Powers of Ten Thousand: Navigating in Large Information Spaces

Henry Lieberman

Media Laboratory
Massachusetts Institute of Technology
lieber@media.mit.edu

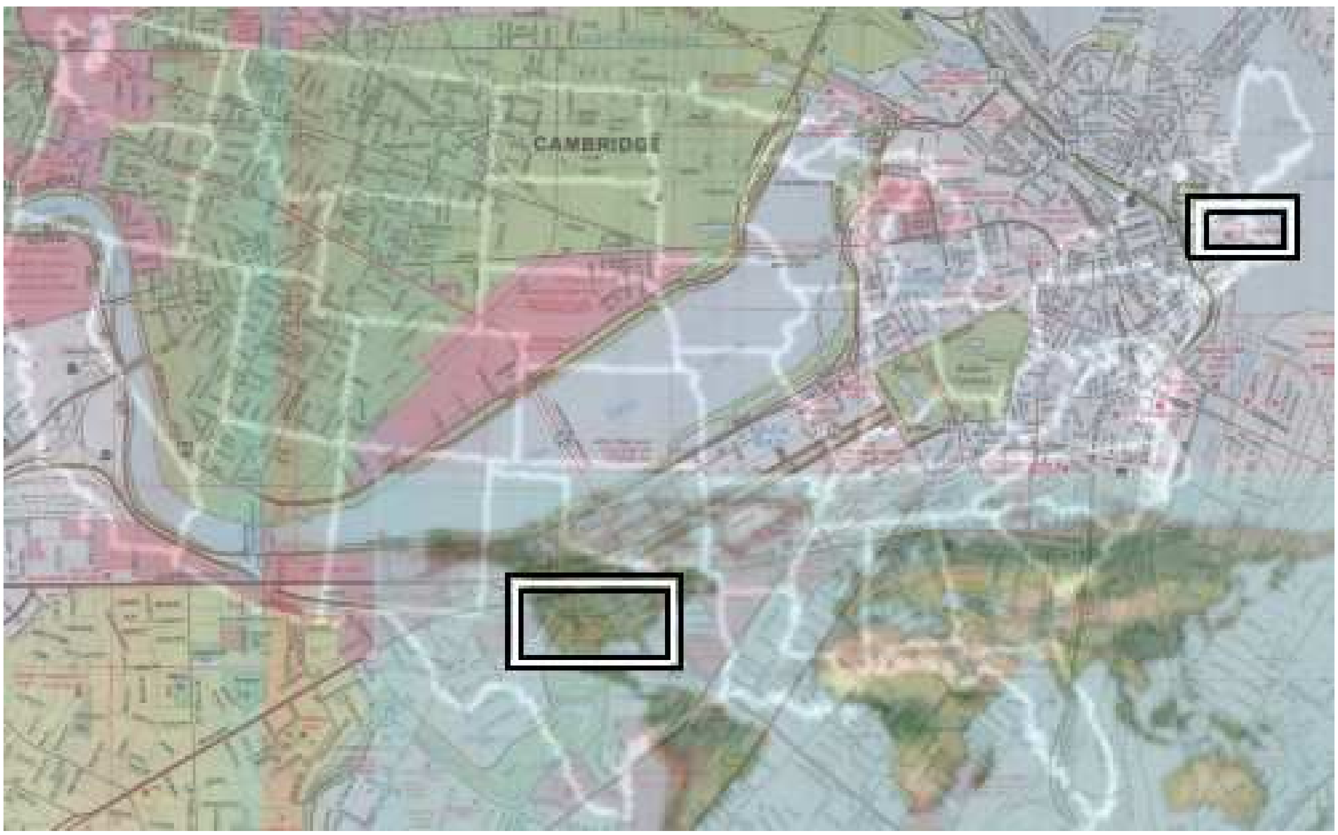

ABSTRACT

How would you interactively browse a *very* large display space, for example, a street map of the entire United States? The traditional solution is *zoom* and *pan*. To get from a corner in Cambridge, Mass. to a corner in Palo Alto, California, the user would zoom out the Cambridge street map until the entire US fit on the screen, pan it over to center on the San Francisco Bay Area, then zoom in on the Palo Alto street map. But each time a zoom-in operation takes place, the context from which it came is visually lost. A close-up zoom of a street map looks roughly the same, whether in Massachusetts or California. When changing locations, sequential applications of the zoom-in and zoom-out operations may become tedious.

This paper proposes an alternative technique, the *macroscope*, based on zooming and panning in multiple translucent layers. A macroscope display should comfortably permit browsing continuously on a single image, or set of images in multiple resolutions, on a scale of at least 1 to 10,000.

Powers of ten thousand

The book and film *Powers of Ten* [Morrison and Eames 82] try to instill an appreciation of the scale of the physical world by a succession of images each on a scale differing from the next by factors of ten, from atomic to galactic perspectives. Each image shows the scale of the next smaller image by a rectangular viewfinder placed at its center.

The power of *Powers of Ten* is that it brings together in a single work phenomena that occur over a wide range of spatial scales, and forces us to think about the relationships between them. But a book is a static artifact, and a film is not interactive. Can we provide interactive tools that help people grasp phenomena that occur over widely disparate spatial scales?

In this paper, I will explore a new technique, the *macroscope*, that involves taking the visual device of *Powers of Ten* and compounding it, using the computer's ability to combine, change and display images interactively

The macroscope: translucent zooming and panning

Let's start by analyzing a typical zooming operation. The user can choose a smaller subset of the screen [with the same aspect ratio, and here, a tenth of the area], which we will call the *viewfinder*. The zooming operation blows up the viewfinder to fill the entire image. A zoom out operation does the inverse transformation.

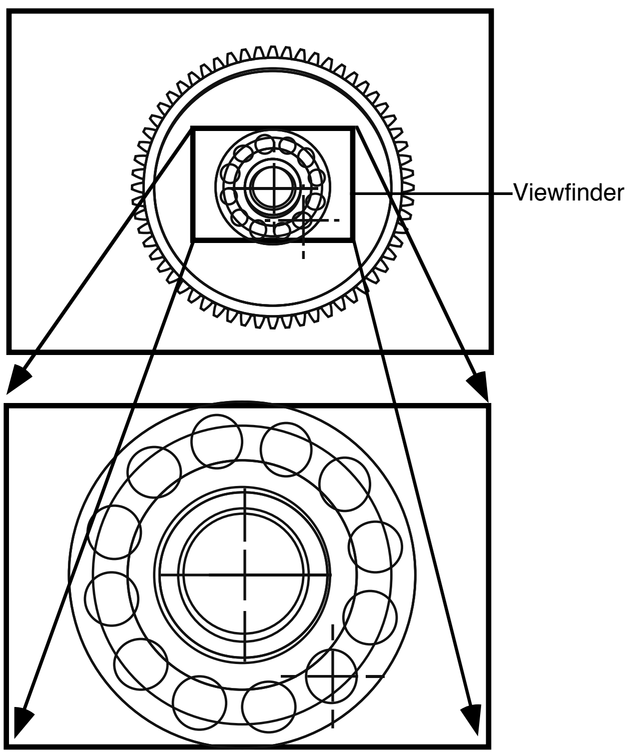

One problem with this is that after the zoom operation, the viewer loses the context of where the blown-up image came from. Our solution is to make the zoomed-in and zoomed-out views share the same physical screen space by displaying them in multiple translucent layers.

Recent experiments [Colby, Scholl 91] have shown that it is feasible to combine multiple layers of information on a single display, using translucency, focus and other image processing techniques to visually combine layers while retaining the integrity of the individual components. In one experiment, up to 39 layers of information can be dynamically controlled and presented interactively without loss of clarity.

Here, we display the zoomed-out view [where we came from] faded-out in the background, with the zoomed-in view [where we are now] in the foreground.

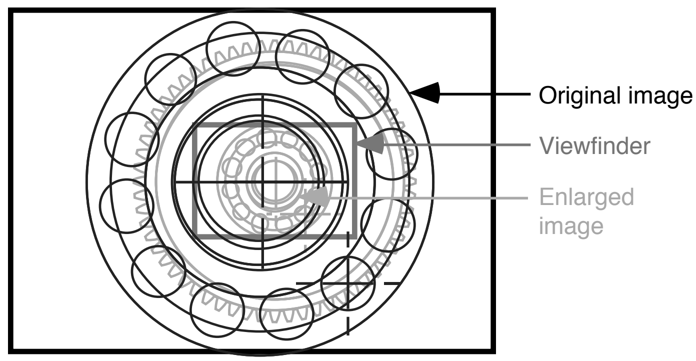

This effect can be repeated in multiple layers. If, as in *Powers of Ten*, the viewfinder is ten times smaller than its source image, four layers give us a scale of 1 to 10,000.

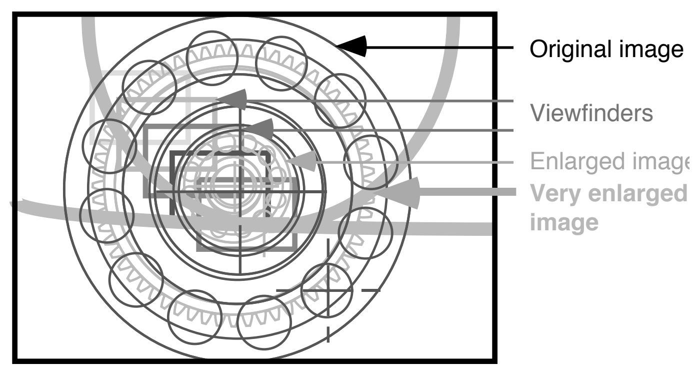

Each layer need not be strictly an enlarged subset of the previous layer. In fact, the technique is probably most effective when the data space has information at multiple resolutions, and in multiple formats. Charts or graphs can selectively reveal more data as their level of detail increases. Maps could be stored at multiple feature scales, and in formats such as political, topographic, or resource maps. Moving to higher-resolution layer can change the format and resolution of the data. The differences in resolution and format help to enhance the visual distinction between layers.

The illustration at the beginning of this paper shows an example of a three-layer zoom from the world to a street map of Cambridge, Mass. The street map is in the foreground. Superimposed is a white outline map of the states of the US, and in the lower-right corner, a topographic map of the world. The US and world maps are blurred to make them recede into the background. There are two viewfinder rectangles, one choosing the US from the world map, another the local Boston area from the US map.

Interactive control of the macroscope

But the real payoff is in interactively controlling the multiple layers. Change of position of the rectangle corresponds to a pan operation; change of size of the rectangle corresponds to a zoom operation.

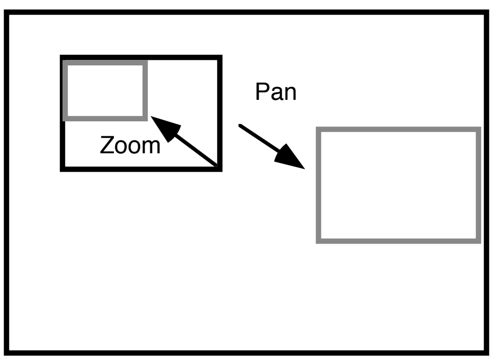

The visual effect of the pan operation will be that of imagining the physical screen as a window onto a virtual space many times larger. When you drag the viewfinder it is as if the virtual space is sliding behind the physical window. As parts of the image "scroll off" the screen your visual memory still imagines them there.

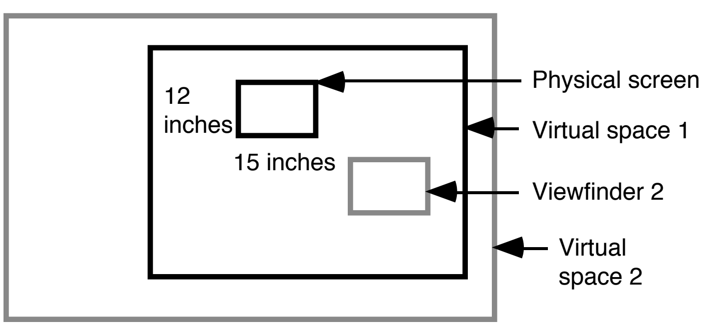

It is as if, given our four-layer case, the screen were a window onto a space physically measuring about two by three miles.

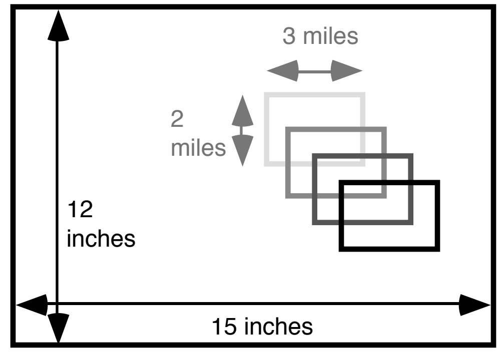

The fact that multiple scales can be seen simultaneously will making navigating in such large-magnification spaces much easier. The viewer is always oriented in the space.

With conventional scrolling, large changes in position at high magnifications are disorienting. The alternatives are either to scroll across a large distance, which may take time; set the scrolling speed to be very rapid, which leads to disorientation; or do sequential zoom-out-then-in operations, each of which loses visual context. With the macroscope, you just select the viewfinder in the layer that is at the appropriate scale and adjust it, without losing sight of your position at any magnification.

The system can make the correspondence between viewfinders and their layers more apparent by automatically highlighting the layer to which a viewfinder belongs as soon as you touch any of its sensitive points, or de-emphasizing the viewfinder layers that do not participate in the operation interactively underway.

The smoothness of the visual illusion can be preserved even if some of the most detailed layers are very sparse. Missing data in sparse layers simply drops out of the multi-layer display, without disturbing the context at other layers. Since most users would not ask for this data, they would never notice its absence. Street maps of the 100 top US cities could be included without having to store maps of every Kansas cornfield.

Dynamically adjusting the translucency levels between layers, as in [Colby and Scholl 91] is an important tool for selectively emphasizing or de-emphasizing information at user control. Layers distant from the user's present position should be de-emphasized but remain accessible. A few layers can be viewed at a time, chosen out of a large set of potentially accessible layers. Allowing user control of the translucency of individual layers through a set of sliders can help selectively reduce the visual clutter of multiple layers, especially when the visibility of the layers depends upon the nature of the image displayed.

Examples of macroscope images

Here are some examples of images from an interactive macroscope implementation for the Macintosh.

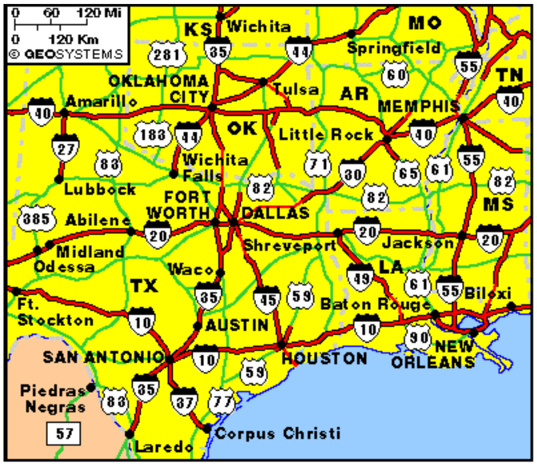

We start with a map of the south central US. Selecting a rectangle from this map, we construct a two-layer macroscope that focuses in on the Oklahoma-North Texas area. A viewfinder rectangle outlines the area selected.

Superimposed on the original map is an enlarged image of the viewfinder area. This creates some visual clutter, but since the sizes of features [roads, city names] differ in the two layers, each can still be reasonably distinguished. In this image, the poorer resolution of the enlarged map actually helps in distinguishing the enlarged layer from the background. Even when multi-resolution images are used, and enlarging an image preserves resolution, it is worth artificially changing the perceived resolution by blurring or fading to aid in separating the layers.

Indeed, the ability of the human visual system to discern and associate features at different scales is what makes the visual language of maps effective at conveying information [Wood and Fels 92]. This is a visual analogy to the well-known "cocktail party effect" that allows people to pick out single voices in a room full of simultaneous unrelated conversations.

The static images don't fully convey the effect of seeing the layers appear interactively, which also enhances legibility. Seeing the moving zoomed-in layer against the static zoomed-out layer acts to visually distinguish the two layers. As the viewfinder is dragged, the scale of the zoomed-in view changes size, appearing to "stretch" a translucent sheet across the background map.

Next, we show the effect of panning the viewfinder rectangle to the New Orleans area. The background remains the same, but the superimposed layer changes. Interactively, if the machine is fast enough, the visual impression is that the top layer "slides" across the background.

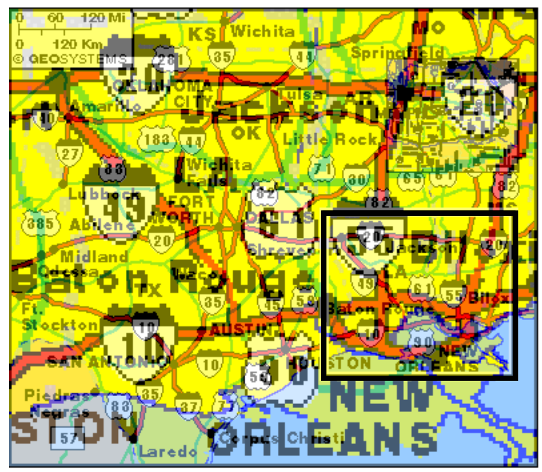

The relative translucency of the layers can be dynamically adjusted to emphasize either the higher or lower layers. In this way, attention can be shifted between looking at the focus of interest and the general context. Below, the macroscope above emphasizing foreground and background layers, respectively.

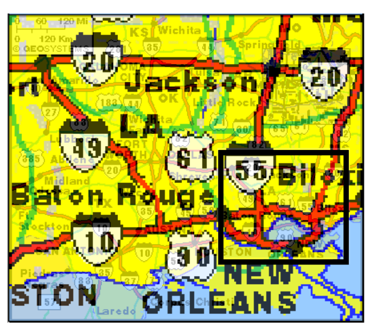

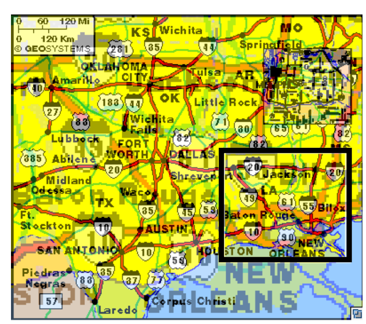

Next, we show the effect of zooming the rectangle centered on the Oklahoma-North Texas area. This creates a denser display since the two layers are now closer in size.

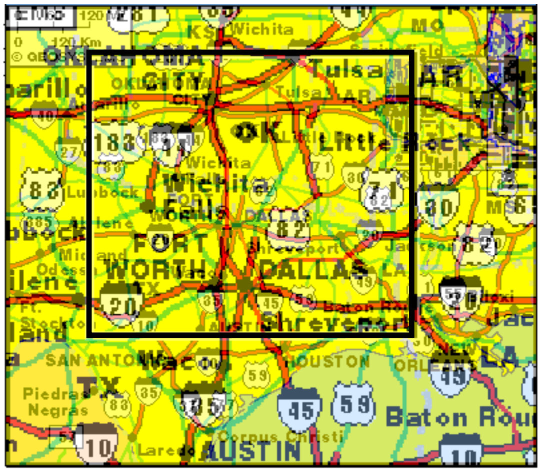

But why stop at just two layers? Selecting a rectangular portion of this image produces a three-layer macroscope, focusing in on the Dallas-Fort Worth area. Here, while the three layers are still discernible, the limits of the resolution in the map [and of black-and-white reproduction] are starting to become apparent.

Both viewfinder rectangles can now be zoomed or panned independently. Changing a viewfinder in a layer other than the topmost automatically invalidates all layers above the one chosen, on the rationale that each level's choice is dependent upon those below it. [However, we might to choose instead to retain previous viewfinder choices.] The viewer can distinguish which viewfinder belongs to which layer since rectangles that appear in lower layers have a more faded appearance.

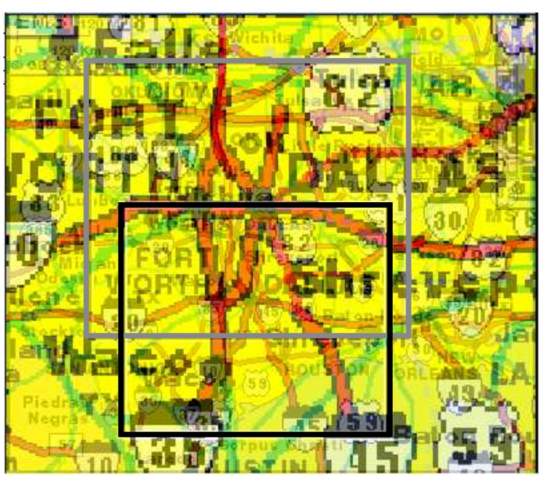

The preceding example showed a macroscope where the map was at a single resolution. We can also use multi-resolution maps, so that zooming into a map brings up a map of higher resolution.

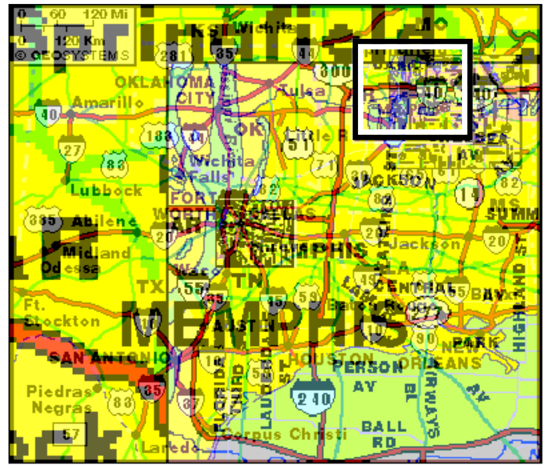

Here, as we zoom into the Memphis area, we bring up a map of the metropolitan area, showing the river adjoining the city, the beltway around it, and major thoroughfares.

The small inset map is a still-higher resolution map of the downtown Memphis area. By dragging out just two rectangles, we are able to go from a map showing a significant area of the United States, to a map showing every downtown street.

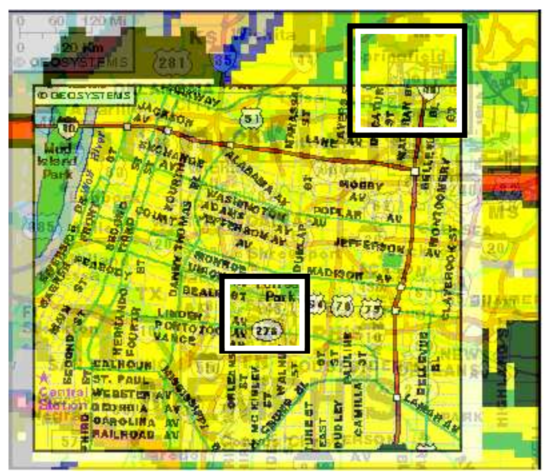

Not all kinds of images are equally amenable to this technique. Photographic images may not fare particularly well, since the superimposition may disturb subtle contrast changes necessary for recognizing the image. Maps are particularly good because their features [lines for roads, dots for cities, text labels] are relatively sparse, and high contrast on solid backgrounds.

Maps are often themselves designed to present multiple layers of information to the viewer, taking advantage of the perceptual tendency to group similar colors or features into a virtual "layer". We also intend to experiment with maps in which these layers have been separated out and can be independently manipulated.

Another example: Graphic display of a hierarchical file system

To illustrate the macroscope concept in another domain, we take the Macintosh Finder, a graphical display of a hierarchical file system. Files and directories are represented by icons, which are laid out on a finite two-dimensional plane. In the conventional Finder, each directory [folder] icon can be "opened", which pops up a window which contains icons of each of the files or folders within it.

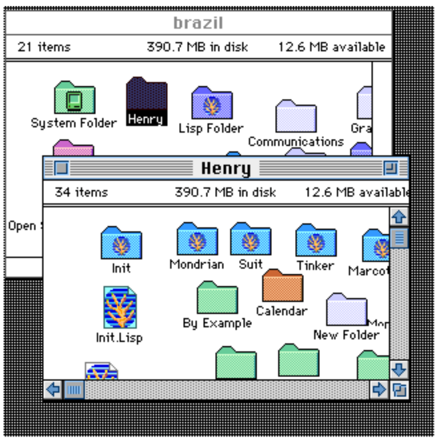

In the macroscope version, we can replace the "opaque" folder icon with an icon that graphically contains all of the files and folders within it, at a much reduced size.

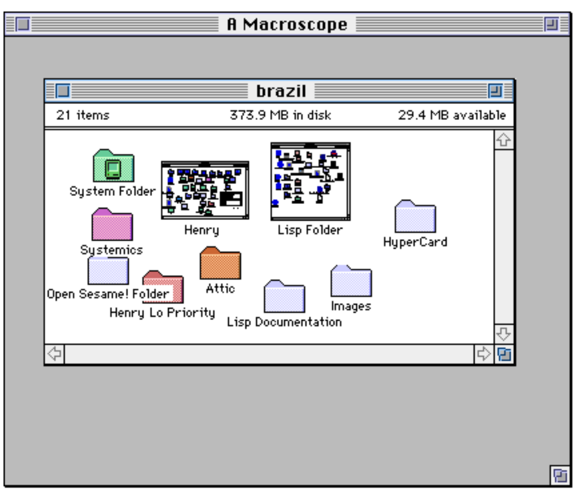

There is no opening or closing of folders, just zooming into the contents of a folder, in the spirit of the original Spatial Data Management System [Donelson 78]. As with the macroscope for the map, zooming into to a subset of the file space results in a translucent display of the region of the viewfinder. Here the larger scale of the zoomed-in icons clearly separates them visually from the background.

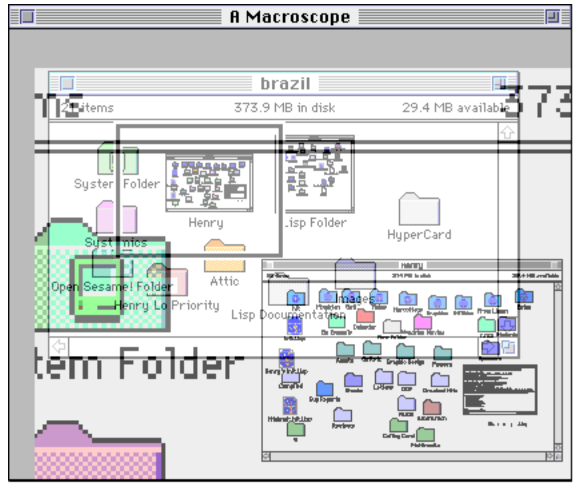

At yet another level, we can zoom into the contents of an individual file containing text, in this case, the very text of the program that is performing the display!

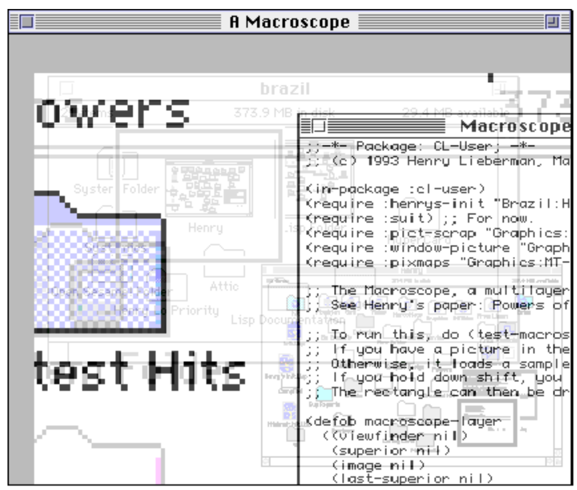

In some respects, this example works better than the map, in that the features of the icons are much more sparse, allowing greater leeway for viewing other layers translucently through the icons.

Future work

One line of future work is to implement the interactive controls over multi-layered display, and image processing techniques as described in [Colby and Scholl 91]. This will greatly increase the number of layers that it would be practical to deal with in the implementation.

Three-dimensional versions of the macroscope are also envisionable. A "two-and-a-half dimensional" macroscope would use the third dimension to separate the individual macroscope layers. The viewfinder in the 3D case would be a cube. It would be highly desirable to have true 3D stereo viewing and a 3D input device. Since the macroscope "overloads" the 2D image, a 2D perspective display of a 3D image and a 2D mouse would probably make perception of the individual layers and control of the 3D viewfinder too difficult.

User testing

We have not yet performed any formal user testing on the macroscope versus conventional zooming and panning. The chief obstacle to immediate user testing is the speed of the operations. On our current machine, a Macintosh 840av, the operations are fast enough to demonstrate the idea of interactive control [ranging from 4-5 frames/second on a small single-resolution 2-layer macroscope to 4-5 seconds/frame for a large 3-4 layer multi-resolution macroscope] but not fast enough to permit direct comparison with traditional zooming and panning. The current implementation uses Macintosh Pict graphic structures, and requires creating two full-image sized 32-bit picts per macroscope layer per mouse movement. Another implementation is in progress on SGI platforms, which should provide smooth real-time interaction, as well as permit stereo and other techniques. In the meantime, the only alternative for testing would be to compare it against zoom and pan that had been artificially slowed to match the pace at which we can currently run the macroscope.

A fair test of the macroscope would be to test the speed of locating a particular spot in the map or graphical space, where the location was not known in advance and had to be visually searched for. This could be directly compared with traditional zooming and panning. Another test would be the speed of changing the point of view from one spot to another, or switching between views of various scales. The macroscope could also be compared with using multiple windows to spatially separate the views.

Related work

Surprisingly, there has not been much work on improvements to the standard zooming and panning operations for navigating spatial information. Use of a 2D plane of spatial information with interactive control of scale was introduced in [Donelson 78]. [Perlin and Fox 93] is a contemporary example that allows a conceptually infinite plane, substituting higher resolution subsets of information as their scale becomes relevant. The use of a single viewfinder for zooming and panning appears in many systems, and a study [Beard and Walker 90] shows its superiority over alternative navigation techniques. The closest work to the present is [Takeda, et. al. 92] but their work was not generalized to arbitrary images and layers. The reference, unfortunately, does not contain pictures of their system. [Furnas 86] is discusses the importance of providing visual emphasis in a computer display for conceptually important information.

[Colby and Scholl 91] explores the use of translucency and other visual techniques, and dynamic control of emphasis, for presenting multi-layer information.

The macroscope can be considered a kind of *see-through interface* [Bier et. al. 93], where the "toolglass" covers the entire viewing surface. Whereas we use translucency to combine layers, Bier et. al.'s interface depends on seeing through transparent pixels in the tool image and the fact that the tool image can be moved by the cursor to avoid obscuring the underlying image on which the tool operates. Their discussion of "magic lenses" suggests that we might generalize the macroscope technique so that the transformations that occur between macroscope layers could be more complex than just the display of a magnified image. We could also generalize the concept of a viewfinder beyond a simple rectangle to more sophisticated techniques for selecting a subset of an underlying image space.

Acknowledgments

Major support for this work comes from research grants from Alenia Corp., Apple Computer, ARPA/JNIDS, and the National Science Foundation. The research was also sponsored in part by grants from Digital Equipment Corp., HP, NYNEX, and Paws, Inc.

References

[Bier, et. al. 93]
Eric Bier, Maureen Stone, Ken Pier, William Buxton, Tony deRose, Toolglass and Magic Lenses: The See-Through Interface, SigGraph Conference, Anaheim, California, 1993.

[Beard and Walker 90]
David Beard and John Walker, The Visual Presentation of Information, Behavior and Information Technology, Vol. 9, No. 6, 1990, p. 451-466.

[Colby and Scholl 91]
Grace Colby and Laura Scholl, Transparency and Blur as Selective Cues for Complex Visual Information, SPIE Conference, March 1991.

[Donelson 78]

William Donelson, Spatial Management of Information, SigGraph 78, Atlanta, GA.

[Furnas 86]
George W. Furnas, Generalized Fisheye Views: Visualizing Complex Information Spaces, CHI '86, p. 16-23.

[Morrison and Eames 82]
Philip Morrison, Phyllis Morrison. C. Eames, R. Eames, *Powers of Ten*, Scientific American Press, 1982. Also Pyramid Films, Santa Monica, CA 1978.

[Perlin and Fox 93]
Ken Perlin and David Fox, Pad: An Alternative Approach to the Computer Interface, SigGraph 93, Anaheim, CA.

[Takeda, et al. 92]
N. Takeda, K. Kawai, M. Koyama, A. Shiomi, and H. Ohiwa, KJ-Editor, An Index-Card Style Tool, in Proceedings of IEEE Symposium on Visual Languages, Seattle, Washington, September 1992, p. 255-7.

[Wood and Fels 92]
Denis Wood and John Fels, *The Power of Maps*, Guilford Press, New York, 1992.