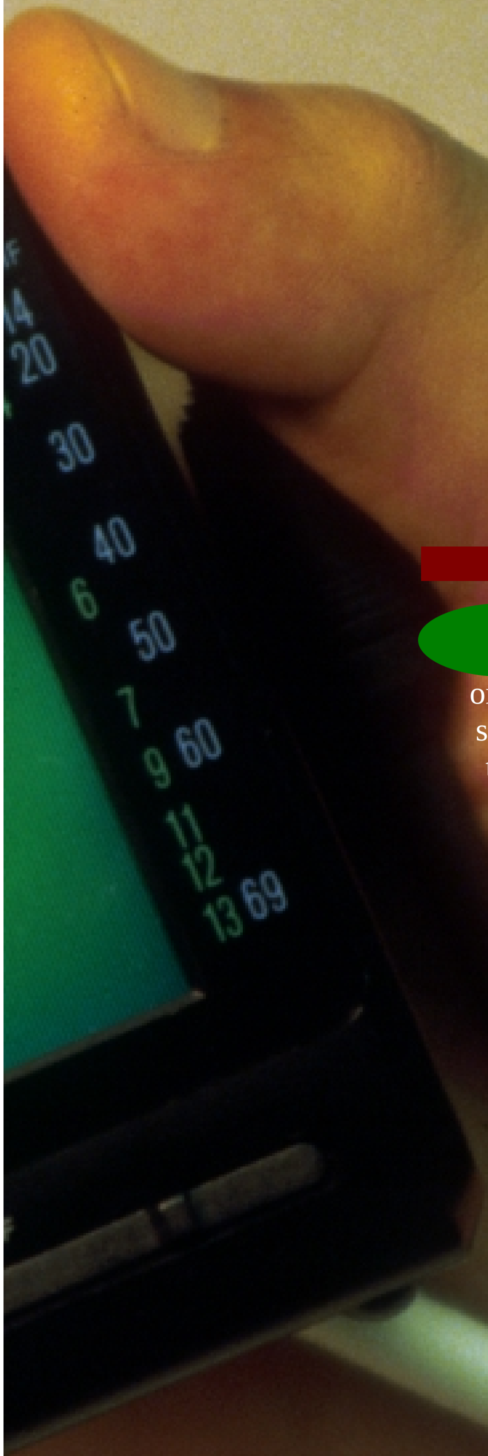

# SITUATED INFORMATION SPACES AND SPATIALLY AWARE PALMTOP COMPUTERS

George W. Fitzmaurice

No longer will we need to be tethered to a stationary computer workstation to browse electronic databases or synthetic 3D information spaces transformed onto a 2D display surface. Instead, we will browse, interact, and manipulate electronic information within the context and situation in which the information originated and where it holds strong meaning. A small, portable, high-fidelity display and spatially aware palmtop computer can act as a window onto the 3D-situated information space – providing a bridge between the computer-synthesized data and physical objects. Our Chameleon prototype explores some of the combined input controller and output display paradigms needed to visualize and manipulate 3D-situated information spaces.

Electronic information spaces are encroaching on our everyday environment. We are increasingly carrying electronic information with us (e.g., floppy diskettes) and also tapping into reservoirs of information via access stations (e.g., automatic teller machines, telephones). Indeed, portable computing allows us not only to carry the information but also to access, modify, and interact with it in a matter of seconds.

Ubiquitous computing (see Weiser's article in this issue) will further these abilities and cause the generation of short-range and global electronic information spaces to appear throughout our everyday environments. How will this information be organized, and how will we interact with it?

Wherever possible, we should look for ways of associating electronic information with physical objects in our environment. This means that our information spaces will be 3D. The SemNet system [4] is an example of a tool that offers users access to large, complicated 3D information spaces. Our goal is to go a step further by grounding and situating the information in a physical context to provide additional understanding of the organization of the space and to improve user orientation.

As an example of ubiquitous computing and situated information spaces, consider a fax machine. The electronic data associated with a fax machine should be collected, associated, and co-located with the physical device (see Figure 1). This means that your personal electronic phone book, a log of your incoming and outgoing calls, and fax messages could be accessible by browsing a situated 3D electronic information space surrounding the fax machine.

The information would be organized by the layout of the physical device. Incoming calls would be located near the earpiece of the hand receiver while outgoing calls would be situated near the mouthpiece. The phone book could be found near the keypad. A log of the outgoing fax messages would be found near the fax paper feeder while a log of the incoming faxes would be located at the paper dispenser tray. These logical information hot spots on the physical device can be moved and customized by users according to their personal organizations. The key idea is that the physical object anchors the information, provides a logical means of partitioning and organizing the associated information space, and serves as a retrieval cue for users.

A major design requirement of situated information spaces is the ability for users to visualize, browse, and manipulate the 3D space using a portable, palmtop computer. That is, instead of a large fixed display on a desk, we want a small, mobile display to act as a window onto the information space. Since the information spaces will consist of multimedia data, the display of the palmtop should be able to handle all forms of data including text, graphics, video, and audio.

Moreover, the desire to merge the physical and electronic worlds requires that the palmtop computer and display have a spatial awareness and understanding of the physical environment along with the ability to visually mimic these environments and individual objects. Thus, the combination of a powerful computer capable of understanding and generating 3D models coupled with a high fidelity mobile display will serve to blur the boundaries of the physical and electronic worlds. Blurring this boundary, and therefore providing a seamless integration of the two worlds, will ease the way in which we interact with them concurrently.

We are investigating the use of an integrated input controller and output display unit to serve as a bridge or porthole between computer synthesized information spaces and physical objects. The research is focused on improving the communication bandwidth and the ease with which users interact with physical and electronic objects throughout their physical environments. Because ubiquitous computing will offer situated information spaces everywhere in our daily environments, we are exploring designs which facilitate a seamless integration of computer augmented data and physical objects in a highly portable tool.

## Background

Computer-augmented environments [7] offer a synergistic merging of computers and common physical objects, which can radically change the way we define human-computer interfaces, construct software applications, and design hardware systems to fit this new model of computing. While this model is beginning to take shape, it is largely unspecified. At least two approaches, however, have begun to emerge: the notion of ubiquitous computing and overlay techniques. The ubiquitous-computing model advocates embedding many small, highly specialized computers within our everyday environment. Researchers investigating overlay techniques are dealing with issues in merging two or more media types into one integrated and composite medium that offers the strengths of the combined media.

## Ubiquitous Computing

The ubiquitous-computing model [15] suggests that very small computational devices be embedded and integrated into our physical environment in such a way that they operate seamlessly and almost transparently. Not only does this model advocate miniaturizing computers but also suggests that these devices be physically aware of their surroundings. These devices emit information periodically or on request either directly by a human or in response to queries made by companion devices. Some applications of ubiquitous devices are active badges [14], which are small electronic badges, worn by people for automatic personal identification, which communicate via infrared transmitters and sensors. The badges can trigger automatic doors, forward telephone calls, and log a person's location and who that person is meeting [8 , 11].

The ubiquitous-computing model requires that we examine the design of our existing computing systems. The new model, at the very least, will stretch our notion of distributed computing in terms of network and operating system requirements.

In addition, ubiquitous computing suggests separate displays for each unit varying from no display, an audio display, to low- and high-fidelity displays. Users will need to continuously scan the environment for the appropriate display that corresponds to the information they are seeking. For example, in order to check if new messages have been left

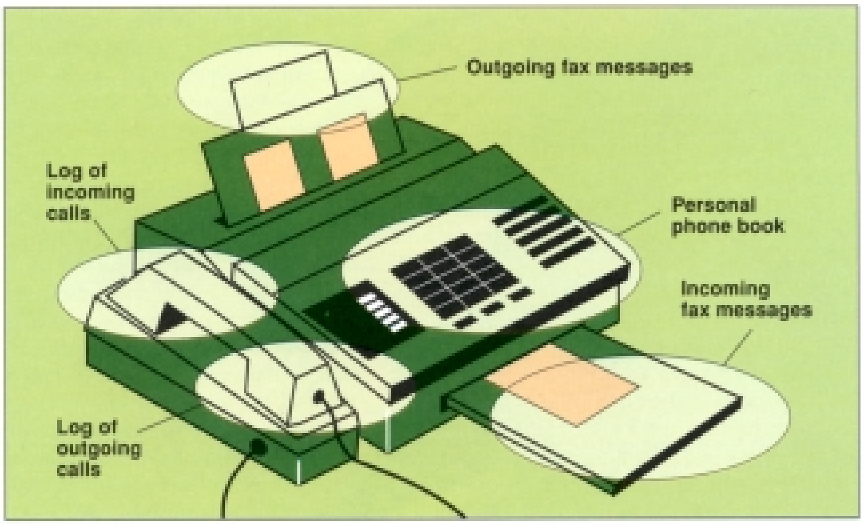


**Figure 1.** Potential information "hot spots" surrounding a fax machine: (a) log of incoming calls, (b) log of outgoing calls, (c) personal phone book, (d) outgoing fax messages, (e) incoming fax messages.

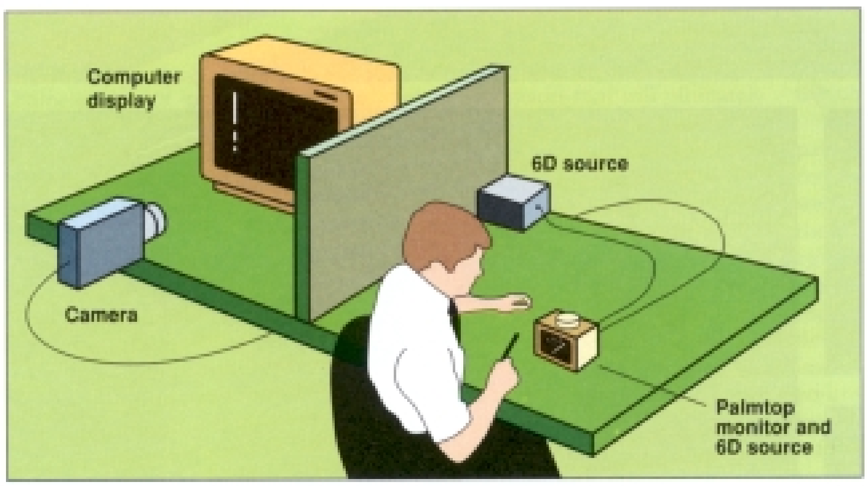


**Figure 2.** Configuration of Chameleon prototype.

on my answering machine, I must physically visit the machine and stand directly above the LCD panel to read the display. Perhaps we do not need a separate display for each computational device. Instead of relying on these environmental displays to be visible at all times, people can carry with them a personal display which could be used in conjunction or in absence of environmental displays. Users should not always have to find the embedded information themselves; the information could find the users via their personal display.

From a user's perspective, ubiquitous computers will cause a flood of information to be available—the majority of which will be useful immediately while some information may be useful at a later point in time. The information spaces will be generated by many computational objects within a user's environment. Ideally, this information needs to be viewed and manipulated within the context of the originating situation. That is, the information spaces should not be abstracted and transported to a stationary computer fixed on a physical desktop. Instead, we need highly portable displays and protocols for visualizing and filtering the electronic spheres of information.

## Overlay Techniques

Overlay techniques offer insight into how common physical objects can be augmented with computer-synthesized data to reduce the time needed to complete a task and ease the way in which we interact with the physical objects. The DigitalDesk [10] integrates traditional paper media with electronic media on a combined physical and digital desktop. A digitizing tablet and cordless pen are used for selecting and inputting data. Electronic ink and images are superimposed on paper documents and the desktop surface.

In the DigitalDesk, words in a paper document can be selected, looked up in a dictionary, and presented in an electronic window on the desktop surface. Another example allows users to select a column of numbers from a paper document and transfer them to a spreadsheet program or electronic calculator by gesturing. This design allows users to interact with paper and electronic media in a similar and consistent manner while adding computational functionality to paper documents.


**Figure 3**. Palmtop unit consisting of a video display, a response button at the top of the unit and a 6D input sensor for providing x, y, and z positional information and organization (pitch, yaw and roll).

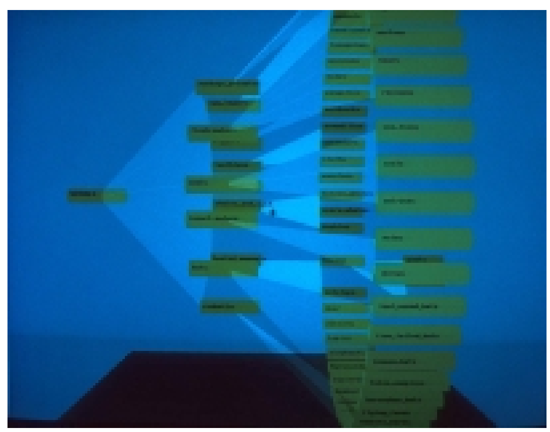

**Figure 4a.** Palmtop unit before zooming with screen shot of what the user sees in the palmtop monitor. The user is browsing a 3D tree hierarchy.

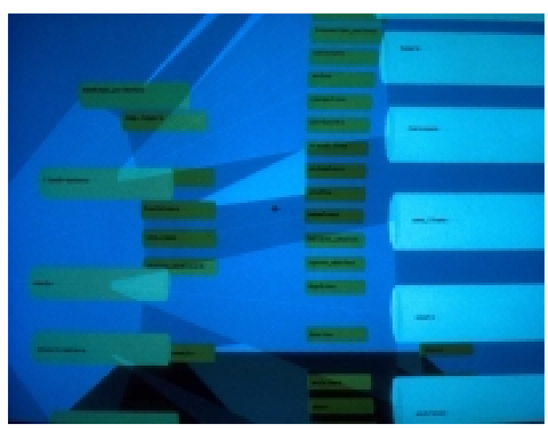

**Figure 4b.** Palmtop unit after zooming with screen shot of what the user sees in the palmtop monitor.

The object-oriented video system [13] offers another example of overlaying media, in this case video data and graphical objects. Users are able to interact with graphical objects and controllers superimposed on live video data. The graphical buttons, sliders, and knob controls are mapped into the physical devices so the user can remotely operate the controls. Selecting video objects results in invoking the corresponding graphical controller. The video shows the context and the immediate effect of manipulating the controls. In this sense, the graphical objects are tightly coupled with the physical objects. Users are not just annotating the video images; they are interacting with them. Note that these two systems have dealt with augmenting physical objects that are viewed in a 2D setting. Additional enhancements must be designed to work in a 3D world.

## Prototype

Chameleon is a prototype system under development at the University of Toronto [5]. It is part of an investigation on how palmtop computers designed with a high-fidelity monitor can become spatially aware of their location and orientation and serve as bridges or portholes between computer-synthesized information spaces and physical objects. In this prototype design, a 3D input controller and an output display are combined into one integrated unit. This allows the palmtop unit to act as an information lens near physical objects. For example, consider a geographical wall map used in conjunction with the Chameleon. The palmtop is always aware of its own physical position and orientation relative to the map; the contents of the display can respond directly to the user's gestures and movements. That is, the user sees information about Toronto or Boston while the device is positioned over the respective city on the map. Varying levels of detail and classes of information could be viewed according to different gestures made by the user.

The design achieves 3D comprehension with a very small screen size. It does this through the use of movement as a compelling depth cue. In addition, the act of movement provides a great deal of 3D sensation, as suggested in the motor theory of space perception [6].

In the prototype configuration (Figure 2), a small 4-inch color, LCD based hand-held monitor acts as a palmtop computer with the capabilities of a Silicon Graphics 4D/310 Iris workstation. A video camera is currently being used to capture the contents of the large workstation screen which is fed into the small hand unit. To facilitate input controls, a response button at the top of the device and a 6D input device (the Ascension bird [1]) are attached to the small monitor.

This design allows the system to detect user gestures, (x, y, and z positional data as well as pitch, yaw, and roll orientation from the bird) and input selections (via the response button) for issuing commands (see Figure 3). The Silicon Graphics workstation is programmed to generate a variety of 3D models and information spaces that will ultimately be positioned near physical objects.

Translation (x and y axes) and zoom (z axis) controls are available on the prototype to navigate through a 3D workspace roughly equivalent to a 3-foot cube. The net effect is that the palmtop unit acts as a window into the 3D workspace. The system is modeled after the "eye-in-hand" metaphor. For example, as the user translates the palmtop unit to the left, she or he moves toward the left wall of the 3D workspace.

Figures 4a and 4b show a before--and-after view of the palmtop unit and the contents of the palmtop screen as the user zooms into the 3D model along the z-axis. The model consists of a 3D cam tree representing a hierarchy of information where the rectangles are nodes on the tree [2, 3, 12]. Note that as the palmtop is moved from the start to the end position (Figure 5), the user sees a smooth zooming animation during the traversal of the path.

In order to allow users to select objects within the virtual world, a cross-hair cursor is fixed in the center of the screen. Users line up target objects over the cross-hair and click (or in some situations double-click) on 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 encountered is selected.

Application-specific controls can be built into the system. For example, in the 3D cam tree, users are 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 upwards causes the tree to rotate in the opposite direction.

## Lessons learned

Preliminary experimental evaluation of the Chameleon design suggests that the small palmtop screen size and gesturing-input combination offer equivalent depth perception compared to large (21 inches) static 3D displays [5]. Additional experiments will further characterize the performance of this combined input and output device. Nevertheless, the experiment revealed a few interesting observations.

When first interacting with the palmtop unit, approximately 25 percent of the users felt that the controls were completely backwards. That is, they had an object view instead of an egocentric view for the input controls. Selecting an object and gesturing to the left were incorrectly believed to move the object or the entire scene to the left. Egocentric controls dictate that gesturing to the left causes the user to move to the left, keeping the object or entire scene in a relatively fixed location. To convert users from object centered to egocentric controls, we first had them keep the palmtop at a fixed distance from their body. Then we asked them to rock forward and backward and then left to right. These physical actions directly corresponded to the navigation controls and caused the user to make a quick transition to the egocentric model.

While using the device, we noticed that users seemed to have a maximum cruising speed while moving the palmtop to browse the 3D workspace. This cruising speed was slower than we

anticipated. Moreover, none of the users complained about or observed a lag in the system. This leads us to consider alternative explanations. Specifically, we are looking more closely at the human visual processor and physical motor skill limitations. One theory is that novice users of the system who are unfamiliar with the contents of the 3D workspace may be browsing in a continuous "focused path-tracking" mode. Expert users who are very familiar with the workspace may alternatively browse in a "focused endpoint-tracking" mode. That is, experts know their final destination and are only interested in the before-and-after views, while novices are interested in observing the entire traversal path to the final destination. Future research will more closely investigate this cruising-speed phenomena.

In terms of palmtop display quality, the resolution of the small LCD monitor was considerably inferior to the large SGI monitor. While ghosting images are not a problem, the resolution of the LCD display does not easily support text. We are investigating the use of scalable fonts and alternative text presentation styles. More sophisticated rendering models to preserve the 3D scene are also under investigation. McKenna [9] has begun to explore the benefits for tracking both the head and a mobile display surface. Nevertheless, we believe the display characteristics and quality will improve over time.

**Figure 5.** The user sees a smooth animation on the palmtop screen while browsing an information space.

Improvements in tracking technology will allow the palmtop to operate in larger spaces. The current use of the Ascension bird provides detailed 6D tracking on the order of a 3-foot cube. Researchers are investigating methods for tracking on a much larger scale. The optoelectronic system being developed at UNC Chapel Hill is designed to provide similar tracking performance to the bird, but on the order of 16-by-30 feet and scalable to larger dimensions (see Azuma's sidebar in this issue).

Notice that more technologically robust configurations can be used for the Chameleon prototype. For example, using an NTSC output channel for the SGI machine instead of the external camera would be an improvement. However, using the camera allows us to easily switch between host machines (e. g., developing on a Macintosh instead of an SGI) and provides a means of quickly altering the quality of the video image. The idea is to allow for rapid prototyping and rapid alterations to the prototype framework for improved design exploration.

Although the 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 we anticipate will be available in a few years in a highly portable form. Three applications are described. The Active map is currently being investigated and prototyped while the other two are still on the drawing board.

## Applications

Given the notion of situated information spaces and the design of the Chameleon prototype, new applications are being defined to explore and reveal issues in computer-augmented environments. The intention of describing the applications is to uncover issues, problems, even styles of user interaction. Our goal, initially, is not to determine the feasibility of implementing each application but to identify interesting characteristics of the applications and how they may influence future system designs and prototypes.

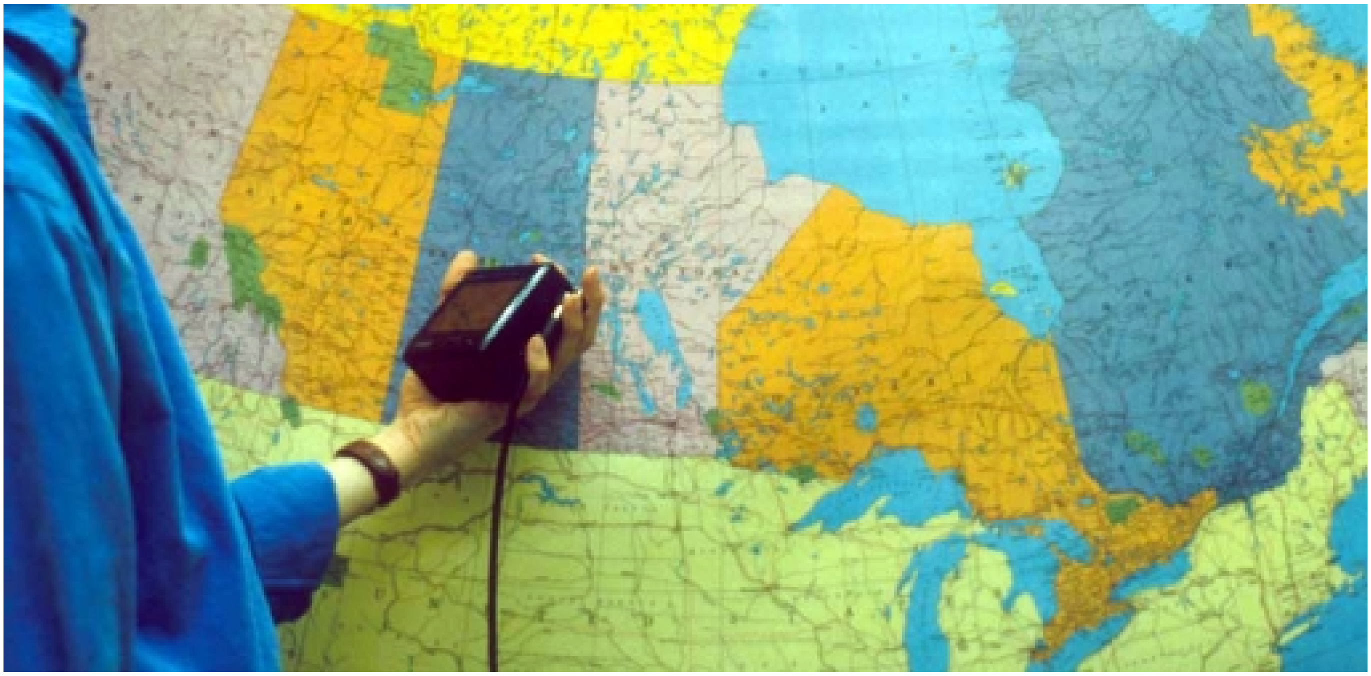

Figure 6. An active map which emits various layers of information is quickly accessed by the palmtop unit. A map of Canada is shown and the user is requesting detailed weather information for a region on the map by moving over the region of interest.

## Active Maps and Paper

The Chameleon system can serve as an electronic information lens when used in conjunction with electronic or paper-based displays. A paper display (e. g., large posters, diagrams, or maps) or a computer monitor (e.g., large 21-inch displays or even a large rear-projected screen) serves as a stationary object containing the dominant information source. As the hand unit is positioned closer to or further from the stationary artifact, varying levels of information or detail are shown within the display of the palmtop unit. For example, suppose we wanted to browse a map of Canada. A large poster or a computer monitor displays a map of the entire country. A user first selects a region by using the palmtop as a pointing device. Zoom controls are available to access more detailed information about a region, causing the palmtop display to change; the map being presented remains the same for contextual awareness and orientation (see Figure 6). A variety of information could be presented to the user depending on the orientation of the hand unit. For example, weather information, travel itineraries, and geographical points of interest could be easily accessed. More sophisticated models will allow the Chameleon and an electronic display surface to interact and react to each other.

A similar design concept could be investigated for individual sheets of paper. Suppose a paper document contains sensitive data such as a fiscal budget. The numbers in the spreadsheet will not be printed in ink but appear in electronic ink only when a user, with the proper authorization, positions the Chameleon device onto the columns containing the sensitive data.

Information spaces will be constructed not only with a single computational unit operating in isolation with physical artifacts but often with collections of computational units situated within a common environment. The computer- augmented library and portable surrogate office are two example applications that illustrate this idea by extending the model of computational objects into a computational environment consisting of cooperating objects.

## Computer-Augmented Library

As an example of a more widespread electronic information space, a computer-augmented library could offer significant improvements over a traditional library. Suppose we were searching for books and information on 19th Century music composers. Searching in a traditional library often involves two separate procedures: (1) accessing and querying a card catalog or an electronic bibliographic database and (2) tracking down books and browsing the bookshelves. In a computer-augmented library, these separate acts and procedures are combined into an integrated process. The electronic database of information is situated around the bookshelves. The shelves and the books emit navigational and semantic information. As we walk through the music section, books on the topic of interest as well as related material will be highlighted by indicator lights. Figure 7 shows a computer-augmented bookshelf with touch-sensitive LCD strips along each shelf. Regions of the strip are highlighted under books that match a user's query. Alternatively, our awareness may be directed by non-speech audio that we can hear through miniature headphones.

To aid our search, selecting a book and tapping on the book spine or on the corresponding touch-sensitive LCD strip under the book (see Figure 7 inset) will cause the search engine to focus on material similar to our selected book. The LCD strips provide users with peripheral awareness of information regions, supplying very coarse levels or pieces of information (e.g., specifying a location or displaying a word).

The Chameleon unit serves as the main focus of high-fidelity information exchange. Textual and graphical data will be presented to further enhance our experience. When requested, the table of contents of individual books will be rapidly presented to us, or a dynamically merged table of contents could be formed from a group of selected books. While many such possibilities exist, note that this computer-augmented environment does not prevent us from browsing the bookshelf in the traditional fashion. Our experience is enhanced through the situated electronic information spaces and the embedded-computer infrastructure, without sacrificing the beneficial aspects of existing technology.

### Portable Surrogate Office

The Chameleon system can be used to offer remote access to a physical environment as well as provide additional functionality. The envisioned design allows users to electronically annotate all objects in the physical environment as well as access and control "mediator objects" that have compatible interfaces working between both the physical and virtual environments.

To capture the contents of an environment, in this case an office environment, a camera is placed in the center of the room. A 360-degree panoramic video image is taken to capture a visual representation of the office (Figure 8). This image, stored

**Figure 7.** Computer-augmented library. The electronic database of information is situated around the bookshelves; the shelves and books emit navigational and semantic information. Touch sensitive LCD strips run along the shelves to select books of interest to refine the search.

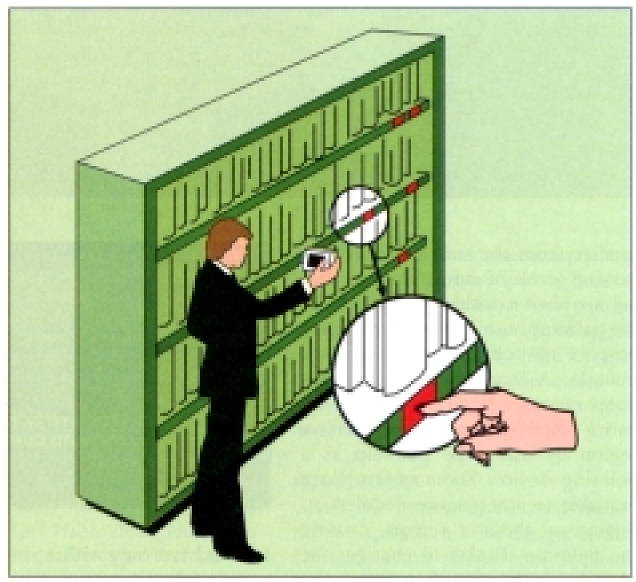

**Figure 8.** A proposed office environment with a camera mounted on the ceiling for periodically capturing the contents of the office to be stored in the palmtop. The office contains touch sensitive LCD strips (appearing in green) for accessing electronic annotations.

on the palmtop unit, allows users to browse the contents of their office environment while not physically present. Note that a more robust design would support real-time video access to the environment given access to high-capacity communication channels.

To facilitate accessing the contents of the video images, a spatial mapping is introduced in the retrieval mechanism. Since the users are very familiar with the spatial layout of their office, the proper portions of the video images are accessed by positioning the palmtop in the direction of the target objects in the physical environment. For example, in an office a desk may be found in the center of the room with a bookcase on the right-hand wall, a whiteboard on the left-hand wall, and a calendar on the front wall (Figure 8).

To view the contents of your wall calendar while you are away from the office, one would first imagine sitting in the office chair. Since the calendar is located directly in front of you, you raise the device until you see the calendar (Figure 9). To see the bookshelf, you would swing the device 90 degrees to your right; the palmtop is providing a window into this office environment. This design takes advantage of the user's persistent mental model of the office environment and provides a constant analogy to the physical interface for accessing or viewing objects.

Not only can users remotely browse the contents of the office, but additional functionality can be offered such as voice and graphical annotations. While browsing the office, users select objects from the video images by lining up the target object into the center cross-hair on the palmtop unit and clicking on the response button. At this point the user can attach a voice annotation to the selected object. A graphical "post-it" note is superimposed on the video data to remind the user of the presence of a voice annotation. The note consists of a voice annotation icon, the date, and a timestamp. For example, one could select a day in the wall calendar and leave a reminder message "Department budget meeting at 2:00" (Figure l0a) or select a book from the bookshelf and leave a voice annotation such as "Return this to Jim" (Figure lOb).

The graphical notes can also serve as anchors for making hypermedia links or associations among objects within the environment. Clicking and holding on a note will define the beginning anchor for a link. Once the device is positioned to the target end anchor, a graphical link line is superimposed on the video images. The appearance of a piece of string physically connecting the two anchor points is produced. Figure l0a shows a link made between a day in the wall calendar and the telephone.

Extending the design of the Chameleon system to support pen

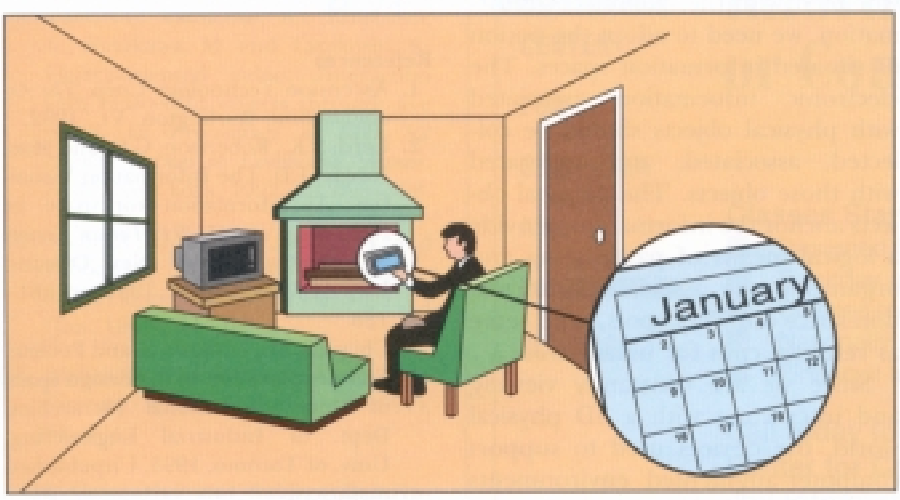


**Figure 9.** Browsing the content of your office environment while at home. The palmtop unit acts as a small window into your remote environment and makes use of spatial organizations and memory for quick access.

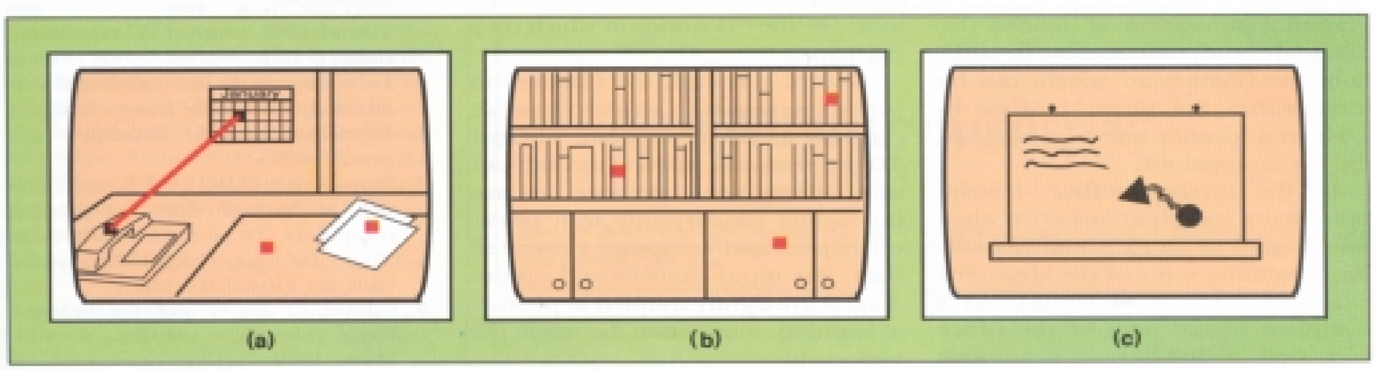


**Figure 10.** Sample views of the palmtop unit showing the combined physical and electronic information. The blue rectangles represent electronic post-it notes which have voice annotations attached. (a) shows two notes left on the desk and a hypermedia link between the phone and the wall calendar. (b) shows three notes left on the bookshelf and (c) shows the contents of the electronic whiteboard.

input allows users to scribble electronic notes or annotations on the office walls or on objects in the physical environment. A common practice may be to leave notes on the office whiteboard.

When the whiteboard is promoted to an electronic "mediator object," the notes made on the palmtop unit can automatically be transferred and appear on the electronic whiteboard. That is, when the palmtop unit is positioned on top of the virtual whiteboard, the palmtop surface acts the same as the electronic whiteboard surface in the physical environment (Figure lOc).

Light switches, thermostats, and even the telephone can be promoted to mediator objects. Users can turn their lights on or off, raise or lower the office temperature and check whether any telephone voice messages have been left on their machine.

In short, mediator objects can react and interface between the physical and computational environments. They offer the best example of integrating computers and sensors with familiar objects in a way that is compatible with the real world and current work practice.

While working in the physical office environment, users need mechanisms for remembering and accessing their electronic annotations and links. Indicator lights similar to ones found on telephone units can be used to serve the same purpose. A more complete solution involves touch-sensitive LCD strips which are placed along the office walls, the desk, and shelves in the bookcase. The LCD indicators highlight regions near annotations or hypermedia link anchors. Pressing on the LCD indicator accesses the voice or graphical annotation or follows the link to the endpoint anchor. In addition, the Chameleon system can be used within the physical office to serve as a scanner unit to locate and display annotations.

In the surrogate-office description, many complications arise when one begins to think about actually implementing some of the ideas. For example, how can we reconstruct a complete spatial map of the office with only a 360-degree panoramic video image? The resolution of the image will be crucial. Text at oblique angles may be unreadable given the video resolution. Another issue is how can we properly align the physical objects with the video data and annotations? Many of these and other issues should be explored to gain a better understanding of the requirements for situated information spaces and computer-augmented environments.

## Conclusions

Ubiquitous computing requires a qualitative change in the way we think and interact with computers. Instead of viewing and manipulating a computerized world through a large stationary computer and display, we want to shift to a new model in which we carry around a very small palmtop computer that acts as our personal display onto information spaces. These displays are aware of their surroundings and change depending on the situation in which they are immersed.

Electronic information will be available everywhere. In order to avoid being flooded and overwhelmed with the quantity of information, we need to adopt the notion of situated information spaces. The electronic information associated with physical objects should be collected, associated, and collocated with those objects. The physical objects anchor the information, provide a logical means of partitioning and organizing the associated information space into "hot spots," and serve as retrieval cues for users.

Since we are constantly viewing and interacting with a 3D physical world, the devices used to support computer-augmented environments need to be just as responsive and aware of the 3D worlds in which they participate.

Toward this end, the Chameleon prototype begins to explore how we can access and manipulate situated 3D information spaces throughout our environment. Because we are interacting concurrently with physical objects and computer-generated data, our input controls and output display are tightly coupled to provide a seamless integration between the two worlds.

This article has explored and uncovered a wide range of issues surrounding computer-augmented environments. The Chameleon prototype and a set of computer-augmented applications were described. Future research will further explore the concepts of situated information spaces, mediator objects, and user interaction styles using the Chameleon design.

## Acknowledgments

The author greatly appreciates the contributions of Mark Chignell, Shumin Zhai, Ferdie Poblete, Sarah Zuberec, Ron Baecker, Bill Buxton, Marilyn Mantei, and Paul Milgram for their suggestions and support. Special thanks are due also to Robert Gilbertson for his initial artwork and photography which aided in visualizing the concepts presented in this article.

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 Digital Equipment Corp., Apple Computer, and Xerox PARC.

**References**

1. Ascension Technology Corp. *The Ascension Bird.* Burlington, Vt., 1992.
2. Card, S.K., Robertson, G.G. and MacKinlay, J.D. The Information Visualizer: An information workspace. In *Proceedings of CHI '91 Human Factors in Computing Systems* (New Orleans, La.). ACM, New York, 1991, pp. 181188.
3. Chignell, M., Zuberec, S. and Poblete, F. An exploration in the design space of three dimensional hierarchies. Dept. of Industrial Engineering, Univ. of Toronto, 1993, Unpublished manuscript.
4. Fairchild, K.M., Poltrock, S.E. and Furnas, G.W. SemNet: Three-dimensional graphic representation of large knowledge bases. In *Cognitive Science and its Application for Human-Computer Interface.* Erlbaum, Hillsdale, NJ., pp. 201-233.
5. Fitzmaurice, G.W., Zhai, S. and Chignell, M. Virtual reality for palmtop computers. To be published in *ACM Trans. Inf. Syst. (July* 1993), special issue on Virtual Worlds.
6. Gibson, J.J. *The Perception of the Visual World.* Houghton Mifflin, Boston, Mass., 1950.
7. Kruger, *M. Artificial Reality II.* Addison-Wesley, Reading, Mass., 1991.
8. Lamming, M.G. and Newman, W.M. Activity-based information retrieval:

Technology in support of human memory. Tech. Rep., Rank Xerox EuroPARC, 1991.
9. McKenna, M. Interactive viewpoint control and three-dimensional operations. In *Proceedings of the 1992 ACM* Symposium on *Interactive 3D Graphics. Comput. Graph.* (1992), 53-56.
10. Newman, W. and Wellner, P. A desk supporting computer-based interaction with paper documents. In *Proceedings of CHI '92 Hunan* Factors *in* Computing Systems. ACM, New York, 1992, pp. 587-592.
11. Newman, W.M., Eldrige, M.A. and Lamming, M.G. PEPYS: Generating autobiographies by automatic tracking. In *Proceedings of the 2nd European Conference on CSCW*, 1991, pp. 175187.
12. Robertson, G.G., Mackinlay, J.D. and Card, S.K. Cone trees: Animated 3D visualizations of hierarchical information. In *Proceedings of CHI '91 Human Factors* in Computing Systems. ACM, New York, 1991, pp. 189-202.
13. Tani, M., Yamaashi, K., Tanikoshi, K., Futakawa, M. and Tanikoshi, K. Object-oriented video: Interaction with real-world objects through live video. In *Proceedings of CHI '92 Human Factors* in Computing Systems. ACM, New York, 1992, pp. 593-598.
14. Want, R., Hopper, A., Falcao, V. and Gibbons, J. The Active Badge Loca tion System. *ACM Trans. In* f. Sys. 10, 1 (Jan. 1992), 91-102.
15. Weiser, M. The computer for the 21st Century. *Sci. Am. 265, 3* (Sept. 1991), 94- I 04.

**CR 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*; **1. 3. 1 [Computer Graphics]** Hardware Architecture—*three-dimensional displays;* **1.3.6 [Computer Graphics]** Methodology and Techniques—*interaction techniques*

**General Terms:** Design, Human Factors

**Additional Key Words and Phrases:** Information access, palmtop computers, 3D control and displays

**About the Author:**

**GEORGE W. FITZMAURICE** is a Ph.D student in computer science (HCI) in the Dynamic Graphics project at the University of Toronto. His research interests include computer-augmented environments, ubiquitous computing, 3D input devices, and interactive 3D information visualizations. **Author's Present Address:** University of Toronto, CSRI, 6 King's College Rd., Toronto, Ontario, Canada, M5S 1A4, email: gf@dgp.toronto.edu