GAZE-ORCHESTRATED DYNAMIC WINDOWS

Richard A. Bolt
Architecture Machine Group
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

ABSTRACT

Consider a large-format display before the user, bearing a multiplicity of "windows," like little movies, the majority dynamic and in color. There are upwards of 20 windows, say, more than a person can ordinarily absorb at once. Some of the windows come and go, reflecting their nature as direct TV linkages into real-time, real-world events. Others are non-real-time, some dynamic, others static but capable of jumping into motion.

Such an ensemble of information inputs reflects the managerial world of the top-level executive of the not too distant electronic future: a world of brevity, fragmentation, variety, above all one of an overwhelming onslaught of events.

The multiplicity and simultaneity of such a display situation ordinarily would make coping with it untenable. The intent of the reported research is to introduce order and control, through the creation of a dynamic, gaze-interactive interface.

Making the behavior and reactivity of the "windows" contingent upon measured eyemovements - the point-of-regard of the observer - aims both to help the observer to cope with the onslaught of events on the one hand, yet enable on the other hand continuing close contact with that everchanging ensemble.

A simulation of such a world is described and demonstrated in the composite medium of computer, videodisc, and video special effects. Eye-tracking technology, integrated with speech and manual inputs, controls the display's visual dynamics, and orchestrates its sound accompaniments. All elements are combined to form a testbed for the conception generally, and to explore the associated human factors and stagecraft.

Key Words: Eye-tracking; videodisc; man-machine interfaces; graphical interface

The work reported herein has been supported by the Cybernetics Technology Division of the Defense Advanced Research Projects Agency, under Contract No. MDA-903-77-C-0037.

INTRODUCTION

The managerial world of the top-level executive or commander is one of brevity, fragmentation, and variety. They spend little time on any one activity, deal with a great number of problems in the course of a day, and the problems, and decisions made about them, are of many types (Cf. Keen & Scott-Morton, 1978).

In contemplating what might be a graphical information systems interface corresponding to such a management world, we envisaged a large-format display whereon are upwards of twenty, at times even up to fifty, simultaneous mini-displays confronting the user, the majority of them dynamic and in full color.

The multiplicity and simultaneity of the displays would make coping with such a visual situation untenable. Our goal would be to introduce order and control, primarily through gaze-directed commands.

The dynamic, gaze-interactive interface we wanted to assemble would function both to enable and protect. We aimed to help the observer to keep up with the onslaught of events on the one hand, performing a "filtering" function - or, more accurately, making it possible for the observer to protect himself - but, on the other hand, to enable the observer yet to keep in touch with a complex world of on-going events.

The display situation described below, which we have dubbed "Gaze-Orchestrated Dynamic Windows," offers to the observer many visual events, while at the same time permitting built-in modes of selective attention to operate. In the midst of the welter of impressions and events, the mode par excellence of directing attention is vision, specifically, where you are looking (Cf. Marks, 1978).

Eye-tracking, the knowledge of where the observer is looking, has been the object of much study (Cf. Senders et al, 1978; Monty & Senders, 1976). However, eye-tracking has been done primarily in the context of behavioral and psychological studies, controlled and in the laboratory context, rather than using an eye-tracker as a system component, with the derived point-of-regard as part of the essential information that an interactive system should have concerning the human user.

Think of it as "eyes as output."

TRACKING TECHNOLOGIES

Of eye-tracking techniques commercially available the technique least obtrusive to the observer is the "pupil-center, corneal-reflection distance" method (Young & Sheena, 1975). This method of tracking measures the fixation point of the observer, the point-of-regard as it falls in the surround, not simply the position of the eye relative to the head.

In this method, advantage is taken of two features of the eye that only change with rotation of the eye, but not with lateral or vertical displacement: corneal-reflection; center of the pupil. The distance between the two as measured are monotonically related to the observer's point-of-regard.

An invisible infra-red light source for the corneal-reflection is used together with a TV camera sufficiently sensitive in the infra-red region to detect this light easily. The TV image, which is zoomed-in close upon the pupil of the eye, is computer analyzed; co-ordinates of the pupil center and of the infra-red spot reflection are determined by examining the timing signals of the video scan.

The infra-red spot is not at all annoying to the person being monitored, and the small TV monitor and associated apparatus can be situated about a meter distant from the observer, or even further away, depending upon the configuration and the use of special lenses. To a considerable extent, the observing TV camera and the infra-red light source can be so housed and/or "dispersed" that only minimal parts of the apparatus are directly in the presence of the person being tracked.

Currently, there are two commercially available systems which use this method: the Honeywell Occulometer Mark II, which allows a cubic inch of head motion, and Mark III which allows a cubic foot of head motion; and the Gulf & Western Model 1998S, which allows up to a cubic foot of head motion. The Honeywell unit tracks ±30° horizontally, -10° to +30° vertically, to ±1° accuracy. The Gulf & Western tracks ±20° horizontally, ±15° vertically, also to ± 1° accuracy. (Cf. Young & Sheena, 1975; vendor literature).

Both units provide pupil diameter as an output as well as point-of-regard information. (Cf. Kahneman, 1973, for a discussion of the relevance of this measure of eye activity)

Another promising approach to point-of-regard tracking is the following combination: 1) a miniature corneal-reflection tracking system housed within or on conventional glasses frames (Cf. Rinard & Rugg, 1976, 1977), combined with 2) small magnetic space-sensing instrumentation (Cf. Rabb et al, 1978). The corneal-reflection tracking system provides a measure of the orientation of the eye vis-a-vis the glasses, and the space-sensing instrumentation mounted to the glasses frame provides a measure of the orientation of the glasses to the environment; the two measures combined provide the observer's point-of-regard re the environment. Such a tracking set-up is described by Foulds (1980).

While a glasses-mounted system is much less unobtrusive than the remote corneal-reflection pupil-center distance measuring systems, it does have the advantage of unrestricted vertical and horizontal range, and free subject motion in z, as well as a substantial price advantage.

We are currently investigating advantages/disadvantages of both systems for our purposes. But the chief point now is the presence of relatively unobtrusive eye-tracking in place at the interface to capture the user's point-of-regard: an interface that knows where you're looking. Which in turn implies a graphic interesting and rich enough to be worth looking at, and where direction of gaze means something in terms or its orchestration.

VIDEODISC GRAPHICS: the "WORLD OF WINDOWS"

The actual physical premises for our large-format display is a special room at MIT's Architecture Machine Group's laboratory, which we have dubbed the "Media Room." The room is about the size of a personal office. An entire wall at one end of the room is actually a large screen, thirteen feet on diagonal, which can be back-projected upon by a light-valve television projection system stationed in an adjoining room (Cf. Figure 1)

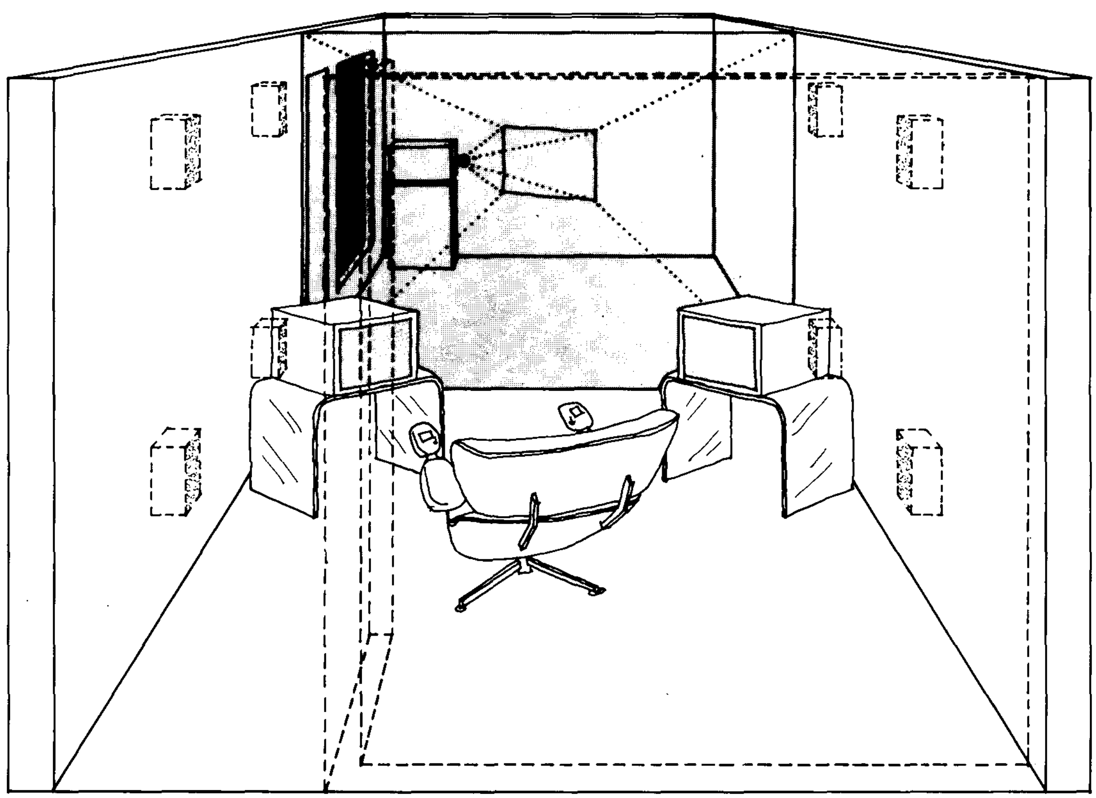

Figure 1

The observer views the screen while seated in a chair about ten feet back from the screen. We may note, because they are mentioned later on in this paper, that there are two tiny pressure-sensitive joysticks mounted in either arm of the chair.

Before the observer on this large display is a multiplicity of "windows," in full color, the majority of them dynamic with actions proceeding at full tilt, and a few of them static.

This is the "World of Windows" (WOW).

There are upwards of 15 windows, that is, more than a person can ordinarily "absorb" at one time. Some of the windows will come and go, some appear and persist, others disappear. Thus, the actual number on display at any moment will vary: 15,20,25,30 . . . (Cf. Figure 2).

The arrangement of the windows is non-regular, though we recognize that there may be circumstances where a regular grid of windows might be appropriate. The size of any window reflects the density of the information in the window, and perhaps the importance of its subject matter. There is no "structure" built into our testbed WOW, in that the subject matter of any window bears no necessary relationship to that of neighboring windows; however, there is no reason why neighboring windows might not be related by theme.

There is implicit in the fact of such a display some kind of information-gathering network "behind the scenes," deciding what is to be gathered, gathering it, placing it on view in certain spots of the large display, according to some protocol, with or without alerting messages to the observer. In an enlarged scenario, perhaps the observer somehow interacts with this human and/or automated "information-gathering network" to obtain certain requested inputs, etc. Our simulation in effect presupposes some such behind-the-scenes network, or at least a greater supportive database with which the observer can interact; however, our immediate interest is with interaction of the observer with the interface as such rather than with the details of supportive facilities.

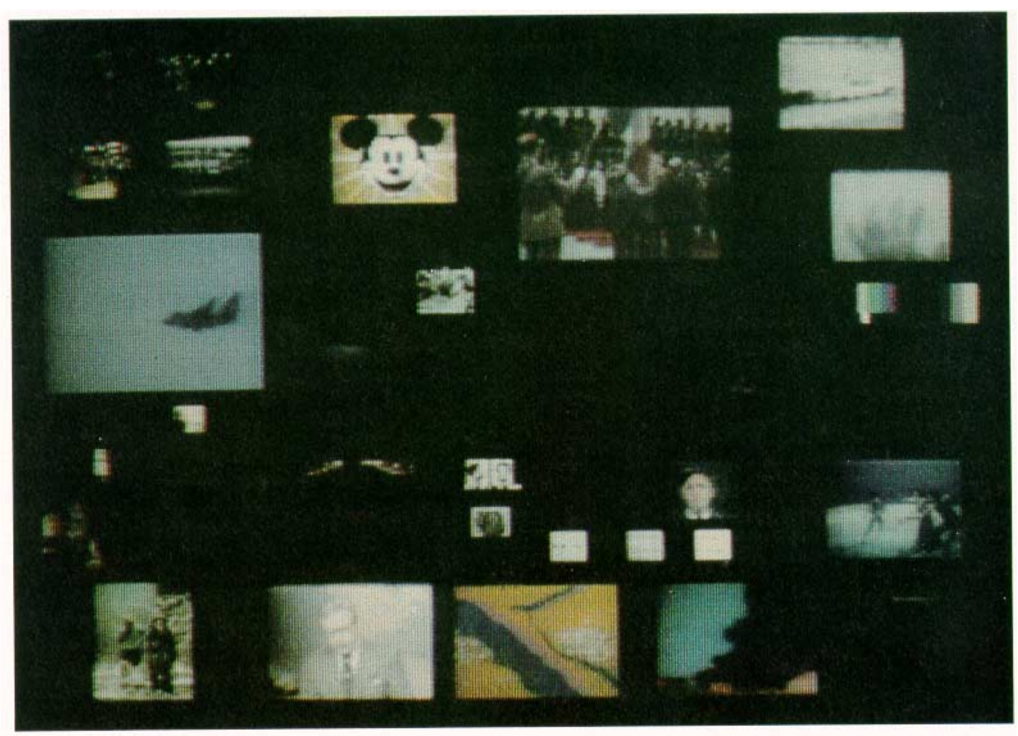

Figure 2

A representative frame from the World-of-Windows

In any case, the result of the collection of images is a realtime stream of parallel audio-visual "events" laid out in cross-section for the observer. The broad frontal display in the x,y plane of the large-screen is what offers multiplicity to the observer. The observer selects from amidst this multiplicity by where he is looking.

The Videodisc Simulator

We have made a videodisc which simulates such a real-world timestream of simultaneous on-going events:

- -conversations
- -groups of people
- -large-scale events: e.g., military maneuvers, sports events
- -newscasts
- -dynamic graphs
- -radar maps

and the like. Additionally, there are a few windows into non-real-time events: "movies" and "slides."

While the videodisc material was intended to impersonate a welter of live, real-time remote channels into ongoing events, in a more general sense, the format of the disc represents a collection of events that occur in parallel mode, juxtaposed, at a reduced scale. The events could be actually real-time events, or not, or some combination of both classes of events. In addition to being represented on the disc at a small and "simultaneous" scale, each event was to be stored on disc as an independent single event at its own level at full-screen scale.

From a graphics point of view, the structure of the disc to simulate the real-time WOW is of more interest and importance than the content of the images as such. From the point of view of interacting with videodisc, the focus becomes: 1) how the disc is used to support the necessary illusion of real-time eye-interaction, and 2) how the eye-graphics interaction itself proceeds.

First, a description of the disc, its layout, and then a description of how simulating the "World of Windows" proceeds.

Disc Layout

As a first step, the collection of movie snippets (on videodisc and videotape) comprising the WOW was selected. The briefest snippet was about 27 seconds, the longest was five minutes. Total running time for all snippets, end-to-end, was about 72.5 minutes.

All of these movie snippets, plus some stills for "slide shows" had to go on the disc at full-frame size so that the observer could be zoomed-in upon them from their little image in the WOW collage.

The length of time for the parallel play of the collage of little versions of the snippets was to be 5 minutes. This 5 minute collage section thus would contain in a spatio-temporally overlapping pattern all of the 72 minutes of the individual full-size snippets.

As one side of a videodisc (MCA Discovision) holds 30 minutes of play, clearly 2 discs at minimum were necessary. However, 3 discs were decided upon, to allow certain "stage-management," to be described below.

The layout of the 3 discs is depicted in Figure 3. All 3 discs have identical 5 minute sections of the small window collage. Then, the approximately 72.5 minutes of material was distributed on the discs, about one-third of the material on each. Only one full-frame version of any snippet appeared somewhere on one of the three discs.

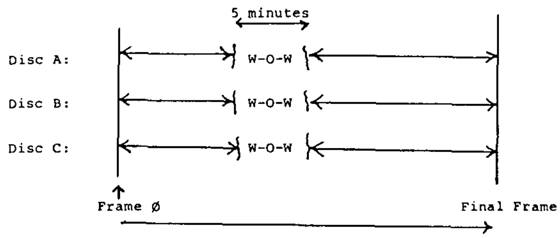

Figure 3

Each videodisc, A, B, and C has the 5-minute collage of the "World-of-Windows" laid out near the center of the disc. The rest of the space on the three discs is used to store the full-frame size versions of the video episodes (72.5 minutes in all).

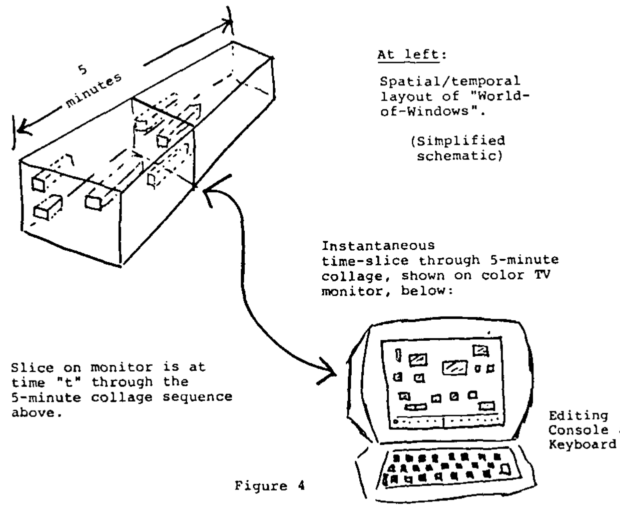

Figure 4

The editor monitor, with keyboard and graphics tablet, permitted editing into place of video snippets comprising the "World-of-Windows."

Stipulating the spatio-temporal layout of the collection of snippets onto the 5-minute collage or composite block was greatly aided by an editing program which could display on a frame-by-frame basis what any time-slice through the 5-minute WOW segment would look like in terms of general spatial layout. The editor could insert, delete, and move about small color-coded and numbered rectangles, representing any snippet. These rectangles were scaled to the desired relative size. The editor could also "play" the spatial mockup to show off the "look" of the layout as editing progressed, allowing on-the-spot changes and corrections. (Cf. Figure 4)

The little rectangles were, as noted, assigned code numbers associated with their corresponding snippets so that, upon completion of the spatio-temporal layout editing, a "recipe" for assembling the snippets at proper relative size and at proper spatial/temporal juxtaposition could be output from the editor.

Pre-mastering Assembly

Each source video snippet was then reduced to its proper relative size, and "printed" individually on videotape in its proper relative x,y position. At the conclusion of this step, each snippet was in place and sized properly, but existing each by itself on a separate tape.

The next step was to merge the separate tapes into one. The strategy was to do a pair-by-pair generational merge, to minimize image degradation by minimizing the number of generations necessary. In essence, single episodes A and B were merged with proper temporal offset, to form tape A', episodes C and D to form B', E and F to form C', etc. Then, in generation 2, tapes A' and B' were merged, again with proper temporal offset, to form tape A", C' and D' to form B", etc., and so on until through about five actual copying generations all the tapes were merged to a single five-minute tape holding all the tiny spatially-deployed and temporally-staggered episodes.

From this 5-minute composite WOW tape, plus the aggregate 72.5 minutes of material at full-frame scale, premastering tapes were made as input to videodisc mastering.

Playback Strategy

As noted above, 3 discs were made. The reason for this is due to the effects desired when the WOW is being observed.

Real-time windows are dynamic: the images are of people talking, gesturing, walking about, of planes diving, cars careening, etc. The reactivity to these dynamic windows to looking is to "freeze" when you look at them, then resume action. The "freeze" is feedback that they know you are looking at them.

The non-real-time windows e.g., a "movie", in contrast, start up when you look at them from a stopped position. You may look away and then come back; the movie resumes where you left off looking. Unlike the supposedly real-time windows, the timestream of the event waits for you.

Now, at the commencement of a demonstration of the WOW, necessarily a five-minute sample of what the experience is like, all 3 discs are played together, each 5-minute WOW segment being played in precise synch.

Should the observer look at some particular small window for a "supra-threshold" duration of time, we want to be able to freeze that window, while not the others. How??

Suppose the WOW image currently seen on the large-screen is coming off disc player A, and further that the full-screen size version of the episode resides on disc B on another player. What we do then is keep exhibiting the video stream from disc A, but matte through a freeze-framed image of the looked-at window from disc C. We do not take this freezed window from disc B as we want now to begin on B a "seek" to the full-size version of the episode in question, in particular to its start frame, plus an offset reflecting how much of the episode has already, in miniature amidst the WOW, gone by. Now if the observer acts such as to cause a zoom-in to the episode at full size, the discs (especially disc B) are set up to do so.

In general then, there are 3 discs having a central section holding the 5-minute WOW composite so that: 1) we are able to freeze (or start up, in the case of stills) some eye-selected episode by matting-through from some second disc, and 2) we can perform the 2-disc matte and keep free for "seeking" on whatever disk has the full-scale version of the selected episode for zooming-in, if required.

THE "PERSONALITY" OF THE SYSTEM

Interaction with Windows

Before getting into discussion of the visual side of reactivity, we should note that an important aspect of the reactivity of windows is rooted in sound.

One scenario is that only the looked-at window can give off its associated soundtrack, or that the soundtrack of the looked-at window is markedly louder, against a muted background of the composite sound of the other windows.

Another scenario re sound is to play the track of that window currently nearest the observer's point-of-regard, plus those of its immediate neighbors in stereo (with appropriate fall-off for distance). This tactic would make the observer's auditory locus the looked-at spot in the x,y plane. This procedure is, for the present, untenable due to videodisc addressing considerations arising out of the way we have organized our simulation.

An alternative rendering of "neighboring" at least 2 sounds, however, is to construe proximity not on a spatial basis, but on a temporal basis: images looked at most recently are juxtaposed auditorially. For example, you are looking at a newstape of President Reagan. His soundtrack amplifies, the others attenuate. You then glance over at a newsflash about the Near East. That soundtrack loudens, but you still hear a relatively loud Reagan soundtrack (the window you just left). If you return to the Reagan newsclip, then that soundtrack restores to dominant volume; but, if you stay at the Near East window, the Reagan soundtrack gradually attenuates to "background" level. ("Background" can be either silence, or the muted composite of all the other tracks together.)

Now, where and how you look will get you "zoomed-in" upon some window of interest.

Zooming Techniques

At first pass, we simply "cut" to the full-screen view: a "degenerate" form of zoom. Next, we plan to have a "special-effects" type of zoom, where the frame of the little dynamic window expands to full-screen scale. The little image itself does not expand, but stays in place at its original size (Figure 5). In the "wake" of the expanding frame, the full-size version of the contents of the little screen is revealed. The image contents of the little window at the same time fades, either to disappear entirely, or remain at the same general level of background brightness as the WOW doubled-exposured with the zoomed-in-upon full-screen image (See section below on "Double Exposure").

With appropriate new hardware, we will go into "squeeze-zooming," where the little scene in the zoomed-in-upon window expands as well, not just the frame moving outward. This is the style of zooming commonly seen on Saturday football TV.

Tiny image of, e.g., General Jones, speaking amidst the "World of Windows":

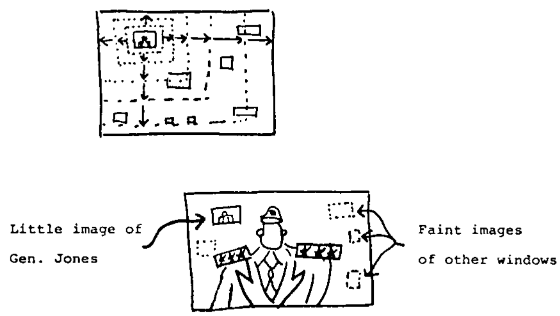

For example, as the full-screen episode proceeds, we see a composite double-exposure: we see both the full-scale image of General Jones speaking before us, plus a ghost-like image of the World-of-Windows video-mixed in. The particular little window from whence we came (the tiny image of Gen. Jones) will be among the fainter set of windows, yet distinctively highlighted to set it off from the rest.

Figure 5

Zooming in on a Window

There are at least two competing philosophies about what should cause you to be "zoomed-in" upon some window you are looking at: 1) zoom-in automatically, based upon some timing-out of how long you look at ("stare at") some certain window; 2) zoom-in upon the window you are eye-addressing contingent upon a deliberate action via an independent modality (e.g., joystick action; or a word spoken to the system via our Nippon Electric Company (NEC) DP-100 Connected Speech Recognizer).

There is appeal to both routes. The triggering of zoom by "stare" conjures up ultra-automaticity. There is a trade-off involved: the system wants to wait to be "sure" that you want to do in, yet it shouldn't hesitate too long.

However, it can be argued that while the selection of what window you want to enter is properly given in the x,y plane by the eye's point-of-regard, any decision by the system to now zoom-in, as opposed to simply letting you view the window from afar, is best given by a specific indication on your part in an independent modality. That is, "staring" to effect some action may not be a good assignment of function; it is not a use of the eyes, but a non-use over time that ties them up. We are trying alternative combinations to see what seems best.

One combination will be: if you keep looking at a window "long enough" (e.g., 4 - 5 seconds), you hear its soundtrack only, not the general sound hub-bub. Then, if you look longer, you begin to be zoomed-in. You can abort the zoom-in (if you decide you only want to listen-at, not zoom-in-upon, the window) by pulling back on the joystick of our Media Room's chair. If you decide that now you want to zoom-in, then a push of the joystick is required to over-ride your previous over-ride of the zoom-in.

If, for any window, you immediately know you want to go zooming-in, you needn't just tolerate the system's timing-out of your looking. You push forward on the joystick to say you want to zoom in.

This general scenario for zooming would seem to be compatible with a "browsing" or perhaps moderately leisurely stance or frame-of-mind with regard to the set of windows. In a more "active" or very quick-and-urgent mode, the overall system response might want to be one where you get immediately zoomed-into the window of your choice (by looking) unless you keep holding back on the joystick. The continuous holding back on the joystick to prevent you popping in, or to pop you in there by a simple "release" action, corresponds to alert, "up" stance generally.

The search for a comfortable and compatible protocol for zooming (and not zooming) is part of the research, though, and the above scenarios represent explorations, not conclusions.

Leaving a Window

There are two ways in which you "leave" a window you are currently zoomed into: 1) for a moment (temporarily), for some reason irrelevant to the display, e.g., you turn away to greet someone who has entered the room; 2) when you leave some window for the display-at-large, or for some other window.

Whether you are looking away from the entire display, or away and at another window can, of course, be detected by eye-tracker, and the system's action should be appropriate to either case.

If you look away from the display, or if looking is non-specific (i.e., "meandering about"), a reasonable guess is that you are somehow elsewhere (and thus elsewise) engaged, and the display should not change state. Thus, if you had been zoomed-in upon some particular window, the zoom should hold so that you may resume looking at that window without having again to effect a zoom-in upon it.

Or, if you are tuned into some specific window's soundtrack, but not yet visually zoomed-in, then that sound status should be held until and if you return, as indexed by eye-tracker, from beyond the display back to the same window. Of course, if you return to the World-of-Windows under such circumstances, but after some seconds (number to be determined) are not refastened to that specific window which was not zoomed-in on, but whose sole soundtrack was playing, then the soundtrack reverts to the general melange of $\underline{n}$ sounds.

In sum, the simple fact that you have left the currently attended window does not necessarily mean that it is to be "closed down" (zoomed-out from; or, if not zoomed-in, its sound replaced by all the sounds or silence). It depends upon where your eyes go: yes, closed, if you go to another window, or to the display-at-large (non-specifically about the display area); but not closed if you leave the display entirely (but temporarily) - it depends upon what you do upon your return, as discussed above.

Zooming-out

Because the observer can leave a window temporarily for some reason, and wish to stay zoomed-in for when he comes back, we need a way independent of the sheer fact of looking away to get the observer zoomed-out when he really wants to zoom-back to return to the overall set of windows. This is by pulling back on the joystick. (Or by voice command.)

The visual action need not be an explicit zoom-out, the reverse visually of zooming-in, but could be a direct return to the World-of-Windows (perhaps via a "fade"). The precise form that seems best (read: most visually plausible and pleasing) is an object of our research.

"Double Exposure"

When zoomed-in upon some window, now at full-screen scale, it will be possible yet to see the WOW in its entirety via a video "double-exposure" mix, variable from 100/0%, through 50/50% to 0/100%. The degree of mix is an independent parameter, controllable by the observer via the remaining joystick of the Media Room chair.

Thus, the observer, when in close upon some window, can yet be in touch, in at least a "pre-attentive" way, with the World-of-Windows at large: the observer pays attention to the large scene, now to the WOW, then back to the large scene. (Cf. Neisser & Becklen, 1975.)

In this double-exposure situation the observer will be able to zoom-in upon some little episode seen in the faint view of the World-of-Windows by conjointly looking at it <u>and</u> pressing forward on the zooming joystick. A press on the joystick is mandatory in <u>this</u> context to avoid the situation where an active World-of-Windows episode just happens to be lined up on the screen on the same spot where the observer is intently looking in the large image (the other half of the double-exposure).

In this diving-through-the-double-exposure situation, the action will first be the cutting out of the big image to leave the World-of-Windows only, then a zoom-into the new episode. The new episode will itself be double-exposed with the World-of-Windows at the same degree of image mix, unless the observer resets the mix level.

Non-real-time Windows

The primary difference in observer interaction with this class of window is that the event starts up when looked at.

The features of zooming-in, where you first hear only that sound from the window you are looking at, then get the full-screen picture, are the same as for real-time windows.

However, if you look away from the full-size window of one of these events, what occurs depends on which sort of two general classes of non-real-time event it is.

If it is a "slide show" type, where the slides have been changing all along, then they will go on changing independent of whether or not you are looking at them. (This is an arbitrary decision; there is no reason why they might only "flip" when you go in, and then only when you look at them . . .).

If it is a "movie" type, then it only proceeds when you are looking at it; that is, it always waits for you.

CONCLUSIONS

The world of interactive videographics has been dramatically enlarged by the advent of the optical videodisc. The planning, layout, and formatting of the videodiscs to support the simulation of a "World-of-Windows" described herein conveys the sense of programming in images.

At the user/observer interface level, interactivity is extended to incorporate where the user is looking, making the eye an output device.

While the current application of eyes at the interface is in a sense a special case, eyes-as-output has a promising future in interactive graphics, especially where the graphics contemplates the possibilities arising out of the orchestrated use of videodiscs.

ACKNOWLEDGEMENTS

Thanks are due to Chris Schmandt and Eric Hulteen, who made substantial contributions to the concepts described as well as implementing them.

Thanks are due also to Mike Naimark, who selected and edited the materials for our World-of-Windows videodisc, and guided it through its pre-mastering phase.

REFERENCES

FOULDS, R.A. <u>Studies in the use of visual line of gaze as a selection technique in nonvocal communication</u>. Proposal to National Science Foundation by Tufts-New England Medical Center, Department of Rehabilitation Medicine, Richard A. Foulds, Principal Investigator, July, 1980.

KAHNEMAN D. Attention and effort. Englewood Cliffs, N.J.: Prentice Hall, 1973.

KEEN, P.G.W. & SCOTT-MORTON, M.S. Decision support systems: an organizational perspective. Reading, MA: Addison-Wesley, 1978.

MARKS, L.E. Multimodal perception. In E.C. Carterette & M.P. Friedman, Eds., Handbook of Perception, Vol. VII (Perceptual Coding), New York: Academic Press, 1978.

MONTY, R.A. & SENDERS J.W.(Eds.) Eye movements and psychological processes. Hillsdale, N.J.: Erlbaum Associates, 1976.

NEISSER, U. & BECKLEN, R. Selective looking: attending to visually specified events, Cognitive Psychology, 7(4), 480-494, 1975.

RABB, F.H., BLOOD, E.B., STEINER, T.O. & JONES, H.R. Magnetic position and orientation tracking system. IEEE Transaction on Aerospace and Electronic Systems, Vol. AES-15, No. 5, September 1979, 709-718.

RINARD, G.A. & RUGG, D.E. Current state of development and testing of an ocular control device. 4th Conference of Systems and Devices for the Disabled, Seattle, WA, June 1977.

RINARD, G.A. & RUGG, D.E. An ocular control device for use by the severely handicapped. 1976 Conference on Systems and Devices for the Disabled, Boston, MA, 1976.

SENDERS, J.W. FISHER, D.G. & MONTY, R.A. (Eds.) Eye movements and higher psychological processes. Hillsdale, N.J.: Erlbaum Associates, 1978.