# Design for a 3D Spatial User Interface for Internet Gopher

Mark P. McCahill
Distributed Computing Systems
University of Minnesota, USA
E-Mail: mpm@boombox.micro.umn.edu

Thomas Erickson
User Experience Architects' Office
Apple Computer, Inc., USA
E-Mail: thomas@apple.com

**Abstract:** Internet Gopher is a popular Internet-based information system. However, the current Gopher user interface has some problems that we believe can be addressed by a 3D spatial user interface. For instance, collections of large numbers of items or documents are poorly represented by menus since it is difficult to visualize clustering of related items, and the opportunities for representing traces left by other users (and other forms of social interaction) are limited. This paper describes an initial design for a 3D spatial interface for Gopher and some of the user interface tradeoffs necessary to reach an initial design.

## 1 Introduction

The prospect of addressing some of the limitations of the current Internet Gopher interface and observing the reactions to and usage of the interface in a broad community motivated us to design a 3D interface for Gopher [Alberti, 1992]. Although 3D interfaces have been getting increasing attention from researchers, little attention has been devoted to the design tradeoffs and problems inherent in developing a 3D user interface for general deployment on the Internet. Most of the 3D spatial interfaces in existence — particularly those that provide interfaces to information spaces — are in research laboratories; however, the advent of RISC-based personal computers now makes it feasible to develop 3D-based interfaces for general use. We believe that 3D interfaces are important to explore because they offer a number of new prospects for better representing the many dimensions of information and meta-information in a compact and natural way. We also believe that in the long run information spaces will also be used as social spaces, and that 3D scenes can serve as a natural framework for supporting social interaction [Erickson, 1993]. We intend to refine this initial design, implement it, make it freely available over the Internet, and observe the reactions to and usage of the interface.

## 2 Problems and Prospects

While the current Gopher user interface is popular, there are three well known usage problems: the *Lost-in-Space* problem, the *Grouping* problem, and the *Browsing* problem. *Lost-in-Space* users complain of feeling lost after navigating for a while and have difficulty remembering where they found an interesting item. In part, this due to the absence of any global representation of the structure of information hierarchy, and in part because the path followed by a user is either invisible or, at best, implicitly embodied in a stack of directory windows. This problem has been observed in a number of other domains, most notably hypermedia. Users need an overview of Gopherspace, within which they can see their locations and the paths they have followed. If customizable 3D scenes are used to represent Gopher directories, we can provide the user with a stronger 'sense of place' and partially address this issue.

The *Grouping* problems occurs because within a directory it is difficult to show relationships between items represented in a linear list. Some server administrators resort to putting items with blank names in their Gopher directories to group clusters of items. This problem also occurs in lists of results generated by search engines. The results are typically sorted by relevance (as ranked by the search engine), but the current user interface doesn't have a good way to convey their relative relevance. As with directories, it is difficult to show the clustering of related sets of documents. Ideally, both relevance to the search terms and "closeness" to other documents (along a variety of

user-specifiable dimensions) ought to be visible to the user at a glance. The placement of objects in 3D scene could provide an intuitive representation for clustering and relative closeness of documents.

*Browsing* problems occur because documents reflect so little of their content. All that is available is the item's name and the information about the document's type embodied in the icon. The user's only other option is to open the document—often a time consuming process—and see everything in the document. Neither option supports browsing: users need to see more information about the content of a document without there being so much that they are unable to compare and contrast different documents. The surfaces of a 3D object give us a place to display much more meta-information and proxies of the content of the item.

A fourth problem is backward compatibility. There is a large installed base of Gopher servers and clients and backward compatibility of any new client is essential for acceptance. Since it is impossible to change all of Gopherspace overnight, any new client must handle both servers that have stored additional information (e.g., about relative clustering of items in document collections) and old, unmodified servers. This must be done without stepping outside the new user interface metaphor.

In addition to addressing problems, there are a number of intriguing prospects which a new interface could open to exploration. For instance, interaction traces would allow users browsing a large information space could benefit from the activities of previous visitors. Users are often interested in knowing about the relative popularity of documents, as judged by the frequency with which they are viewed or copied. The idea behind interaction traces is that this could be reflected in the visual appearance of the document's icon. A number of researchers have explored ways in which visible representations of computational objects can reflect traces of the ways in which users have interacted with them[Hill, 1991] [Wroblewski, 1994].

## 3 The Preliminary Design

In this section we sketch out the preliminary design for the interface, with some comments on the rationale for various decisions. It is important to emphasize that this is the starting point, not the ending point. We take it as a certainty that as we proceed both implementation constraints and feedback from users will shape the design in major and unforeseen ways. We expect the implementation of the design will proceed through (at least) two stages: first, we will focus on creating a 3D information space; only later will we integrate the functionality necessary for supporting social interaction. For expository purposes we will describe the preliminary design in terms of three levels of representation — the overview, the neighborhood, and the individual item. We will begin with the neighborhood, the representation for the collection of items with which the user is interacting. Next we will examine some details of the 3D icons used to represent individual items. Then we'll look at the representations which provide the overviews of regions and of Gopherspace as a whole. Finally we will describe interactions in Gopherspace, including: opening documents, navigating around a neighborhood, moving from one neighborhood to another, and user-user interaction.

### 3.1 The Neighborhood

The neighborhood is the representation of the collection of items with which the user is interacting. Neighborhoods are either constructed by server administrators (a directory in the server's hierarchy) or generated by search engines in response to a user-entered query. We explored two representations of items in neighborhoods: circles and 'streets'.

Circular arrangements of items have a number of strengths. First, the user will generally have a straight on view of the fronts of several 3D icons, which is valuable because the fronts of these icons will generally contain details of their content. Since the user enters the neighborhood near the center of the circle of icons, the user is always going to be looking at something. A second useful property of a circular arrangement is that it is easy for the user to understand: the user can quickly get an idea of how many icons are in the neighborhood (based on the density of icons and the radius of the circle). The circular arrangement of icons defines an enclosed space that may be used as a collective gathering space for users. The circular arrangement also defines a center point, at which we will place a 3D 'kiosk' icon that will function as the user's link back to the previous neighborhood, and as a sort of information desk for the neighborhood. If we allow for display of user-entered comments from the people who have visited this directory this will also appear on the kiosk.

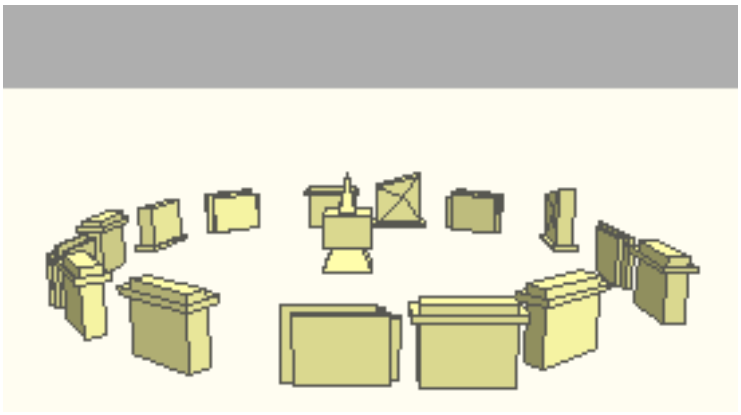

Figure 1. The circular ‘Stonehenge’ icon arrangement.

The street metaphor was investigated and rejected because the user is either facing down the axis of the street and has an oblique view of most of the faces of the icons, or is facing one side of the street and is required to turn fully around to view the next closest icon. It also may be difficult for users to tell how long a street is, and unless the street is short, it really has no center or natural gathering point. Finally, simple mock-ups of streets in a 3D modeling program resulted in arrangements that felt very claustrophobic, since fairly large 3D icons were necessary for information display purposes.

A variant of the circular arrangement of items is the spiral.We intend to use the spiral arrangement for collections of icons generated by search engines in response to user queries. Formally, the spiral has a family resemblance to the circular arrangement so that it too defines an enclosed area with a center point; at the same time, its greater openness and dynamicism seems a good reflection of the transitory nature of most queries. Also, a spiral has directionality, and provides a natural ordering within which the relevance of items to the query can be reflected. The more relevant the items, the closer they are to the root of the query; more generally, a search that returns a large number of very relevant items will have a tightly coiled spiral, whereas one with few relevant items will have a very loose spiral.

We intend to support the use of sound as part of a representation for a neighborhood. Sound is valuable because it can maintain the sense of being in a particular place, even when the place is too big or complex to be shown all at once. Server administrators should be able to define a digitized sound for sound-savvy clients to play while the user is within the directory. Sound can play a variety of other roles. It may be used to reflect activity of other users in the same or nearby neighborhoods [Gaver, 1991] [Cohen, 1993]. Variations in its timbre could be used to give hints as to the size of the neighborhood [Gaver, 1989]. In the physical world, sounds also play an important role in supporting the feeling of a sense of place [Southworth, 1969].

If usage information is available from the server, footprints (or some sort of dirt on the ground) can be used to show which of the items in the neighborhood are the most popular. This is like the worn marks on subway platforms in New York city. You can predict where the doors of the subway train will be when the train stops by looking at the worn spots on the platform. The same sort of cue can be used in a neighborhood to show users who want to follow the beaten path, where that path lies.

### 3.2 The 3D Icons

In our design 3D icons can vary along three dimensions: form, surface characteristics (color and texture), and information about the icon’s content (name and proxy). The basic form of a 3D icon is an approximately rectangular box that represents the type of the object. Box-like icons are keep the scene rendering requirements to a minimum by minimizing the number of polygons and simplifying hidden surface removal. Box-like objects also provide maximum space for mapping textures, drawing items’ names and presenting other information.

The form of the icon indicates the type of object it represents: the constraints on the icons’ forms are that the icon’s type should be recognizable from any direction (and ideally from a distance), that most icons have a large, flat surface area on which its proxy may be displayed, and that the form be relatively simple so that large directories render quickly on low-end machines. The objects represented by the icons are mostly analogous to types of icons in Gopherspace today (documents, search engines, interactive sessions, and gateways to other directories). Two new types are kiosk and person icons: kiosks have been described in the preceding section; person icons are place holders

for icons which will represent users when the environment is able to support synchronous social interaction.

Surface characteristics of icons include color and texture. Our intent is to use color as redundant information, to indicate the type of the icon. This will allow icon types to be distinguished in overviews. Texture is a variable that server administrators will be able to customize; they will be able to define the texture as a bitmap that will be painted onto the surface of the icon, and onto which other information (e.g. proxies) will be mapped. The face of the 3D icon is divided into areas used for the name of the item and the proxy. The title of the item is written across the top. Below the title is where the proxy is displayed. The proxy reflects something of the content of the object represented for the icon: for a picture, it might contain a thumbnail sketch; for a text document it might contain key words; for a new neighborhood, it might contain an indication of the neighborhood's size. Researchers have suggested a number of possibilities for proxies which would be interesting to explore in this context [Houde, 1993] [Wroblewski, 1994].

As the user approaches a 3D icon, the amount of information displayed on the icon changes since there is more screen space for display and since the user is presumably more interested in the item. From a long distance, only the general outline and color of the icon is readily discernible. At middle distance the name of the icon is visible. At close range proxies for the information within the icon become visible; as the user approaches even closer to an icon, more information (such as textures or sound) might become available.

### 3.3 Representing Overviews of Gopherspace

The purpose of an overview is to give users a glimpse of the larger context in which they are situated. In the current design, all overviews take the form of 2D maps, which can be brought up in a window separate from that containing the 3D view. There are three sorts of overviews that seem likely to be of use: an overview of the neighborhood; an overview of the local region of Gopherspace; and an overview of all of Gopherspace. In an overview of the neighborhood the collection of items will either be the current directory or the results of the most recent search. The overviews are 2D projections of the neighborhood as seen from overhead: that is, either the 'Stonehenge' arrangement or a spiral of icons. Colors will enable users to distinguish the type of the icon (since, in a 2D projection this will probably not be readable from the icon's form).

We haven't decided precisely what an overview of a region should be. Our working definition is that a region is comprised of the current neighborhood, neighborhoods directly connected to that neighborhood, and (perhaps) neighborhoods connected to those. An alternate possibility is to define the region as all the directories on the server the user is currently accessing. These possibilities differ because in many cases a gopher item may be connected to a neighborhood on a different server. The first definition seems more reasonable, but it requires on-the-fly querying other servers to generate a region's overview, and in the case where there are many inter-server connections, may take too long. The alternate possibility is much more efficient (and can even be pre-computed), but the results may contradict the user's perception of the organization of the Gopher hierarchy.

Finally, an overview of all (or at least large portions) of Gopherspace would be useful. However there are practical problems. It is difficult to identify the contents of Gopherspace because Gopher is a distributed system without a compulsory server registration. One initial step would be to have the client construct overviews of the areas of Gopherspace that the user has visited. While this will be useful only for seeing where one has been, and not for aiding exploration of new areas of Gopherspace, it would provide place holders in both the user interface for more comprehensive overview information that may eventually become available.

### 3.4 Interaction in Gopherspace

Most people are familiar with navigating 3D spaces since this is something they do in the real world. On the other hand, most people have little experience flying through space, so by limiting the number of options the user has for flying into limbo, we can make the user experience similar to some of the better arcade style games. By limiting the user's navigational controls to forward, back, turn left, and turn right, we can make it easy to learn how to use the interface. By limiting the height of the 3D icons we can ensure that the icon's proxy will usually fit within the user's field of view, and thus we needn't worry about providing ways to look up or down.

It is necessary to be able to select and open items and to start, stop, and change the direction of movement in Gopherspace. One possibility is to ignore selection, and to allow items to be opened by driving into them [O'Sullivan, 1994]. Alternatively, we can allow users to navigate around the space, and then select items and open them. In this case, either screen controls will be needed for one or both options, or use of a command key with the mouse will be necessary to disambiguate between the steering and opening modes. Onscreen controls seem to be the

better option, both in terms of ease of learning and use, and because they could support a wider range of low level operations. For example, a single click might be used to choose an option to scan the neighborhood, taking the user on a circular path around the inner periphery of a neighborhood — if clicks, and click-and-holds were used to steer, this could quickly become physically tiring).

In most cases, opening an item would result in a new window being displayed on the screen. If the item is a document, the window would of course contain the contents of the document. If the item was an interactive session or search engine, it would contain whatever prompts are generated by session host or search engine. If the item is the kiosk icon in the center of a neighborhood, it returns the user to the previous neighborhood. If the item is a gateway to a neighborhood, opening the gateway results in a visual transition to the new neighborhood. One idea for a visual transition is to give the user a sensation of high-speed travel travel over an anonymous looking grid, culminating with a slow approach to the new neighborhood. It is important that the approach to the new neighborhood allow the users to see the general arrangement of the neighborhood so that they can get a sense of the number and type of items in it. To make this visible, the user's trajectory might be one of swooping down from above (like a plane landing). The user is then deposited near the central kiosk, positioned so that both a part of the kiosk and some of the items in the neighborhood are visible.

### 3.5 Software Architecture

The 3D user interface described here will be rendered and managed entirely by the client software. No changes to Gopher servers will be required, at least for the information environment phase of the implementation. Client software that can synthesize the spatial scene from current Gopher directory and optional item meta-information, allows us to maintain compatibility with all Gopher servers. A pleasant side effect of this approach is that server and network bandwidth demands are minimized since we do not require servers to render scenes and ship bitmaps of the scenes over the network. Backward compatibility issues are also addressed by using the Gopher+ protocol's item attributes to hold meta-information. Gopher+ item attributes provide an open-ended extensible way of associating arbitrary meta-information with items and directories, and methods of accepting information sent from the client, so the user interface we propose will not require a new protocol.

## 4 Conclusions

As this is a report of work in progress, we have no real conclusions, as yet.The next step is to deploy a working prototype on the PowerPC Macintosh. As of this writing, a crude prototype, complete with rendering engine, is running. It is worth noting that even in these earliest stages of design, we have had to deal with issues that lie outside the realm of the disciplines traditionally involved in interface design. What forms should 3D icons have so that they are both simple and yet recognizable from many directions? What types of spatial layouts can help people remain oriented in a large space? What sort of layouts are legible both from the inside and from the outside? What icon and layout geometries are most invariant over scales, allowing them to be recognized from very close, and very far away? How rapidly should a transition into a new scene occur, so that the user can view enough to be oriented, but avoid boredom? How can we simultaneously support the sorts of regional and individual variation of appearance that makes real places rich and inviting, without sacrificing the core of consistency that that allows people to stay oriented and comfortable in real spaces?

# 5 References

[Alberti 1992]. Alberti, B., Anklesaria, F., Lindner, P., McCahill, M., Torrey, D.: "The Internet Gopher Protocol: a distributed document search and retrieval protocol."; University of Minnesota. anonymous ftp from boombox.micro.umn.edu <URL:ftp://boombox.micro.umn.edu/pub/gopher/gopher_protocol/>

[Cohen, 1993]. Cohen, J. (1993) Monitoring Background Activities. *Proceedings of the First International Conference on Auditory Display*. Santa Fe, NM.

[Erickson, 1993]. Erickson, Thomas. (1993) "From Interface to Interplace: The Spatial Environment as a Medium for Interaction." Proceedings of the European Conference on Spatial Information Theory. Heidelberg: Springer-Verlag.

[Gaver, 1989]. Gaver, W. W. (1989) The Sonic Finder: An Interface that Uses Auditory Icons. *Journal of Human-Computer Interaction*, 4:1. Lawrence Erlbaum Associates.

[Gaver, 1991]. Gaver, W. W., O'Shea, T. & Smith, R. B. (1991) Effective Sounds in Complex Systems: The ARKola Simulation. In *Human Factors in Computing Systems: the Proceedings of CHI '91*. Austin,TX: ACM Press.

[Hill, 1991]. Hill, W. C., Hollan, J. D., Wroblewski, D. A., & McCandless, T. P. (1991) "Edit Wear and Read Wear." In *Human Factors in Computing Systems: the Proceedings of CHI '91*. Austin,TX: ACM Press.

[Houde, 1993]. Houde ,S., Salomon ,G. (1993 ) Working Towards Rich & Flexible File Representations, in*Adjunct Proceedings of INTERCHI'93*, Amsterdam, The Netherlands. April 24-29.

[O'Sullivan, 1994]. O'Sullivan, D. (1994). Yorb: An Electronic Neighborhood. Personal communication and unpublished videotape.

[Southworth, 1969]. Southworth, M. (1969) The Sonic Environment of Cities. *Environment and Behavior,* June.

[Wroblewski , 1994]. Wroblewski, D. A., McCandless, T. P. & Hill, W. C. (1994). Advertisements, Proxies, and Wear: Three Methods for Feedback in Interactive Systems. Natural Dialogue and Interactive Student Modeling[tentative title]. Heidelberg: Springer-Verlag, in preparation.