XIAOLONG (LUKE) ZHANG, GEORGE W. FURNAS

18. SPACE, SPATIALITY AND TECHNOLOGIES

Multiscale Space and Place: Supporting User Interactions with Large Structures in Virtual Environments

INTRODUCTION

This chapter is about the exploration of a new design paradigm for virtual space and virtual place. Spatial metaphors have been widely used in support of information access and social interactions. Example designs include Media spaces (Bly et al., 1993) and spatial video conferencing (DeSilve et al., 1995; Sellen & Buxton, 1992). In the design of collaborative virtual environments (CVEs), spatial metaphors are embedded in 3D environments to create a shared space for people to work together. Shared space integrates data and users together, and provides an "explicit" and "persistent" spatial context for collaboration (Benford et al., 1996). The sense of space can facilitate social interactions by allowing the partition of available space and permitting users to apply their spatial social skills, and so on (Benford & Greenhalgh, 1997).

Usually, our interests of space in virtual environments lead toward its implications for social interactions. Our understanding, interpretation, and use of virtual space are still largely based on our experience in real space. For example, CVEs are often designed to support activities we see in the real world: meeting (Greenhalgh & Benford, 1995), social gathering (Lea et al., 1997; Waters et al., 1997), product design (Linebarger & Kessler, 2002), etc. While there is nothing wrong to replicate or simulate our real-world experiences in virtual worlds, it is also needed to consider how the unique characteristics of virtual space can benefit our work. This chapter will examine multiscale virtual environments (mVEs), a new design paradigm that exploits the virtuality of virtual environments. In mVEs, a user can choose to work at different interaction scales: being a giant to see the big picture of a structure and manipulate large objects, or being an ant to examine the details of a particular part of the structure and work on small objects. The user's perception and action in mVEs could be very different from what they have in conventional virtual environments. We will study multiscale space and multiscale place separately by emphasizing their different roles in supporting users' actions. Multiscale space refers to the unique spatial aspects of mVEs, and our interests of it lie in its support for user interactions with virtual objects and virtual environments. The notion of

P. Turner and E. Davenport (Eds.), Spaces, Spatiality and Technolog, 261—20.

multiscale place emphasizes the social aspects of mVEs, and concerns the impacts of multiscale space on interactions among users.

This new design paradigm could have a great potential to address a challenge we are facing in dealing with structures that cross different scale levels. We are living in a world where things have very different sizes and, as Aristotle argued, remain naturally in their proper places. These different sized objects create different structures at different size-scale levels, and demonstrate important characteristics at different levels. Consequently, the physical world looks quite differently from scale to scale. The book "Powers of Ten" (Morrison & Morrison, 1982) vividly illustrates the images of the physical world from the *gigantic* galaxy level to the *minuscule* sub-atomic level. Our understandings of such a complex world did not come easily. It took us thousands of years to discover these structures at multiple scale levels and construct a notion of a hierarchical world. Even so, we still do not understand the world fully, in particular at those scale levels that are beyond our direct perception scale range. This might be related to our limited cognition and interaction capability.

Compared with the vast scale range of the world, the scale range of human beings' cognition and interaction capabilities is very limited. In the real world, our naturally born capabilities do not allow us to see atoms with our naked eyes and to reach a destination thousands of kilometres away easily just with our legs. The size scale of matter in this world spans from the level of 10^{-16}m, the size of the smallest elementary particles, to the level of 10^{26}m, the size of the observable universe (Morrison & Morrison, 1982). This range of forty-two orders of magnitude is far beyond the normal interaction scale range of human bodies, which is only about four orders of magnitude from millimetres (10^{-3}) to tens of meter (10^{1}). This mismatch in scale makes it difficult or even impossible for people to observe and manipulate many object structures directly in the real world. People have to rely on instruments to deal with objects at different scales: using microscopes to observe small structures like cells, and using satellites or space shuttles to examine large objects like the Earth. These instruments and people's naked eyes together build a multiple-scale (multiscale) system that helps to deploy our scarce scale resources across the needed scale ranges.

People face similar problems in virtual worlds, where objects like information structures and model structures are becoming increasingly large and across many different levels. The information world has grown and continues to grow rapidly. Billions of web pages are indexed and distributed on the Internet; the hierarchy of a file system on a personal computer may hold tens of thousands of documents; a corporate web site may connect millions of web pages. Model structures people build in virtual environments are also getting larger and larger. A protein structure may have thousands of DNA sequences. A virtual city can hold tens of thousands of buildings. Managing such huge information structures or model structures could be a daunting task for us. Given our limited cognition resources, we cannot observe all levels of a large structure simultaneously, and we have to focus our attentions on some parts of the structure in interactions. Either local content or global context information has to be hidden. Incomplete information may cause problems in the understanding and management of information structures. Obtaining both detailed contents and sufficient contexts could be a challenge.

One approach to address this issue is to use multiscale technology, which will be introduced soon. In this chapter, we discuss multiscale in 3D virtual environments. It begins with an introduction to multiscale in virtual environments. Then, it forwards to issues concerning user perception and action in multiscale space. Next, the article examines emerging issues related to social interactions in multiscale place. Finally, the implications of the notions of space and place for the understanding and design of virtual environments are discussed.

MULTISCALE TECHNOLOGY

Multiscale techniques have been seen in many research projects. The first exploration of multiscale technology was Pad (Perlin & Fox, 1993), and other systems include Pad++ (Bederson & Hollan, 1994) and Jazz (Bederson et al., 2000). These projects largely focus on 2D interfaces.

A 2D multiscale user interface provides users with a ubiquitous zooming capability which allows users to manipulate the rendered size of objects continuously and dynamically, so it is often called zoomable user interface (ZUI) (Perlin & Fox, 1993; Bederson & Hollan, 1994; Bederson et al., 2000). A 2D multiscale environment can be directly understood as one allowing multiple magnifications of display.

In such ZUIs, there is another aspect of scale which relates not to perception but to action: zooming also sets users' action domains. For example, manipulating objects through the ZUI window or panning that window is typically constant in the display units (e.g., panning one window width per sec, moving a visible object half a window width by direct manipulation of its image). This means that when users are really zoomed in, they are making very small movements in the virtual world, and when they are zoomed out, they are making huge ones – the display magnification turns into a corresponding magnification of action. We will see that this natural correspondence of perception and action in 2D is not so straightforward in 3D.

Our study explores the understandings and design of multiscale technology in 3D. In a 2D ZUI, displayed size and corresponding action scale are mediated by a simple magnification parameter. In 3D, however, one might argue that no such magnification parameter is needed – all we need to do is to adjust the viewing distance to control the image size of objects we see. If you want to see something small, move in close. If you want to see something big, step back.

Imagine the eye as a pinhole camera, and that what we see is the image on the projection plane of the camera. As seen in Figure 1, the projected size can be written as:

$$S_p = \frac{S_0}{D} \cdot d_c$$

where: S_p – the projected size
S_0 – the real size
D – the distance of the object to pinhole,
d_c – the distance of projection plane to pinhole

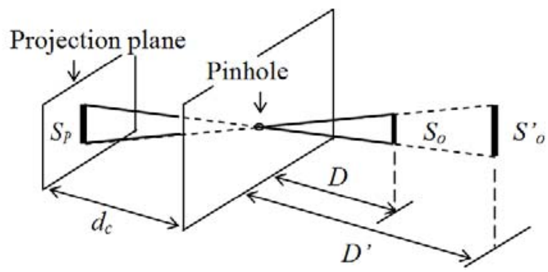

Figure 1. Pinhole Camera

Obviously, the image size, S_p, can be controlled by simply adjusting *D*. However small an object is, a distance can always be found to make a sizable image projection. For example, to see an object with a size of 1nm, users can set a viewpoint at a distance of 1nm to the object , and get the view angle of the object about 53 degrees, plenty large enough to see good details. It is not necessary to use magnification. In fact, one might argue that the idea of magnifying the world would be fruitless. If the centre of magnification is the pinhole, both the size of the object and its distance to the pinhole will increase proportionally, the projected size will be the same. As seen in the figure, S_0 and S'_0 have the same projected image. So why would there be any sensible notion of scale change for interaction in a 3D virtual environment?

The answer comes from the fact that interaction in a virtual environment involves more than providing a single, static pinhole camera view. There turns out to be many other size-dependent aspects of interaction that bring sense to the notion of a 3D multiscale interface to a 3D multiscale world.

More than just viewing distance, there are other perception parameters, as well as several action parameters that need to be adjusted. If a stereo display is available, for example, the eye separation should be at a scale similar to the viewing distance to generate a correct binocular disparity. If the eye separation and the viewing distance do not match well, say a 1nm viewing distance vs. a 1m eye separation, the left and right images would be totally irrelevant, and stereopsis would be impossible. Similarly, having an action domain comparable with the viewing distance is also critical for interaction. Moving with a step size of 1km under a viewing distance of 1nm would make it hard to keep a consistent view. Any movement will change the view significantly.

In this sense, to work on structures at different sizes, users need a whole coordinated set of interaction size parameters, including viewing distance, eye separation (in stereo views), eye-level (if there is a ground plane), navigation speed, reaching distance, and so on. mVEs, by integrating multiscale technologies into 3D virtual environments, give users the ability to control this set of parameters, which can be mediated by a single measure – *interaction scale,* when they work on objects and structures at various size-scale levels. To work on small objects like atomic lattices or a DNA double-helix, users can be something much smaller than an ant, let's say *a nano-ant* or *nanant*, moving and manipulating very precisely to maintain reasonable action accuracy. Similarly, interacting with a virtual model of the Milky Way Galaxy, the relevant distance scale would be at the level of light years, and users need to work like a super giant, a *giga-giant* or *gigant*, to interact properly in the virtual environment. Although transcending the spatial scale of the space in the real world is impossible, users can have such experience in mVEs.

Some research projects on 3D virtual environments have considered scale as an interaction parameter. The auto-scaling tools (Mine et al., 1997) gave a user an extensible arm-reaching distance to help reaching object at any distance. The Go-Go technique (Poupyrev et al., 1996) dynamically rescaled a user's arm to increase the reaching distance. The Image-Plane-Interaction (Pierce et al., 1999) allowed a user to manipulate the projection of objects at any distance. However, these tools only consider one particular interaction parameter – arm-reaching distance, and ignore the need to yoke other perception and action parameters together. Furthermore, these designs, as well as Pad++, focused on single user environments, and gave no consideration to the use of scaling in environments shared by multiple users. The design of mVEs, on the contrast, provides a user a space, where she can adjust her interaction scales in the work, *and* a place, where she and other users can combine their multiscale capabilities to deal with large structures.

SPACE AND PLACE IN MULTISCALE VIRTUAL ENVIRONMENTS

Philosophical arguments on space and place have had impacts on research in Human-Computer Interaction (HCI). Harrison and Dourish (1996) distinguished these two concepts by their different roles in people's life, and understood space and place distinctively by emphasizing the physical and social aspects respectively. This classification implies a way to understand mVEs: studying the physical properties of space in mVEs to understand how mVEs can support user interactions with objects and virtual environments, and examining the social properties of space to see how mVEs may affect user interactions with others.

In the real world, people usually put more interests in place than space. Architectural design is about to build places embedding social structures, rather than just to create physical structures in space. We can build things in space, but can hardly do anything on the physical properties of space itself. Our philosophical understandings on space have been consistently evolving: from the Greek notion of space as an abstract entity, to the conception of space as supremacy held by Newton and Descartes, to Kant's view of space as *a priori* and being subjective, to the phenomenological view of the critical involvement of human beings in the constructions of the notion of space (Casey, 1997). The advance of philosophical conceptualization on space may have affected our views on the relationships between human, objects, and space, but our capabilities to manipulate the physical characteristics of space have not improved too much. Thus, in real-world design, all we can do is to build place based on our improved understandings of the relationship between human and space, and to embed social implications in space.

Such a socialized view of space has been reflected in user interface design. In HCI research, new designs often focus on place, rather than space (Fitzpatrick et al., 1996; Gaver, 1992; Bly et al., 1993). This is not a surprise, given the nature of HCI research, which emphasizes human and social issues in the use of technology. While such a social perspective is important, we may also need to think about how to leverage the spatiality of virtual space to enhance user interactions. Virtual space is artificially made, and how we can interact with it is all up to design. New interaction experiences can indeed be created in virtual space to support people's work. Some research projects have explored this direction by going "beyond being real" (Hollan & Stornetta, 1992). For example, distance in virtual environments can be distorted with fisheye views (Furnas, 1986) so that objects of interest can get more space than those of no interest (Robertson & Mackinlay, 1993; Raab & Ruger, 1996); the available space in 2D workspace could be expandable by using zooming tools (Perlin & Fox, 1993; Bederson & Hollan, 1994, Bederson et al., 2000); navigation in 3D could be simplified by using teleportation, which makes spatial separation in virtual environments seemingly diminish. These creative designs not only help users to better work with complicated data and virtual worlds, but also open the doors to new ways for social interactions.

The design of mVEs is such an attempt that overcomes the constraint on spatial scale and allows users to manipulate the scale of virtual environments. mVEs offer a new type of virtual space in which users can observe important characteristics of structures at different scales and act on objects with different sizes easily. Their perception and action in mVEs could be very different from what they have in conventional virtual environments. At the same time, mVEs also provide users with social places different from what they have in real life and in conventional virtual environments. In multiscale places, users will be able to cooperate from different scales and leverage their different capabilities. Hollan and Stornetta (1992) argued that how people communicate with each other is related to what media they use. The introduction of multiscale changes the characteristics of virtual space, which mediates most of collaboration activities in mVEs, and would affect the ways on which users rely in their interactions with others. Thus, our analysis will be at both the individual and collaboration levels. At the individual level, we focus on user interactions with objects and multiscale space in virtual environments. At the collaboration level, we are interested in social interaction among users in multiscale place.

MULTISCALE SPACE

Our interests of multiscale space concern how multiscale may affect users' understandings of and their actions in virtual space. Multiscale technology can visualize the same structure with different representations at different scales and affect users' visual perception and related action. The focus here will be on the impacts of scale on visual perception and action, in particular the control of scale, spatial perception in multiscale space, and visually-guided activities. A scale-based information visualization technique will also be briefly introduced.

Integration of Interaction Scale Parameters

The first problem in multiscale space is how to give a user sensible control over the scale, or the suite of various interaction parameters. This has not been a serious problem in traditional virtual environments, where users usually stay at one interaction scale level. In these virtual environments, although there are some circumstances where users need to change a particular interaction parameter, such as quick navigation speed or far reaching distance, situations that require users to alter a whole set of perception and action parameters are rarely seen. In mVEs, however, when users choose to work at different scales, this issue cannot be ignored.

Our approach is to use avatars to integrate interaction parameters together. An avatar is a graphical representation of a user in a virtual environment. In collaborative environments, it provides such information as presence, location, identification and the like about the represented user (Benford et al., 1995), and is widely used in CVE design. An avatar is the virtual body of a user in virtual worlds. In real life, the human body ties our perception and action together, and perception and action parameters are at a scale level similar to the body size – the meter. Our step size is about 1m; arm-reaching distance is at the meter level; the shortest focus distance of our eyes is about 0.1m. The human body can be seen as a measure to define our interaction scale in the physical world. In mVEs, an avatar works as a virtual human body to define the user's perception and action.

In 3D multiscale environments, the avatar metaphor can provide a natural way to tie together the suite of size-based parameters needed for interaction: the size of an avatar is linked to the interaction scale the represented user chooses. By changing this scale, users' avatars will be resized accordingly. Figure 2 shows a user's avatar at two scales, which is sized differently. The left image indicates a larger interaction scale of the user, who has a broader overview range, a faster locomotion speed, and a greater manipulation domain. In comparison, the right image shows that the user, who is at a smaller scale, gets a narrower overview range, a slower speed, and a closer manipulation domain. The avatar metaphor yokes a suite of parameters together by tying interaction scale with the avatar body size.

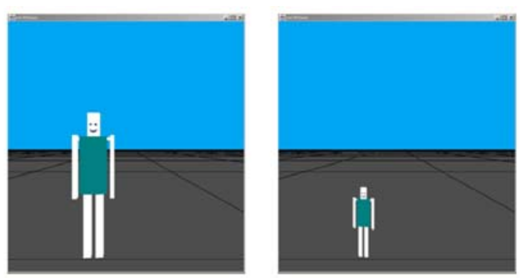

Figure 2. A User's Avatar at Two Different Scales

Interaction scale actually measures the relationship between size parameters of the users and those of the environment. This relationship is relative. Mathematically, to work with a small object, we can either shrink a user's parameters (including observing distance) or grow the world's parameters (maintaining observing distance). Actually these two models would make scant difference for individual users, but when several users work together at different scales, having multiple users independently changing the scale of the world would usually not be desirable. Typically, independent actors would want to change their relation to the world independently, so the implementation of scale control should support the metaphor of each actor changing their own interaction scale parameters. This decision also supports the ant-giant metaphor quite naturally.

The explicit change of interaction scale in 3D relative to the virtual environment will be called generically re-scaling. Scaling-oneself-up, or just scaling-up refers to an increase in the size of a user relative to the virtual environment, and consequently the user's perception and action parameters relative to the virtual environment are increased. Scaling-down decreases the size of the user with respect to the virtual environment, and leads to smaller perception and action parameters.

We notice that other sorts of scaling tools have been used in some research projects. As mentioned previously, the auto-scaling tools (Mine et al., 1997), the Go-Go technique (Poupyrev et al., 1996), and the Image-Plane-Interaction (Pierce et al., 1999) supported scalable arm-reaching distance. However, in order to facilitate the user's interaction with large structures, changing only one specific interaction parameter may not suffice. Other projects, such as the World-In-Miniature(WIM) technique (Stockley et al., 1995) and the CALVIN system (Leight et al., 1996), allowed users to see the virtual world and act on it from different scales. They usually provided users with another view of a scaled-down world at a given scale level, and the dynamic range of scaling is very limited. Also, in these projects, users usually cannot control the scaling process and manipulate interactive parameters. In (Mine et al., 1997), scaling was even designed by choosing a special scaling centre to fool the user's eyes and make the scaling change unnoticeable. Users are passive perceivers of scaling, rather than active controllers.

In mVEs, interaction scale is directly manipulated by users, just as are their other spatial parameters, such as location and orientation. This is analogous to the case of 2D ZUIs (Perlin & Fox, 1993; Bederson & Hollan, 1994; Bederson et al., 2000), where interaction scale is treated as a first class spatial parameter, and users can choose any scale by zooming, as directly as they control panning. A working hypothesis of this research is that a fully controllable scaling tool can help users to understand an mVE better by making users active multiscale explorers in their multiscale world.

Changing a user's working scale has several consequences. Here, we focus on the impacts of scaling on spatial perception and such activities as navigation and object manipulation.

Impacts of Multiscale on Spatial Perception

A key issue in re-scaling the world is the choice of scaling centre. When the centre of projection, usually the location where the virtual eye is located in 3D graphics, is the scaling centre, the virtual sizes of objects get larger, but they get correspondingly further away and hence look the same. The user would not experience any change of the views, similar to the phenomenon illustrated in Figure 1. However, if the user is binocular, re-scaling around one eye will lead to movement of, and hence changes to the image in, the other eye. This will change convergence angles on objects and change binocular disparities, affecting the user's depth perception.

Figure 3 presents the views of a three-box configuration viewed at two different binocular scales. For this figure, the eyes' convergence angle has been kept constant, while the eye separation has been increased. One result is that the eyes now converge on a different point. Conversely, to maintain convergence on the middle box, the eyes would have to be substantially crossed. Since humans cannot converge their eyes more than 20-40 degrees comfortably, users cannot look binocularly at objects closer than a few eye-separations away. Examining small details requires getting closer; doing so binocularly requires shrinking eye separations.

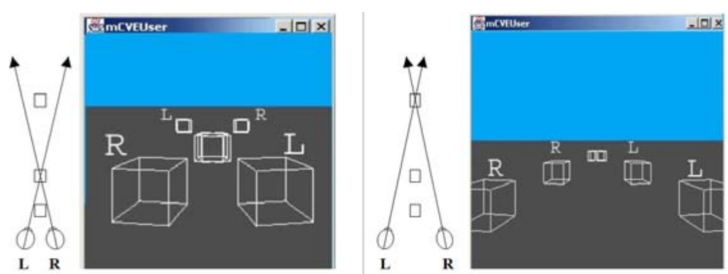

(a) *(b)*

Figure 3. Different Convergences in Scaling: as a user scales up from (a) to (b), the converged object changes from the middle box to the farthest one

The brain uses the differences in relative positions of objects as seen by the two, differently positioned, eyes to compute depth and construct spatial models of a scene. In real life, disparity cues are effective up to 10m (Cutting & Vishton, 1995), which is about 150 eye separations away. Within that range, people make very subtle discrimination, about 12 arc seconds, which translate the depth difference by 1mm at a distance of 1m and 9cm at 10m. Clearly, if users must engage in subtle depth manipulations at some particular scales, they want to bring interaction scale with the corresponding eye separation into that range. For example, if users need to reposition the middle box right next to the farthest one, they can change their interaction scales so that the views converge on the latter. Then, the depth of the former can be adjusted until it is also converged. Examining large structure requires getting farther from them; doing so with good stereoscopic depth perception requires increasing the eye separation.

For small rescalings, choosing the scaling centre at the midpoint between the eyes can minimize disruptive effects without sacrificing the depth cues (Ware et al., 1998). However, if rescalings are large (e.g., orders of magnitude), the eye changes are still dramatic.

There are also other important candidates for the scaling centre. Observing virtual cities or landscapes, a viewer often has virtual feet planted on a grounding plane, making navigation much easier. In such cases, choosing the standing point as the scaling centre will raise and lower eyes, giving the viewer different overview ranges. Figure 4 shows a user's views of a simple landscape, from the same planar location, but at different eye-heights resulting from changing interaction scale. Occlusion patterns change, exposing potentially different explicit information, and altering depth cues affecting the user's space perception.

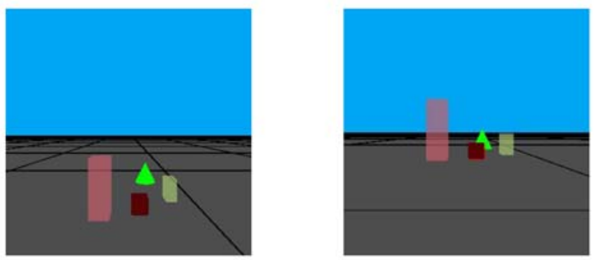

Figure 4. Effects of Eye Level on Spatial Perception

With scaling, users can easily get a better macro-level understanding of the structure of the space by increasing their eye-height level. This may help them to improve navigation performances because of easy access to spatial knowledge (Darken, 1996). Note that this mVE overview technique is still egocentric, and differs from the exocentric one provided by WIM (Stoakley et al., 1995), in which users can see and manipulate their own avatars.

Impacts of Scale on Actions

In the avatar metaphor, intuitively, a user's action should also be proportional to the body size of her avatar. A gigant should have a gigant step and a gigant reach, and a nanant should move with a tiny step and have a much small reach.

Locomotion Speed

It makes sense to have the navigation speed dependent on scale. A nanant with a gigant step would easily get the nanant lost. A gigant with a nanant step would also be a problem, because it would demand tremendous effort to make the view change.

Scale-related speed can help to address a locomotion problem in large space, where a quick speed and an accurate movement near the target need to be balanced. Solutions to this problem include a logarithmic motion function (Mackinlay et al., 1990). In mVEs, this issue can be tackled by the combination of moving and scaling (Furnas & Bederson, 1995), which allows the user to choose a larger speed to approach the target quickly and then to switch to a smaller scale and a slower speed for accurate movement. There is no need to specify the target explicitly as required by logarithmic motions. However, scaling costs time, and could slow down locomotion. Also, users have to be more involved in the control of locomotion.

Scale-dependent speed could have an impact on perception. Motion parallax cues are from moving. A proper motion parallax cue which is consistent with other visual cues, such as elevation and stereopsis, can help users to understand the virtual space better. However, a very fast motion may hurt animation quality, increasing cognitive load in interaction. In addition, scaling will change a user's view of the world, so maintaining a consistent understanding of the locomotion direction at different scales could be a challenge for the user.

Manipulation

Manipulation could also be designed to be scale-dependent, giving a gigant user a larger action domain than a nanant user, just as a taller person can reach farther than a shorter one in real life. However, mimicking reality can miss special opportunities of virtual environments. While some actions may need to be tied with scale, other actions may benefit from the independence of scale. As previously mentioned, the auto-scaling tool (Mine et al., 1997), the Go-Go technique (Poupyrev et al., 1996), and Image-Plane-Interaction (Pierce et al., 1997) all ignore the constraint of the distance, and make it easy for the user to grab any object in sight without considering distance.

While these tools allow the user to select and move any visible object, handling a small and distant object with a small image is still difficult for them. In mVEs, the combination of scale-dependent navigation and scale-independent selection can give some help. Knowing where a distant object is, users can first choose a fast moving speed to approach it quickly. As soon as it is becomes visible, they can select it. Thus, multiscale tools improve the dynamic range of selection (Guiard, 1999).

Scale-Based "Semantic" Representation

We mentioned earlier that our limited visual perception does not allow us to see very big or small objects. Therefore, we have to use various tools to change the sizes of objects in observation. In virtual worlds, we might aspire to do more. That is, in order to be maximally semantically informative, the representation of a structure might change in a non-geometric way with the change of interaction scale. In 2D ZUIs this technique is called "semantic zooming" (Perlin & Fox, 1993; Bederson & Hollan, 1994) and has been used successfully to provide users with scale-based context-sensitive information. In mVEs, a similar tool called scale-based "semantic" representations is provided to help users observe important characteristics of a structure at different scale levels. For example, working on a model structure of a new type of materials, users can see its molecular structures at one scale, and atomic structures at another. With the dynamic and continuous control of interaction scale, users can also know how these two models are related to each other. Applying such techniques in modelling a virtual city, users will easily see the images and features of the city at the levels of city, district, neighbourhood, and street. The smooth transition between these images can help users better understand how issues at different scales relate. The detailed design and implementations of this scale-based "semantic" representation can be found in (Zhang 2004).

So far, we have discussed several aspects of multiscale space, including what interaction scale means to the user, what impacts it may have on user perception and action, and how information is visualized across different scale levels. All these issues concern interactions between individual users and multiscale space. When users work together, multiscale space becomes multiscale places, and concerned issues are turned to interactions among users.

MULTISCALE PLACE

When users gather in a multiscale place and collaborate synchronously, they can choose different interaction scales. This raises many cross-scale collaboration issues, such as spatial relationship, action coordination, awareness, and information sharing. In this section, our focuses will be on cross-scale spatial relationship and cross-scale actions. Cross-scale awareness and cross-scale information sharing issues, which have been discussed in great details in our other papers (Zhang & Furnas, 2002; Zhang, 2004), will be introduced very briefly.

Spatial Relationship

When users choose different scale in a multiscale place, their avatars will be sized differently, and so occupy different volumes of space and have different influence scopes. Consequently, users' understandings of the same space may vary from scale to scale. Their space-mediated interactions could be affected by the discrepancy in the perception of space.

Using space to mediate social interactions is common in real life and CVEs. Under the spatial model proposed by Benford et al. (1994), interactions are mediated by aura, nimbus, and focus, which are all sub-spaces around a user's avatar and move with it. Aura delimits the enabled interaction space of a user. The collision of two auras will trigger social interactions between users. Focus and nimbus control the level of awareness between users. Focus determines the attention range. Nimbus sets the boundaries of the presence of a user, and decides the availability of the user to others.

In traditional CVEs, a user's aura, focus, and nimbus are usually fixed in shape and size. In multiscale places, however, the size of these functional spaces might logically be proportional to the user's avatar size so that her awareness ranges change when she becomes larger or smaller. As the user works as a giant, she will have a larger aura, focus, and nimbus that enable her to interact with more distant users. Similarly, as the user works as an ant, the scope for her interaction and awareness is shrunk so that she can concentrate her work with in a relatively small space. As seen in Figure 5, sphere-shaped auras of two users at two different scales have different diameters, indicating their scale-dependent awareness scopes.

Figure 5. Scale Dependent Auras

Certainly, it is possible to give the user a fixed-size aura, focus, and nimbus, making these subspaces consistent and scale-independent. If so, no matter at what scale she is working, the user will have the same interaction space and awareness range. It could be useful in situations where the user shrinks herself but still hopes to maintain a larger interaction space. However, it could be a problem when the user scales up. Keeping a consistent interaction space will give the user an interaction space smaller than her avatar body size. This means she can only interact with and be aware of those users inside her body, and only those users inside her body will be able to be aware of her. The interaction domain of the user is severely limited. With scalable social interaction spaces, the user will not have such problems. Scalable interaction spaces can help to adapt her actions to her observation scale so that her social interactions (or action domains) can be consistent with the range of perception.

Scalable interaction and awareness ranges could affect users' interactions both positively and negatively. On the positive side, users can modify their interaction distances by simply changing interaction scale. To interact with distant users, a user can increase the aura size by scaling up; to reduce her availability, the user can simply scale down to shrink the nimbus. However, due to the size change of these interaction spaces, a distance which can trigger social interactions at one scale may not do anything at all at another. This may confuse users. Also, when these spaces are re-scaled, users may find some unexpected events. For example, when a user scales up, her aura is increased and touches other users who were outside the interaction ranges previously, triggering new interaction events. This kind of interactions may not be wanted. Similarly, when the user scales down, she may lose the focus of those users who were previously focused. Another problem of scalable aura, focus, and nimbus is that the awareness ranges of users will become asymmetric. In traditional CVEs, users are usually at the same scale level, and their interaction ranges are similar, giving them symmetric awareness ranges. For example, when a user's avatar falls inside another user's focus, it is very likely that the focus of the former can also catch the avatar of the latter. In multiscale places, different sized focuses will give users different attention ranges. Separated by the same distance, a user with a smaller focus may not be aware of other users, whose larger focuses provide full awareness of her. This asymmetric awareness may seriously affect the collaboration if the user does not realize the existence of this problem. Such scale-related spatial issues are not seen in real life and conventional virtual environments, and may distract users from their primary tasks or even interrupt their work. Therefore, these issues have to be addressed in future research.

Cross-Scale Manipulation on Object

Users need to coordinate their activities on objects at different scale. We considered two issues here: how a user's activities may affect others, and how they can work on the same object at different scales. For the first issue, it becomes important when two users are working on the same object at the same time. The deletion of a structure by a user may remove some objects that are important to the other. To avoid this problem, we adopted a design option to use a locking mechanism to check whether two different people's activities may interfere. Although the locking mechanism is not new in collaboration systems (Singhal & Zyda, 1999), what makes this design differ from others is the consideration of the scale factor in locking. When a user begins to work on an object at a particular scale level, the object and related objects at other levels would be locked. For example, when a user is moving an object inside a hierarchy, all its parent objects will have a locking flag, preventing them from being deleted. It should be noticed that what objects should be locked is highly task-dependent. While the above example requires locking all nodes above the interesting objects, tasks like moving a node and all its children to a distant place may need to lock all objects below the interesting node. Therefore, the locking tool allows users to specify which nodes to lock in their work.

For the issues concerning the synchronous collaboration on the same object at different scale, one challenge is to get the object with an appropriate size for users' actions. If two users are at different scales, a target object that is at a right size to one user could be huge or tiny to the other. Without a proper size of the target object, manipulating the object collaboratively could be a problem. One way to address this issue is to provide shared subjective views of the target object to different users by scaling the object temporarily. If it is too small for a user, a scaled-up model can be given. If an object is too large, it can be scaled down. As seen in Figure 6-a, the object of interest, a box, is too small in the view. As the box is selected, it is automatically scaled up (Figure 6-b). Similarly, if an object is too large, it can be scaled down.

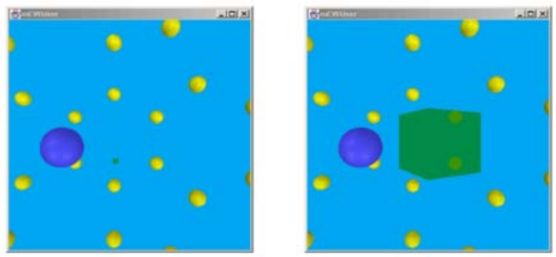

(a) *(b)*

Figure 6. Auto-Scaling in Cross-Scale Object Manipulation

Another design option we provided is to render the shared object with fixed-size representations so that its size does not change with the scale. Users will always be able to see it with the same appearance at different scales. The drawback of these approaches is that if the size of the object somehow becomes important in collaboration (e.g., moving it into a box), having a view with a distorted size could be a problem. To address this, a toggle tool is provided so that users can change the object back to its actual size. Such automatic re-scaling techniques can also be applied in other collaboration actions, such as passing an object across scales.

Cross-Scale Avatar Manipulation

When each participant is aware of others' working contexts at different scales, they can combine their knowledge from different scales and their different action capabilities to complete some difficult or time-consuming work easily and quickly. To achieve this goal, users should be allowed to be involved in others' work directly. One design choice is to make a user's avatar controllable by others. Imagine a situation where an ant-like user needs to take an object to a very distant place. Certainly, she can change her scale back and forth to get this work done. If she could get help from a giant-like partner, they together can do the work in an innovative way: the giant can simply move the ant to the destination approximately, and then the ant fine-tunes her position by herself. Such collaboration will be valuable to actions that require both fast speed and high precision.

The implications of such a direct manipulation of avatars for HCI deserve more attentions. In CVEs, a user's avatar can provide important awareness information, but usually cannot be controlled by others. That is to say, an avatar in general only *affords* providing awareness information in social interactions in CVEs. In other words, the affordance of an avatar object in virtual environments is to support obtaining awareness information.

Affordance, coined by J. J. Gibson (1979) and extended by Norman (1990), refers to perceivable physical properties of objects that suggest activities. Affordances provide people with information about what can be done. A good affordance design should help users to know what objects can be acted on and how to act on them (Norman, 1990).

Some properties of an artefact can also afford social interactions among users. Bradner (2001) called such properties social affordances. Although the social affordance of an avatar primarily lies in its support of collaboration awareness, an avatar, as an object in a virtual environment, can also afford manipulation activities, just like other non-avatar objects. The reason that an avatar is not usually used in such a way might be due to the legacy that people inherit from real life, where it is not often seen that adult people have their bodies moved by others (except in some extreme cases, such as public protests, where people are moved by the police).

Designs to support social interactions in multiscale places do not have to follow real-life experiences. Multiscale virtual worlds are not replications of the real world. Instead, multiscale, as a new technology, provides users with new experiences, and should not be interpreted with old metaphors. Allowing the direct manipulation of avatars by others is an effort to break down the metaphor of “reality”, and to help people work on things they would not be able to do in real life. Of course, this direct interference with other users should be regulated. The user should be aware of the action of her avatar being moved by others and have a control over whether this action can be done. Therefore, before any avatar manipulation activity happens, the avatar owner is warned, and informed of who is going to move her avatar. Her responses determine whether such coordinated activity will go on.

Cross-Scale Awareness and Cross-Scale Information Sharing

In collaboration, a user needs to be aware of other users and their activities. In CVEs, avatars carry awareness information. In multiscale places, however, scalable avatars post new challenges for awareness. A nanant user could be too small to be visible, and a gigant user could be too big to be fully comprehensible. Social interactions could be hurt because of the lack or incompletion of awareness information.

We address this issue by setting a maximum and a minimum limit for avatar size. When an avatar reaches a size limit, it is rendered with a nominal size, and stops growing or shrinking. At the same time, a visual cue appears above the avatar, indicating the actual avatar size is different from what is presented. Toggle tools are provided to allow users to check the actual size of an avatar.

Cross-scale context sharing is another issue for cross-scale collaboration. To collaborate, users need to know their working contexts. In conventional CVEs, sharing information is relative easy, because users are usually interacting with the same set of objects. In multiscale places, however, users may work on different structures and objects at different scales. Their individual views may be very different, and context sharing could be difficult. Imagine a scenario in which two users are working on the same model structure, but one at the molecular level and the other at the atomic level. Seeing different things, how could they share information across scales?

We address this issue by animating the connections between two users' seemingly irrelevant views. The animation is created by interpolating users' view positions, view orientations, and interaction scales. The animated dynamic view transition could help users understand how their views and their work are related to each other.

Cross-scale awareness and cross-scale information sharing have been discussed in our other papers. For detailed theoretical analyses and technical implementations, please see (Zhang & Furnas, 2002) and (Zhang, 2004) respectively.

Our understandings and analysis of multiscale space and multiscale place led to the design and implementation of a Java-based prototype system that support multiscale interactions (Zhang, 2004). Several experiments have been conducted in the prototype system. It has been found that in multiscale places, collaboration between two users allows them to better accomplish search-based tasks in a vast area (Zhang & Furnas, 2003). Another preliminary result shows that multiscale can improve user performances in navigation. A negative consequence of multiscale technology has also been discovered: social distance negotiation in multiscale places could be a big challenge for users at different scales (Zhang & Furnas, 2002).

CONCLUSION

In this chapter, we have presented a novel interactive environment, mVE, and discussed issues concerning multiscale space and multiscale place. We have examined the meanings of scale and multiscale for user interaction in terms of size parameters of perception and action. mVEs give users the freedom to manipulate a suite interaction parameters in their work, determining the spatial relationship between users and virtual space. In mVEs, user collaboration across different interaction scales raises many new issues. Our discussions have been focused on the impact of multiscale on spatial relationship among users and cross-scale manipulation.

In our research, virtual space and virtual place issues have been studied separately. For multiscale space, we have focused on the physical properties of virtual environments, and investigated how the unique spatial features may benefit people's work. Our interests of multiscale place are in the implications of multiscale technology for space-mediated social interactions. Treating space and place separately can help us understand what mVEs can do best to support collaboration. Clearly, 3D space gives virtual environments some edges over 2D space in visualizing and managing complicated objects and data, and multiscale space further enhances this advantage. However, it should be remembered that 3D places may not offer users more benefits than other virtual places like chatting rooms can do in support of collaboration activities that focus on message exchange, because users may have to deal with other issues, such as 3D navigation, not directly related to communications. Introducing multiscale into virtual environments may make this matter even worse, because users have to pay more attentions to scale control, object representations, and spatial relationship in mVEs. This may actually impede people's work if the primary goal of collaboration is to exchange information rather than to manage complicated structures.

Dix has discussed the distinctions between communication-oriented and artefact-oriented collaboration and argued that design practices should focus on suiting technologies with the nature of collaboration (Dix, 1994). However, many CVE projects focused on supporting conferencing and social gathering, in which communications, rather than the management of data structures, is the key. By treating space and place differently, we will have no difficulty in understanding that the potentials of 3D virtual environments in support of collaboration lie more in artefact-oriented collaboration than in communication-cantered collaboration. This is also true for mVEs.

In the future, we will continue our research on multiscale space and multiscale place. For space, we are interested in two issues: how multiscale technology may affect users' access of spatial information at different levels and whether this technology can help users better integrate spatial information across different levels. Another direction we will pursue is to put multiscale places under practical lenses. We would like to see how people like materials scientists and architects would use multiscale places in the construction of cross-scale model structures in their work.

ACKNOWLEDGEMENT

This research was supported, in part, by Microsoft Research. We would also thank the anonymous reviewers for their helpful comments and suggestions

REFERENCES

Bederson, B.B., & Hollan, J. (1994). "Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics", *Proceedings of the ACM 1994 Symposium on User Interface and Software Technology (UIST '94)*, pp.17-26.

Bederson, B. B., Meyer, J., & Good, L. (2000). "Jazz: An Extensible Zoomable User Interface Graphics Toolkit in Java", *Proceedings of the ACM 2000 Conference on User Interface and Software Technology (UIST 2000)*, pp. 171-180.

Benford, S. (1996). "Tutorial on Collaborative Virtual Environments", Retrieved August 3, 2003 from the World Wide Web http://www.crg.cs.nott.ac.uk/people/Steve.Benford/CVEs.html

Benford, S., Bowers, J., Fahlén, L., Greenhalgh, C., & Snowdon, D. (1995). "User Embodiment in Collaborative Virtual Environments Papers: Advanced Media for Collaboration", *Proceedings of the ACM 1995 Conference on Human Factors in Computing Systems (CHI '95)*, pp. 242-249.

Benford, S., Fahlén, L., Greenhalgh, C., & Bowers, J. (1994). "Managing Mutual Awareness in Collaborative Virtual Environments", *Virtual Reality Software & Technology: Proceedings of the VRST '94 Conference*, pp. 223-236.

Benford, S., & Greenhalgh, C. (1997). "Introducing Third Party Objects into the Spatial Model of Interaction", *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work (ECSCW '97)*, pp. 189-204.

Bly, S., Harrison, S. & Irwin, S. (1993). "Media Spaces: Bringing People Together in a Video, Audio and Computing Environment", *Communications of the ACM*, vol. 36, no. 1, pp. 28-47

Bradner, E. (2001). "Social Affordances of Computer-Mediated Communication Technology for Geographically Distributed Work", *Extended Abstracts of the ACM 2001 Conference on Human Factors in Computing Systems (CHI 2001)*, pp. 67-68.

Casey, E. (1997). *The Fate of Place*, University of California press, Berkley, California.

Cutting, J. E. & Vishton, P. M. (1995). Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth, in: Epstein, W. & Rogers, S. (Eds) *Perception of Space and Motion*, Academic Press, San Diego, pp. 69-117.

Darken, R.P. (1996). “Navigating Large Virtual Space”, *International Journal of Human-Computer Interactions*, vol. 8, pp. 49-71.

DeSilve, L., Tahara, M., Aizawa, K., & Hatori, A. (1995). “A Teleconferencing System Capable of Multiple Person Eye Contact (MPEC) Using Half Mirrors and Cameras Placed at Common Points of Extended Lines of Gaze”, *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 5, no. 4, pp.268-277.

Dix, A. J. (1994). Computer-Supported Cooperative Work - A Framework, in: Rosenburg D. & Hutchison, C. (Eds,) *Design Issues in CSCW*, Springer-Verlag, London, pp. 23-37.

Fitzpatrick, G., Kaplan, S. & Mansfield, T. (1996). “Physical Spaces, Virtual Places and Social Worlds: A Study of Work in the Virtual”, *Proceedings of the ACM 1996 Conference on Computer Supported Work (CSCW ’96)*, pp. 334-343.

Furnas, G. W. (1986). “Generalized Fisheye Views”, *Proceedings of the ACM 1986 Conference on Human Factors in Computing Systems (CHI ’86)*, pp. 16-23.

Furnas, G. W., & Bederson B. B. (1995). “Space-Scale Diagrams: Understanding Multiscale Interfaces Papers: Navigating and Scaling in 2D Space”, *Proceedings of the ACM 1995 Conference on Human Factors in Computing Systems (CHI '95)*, pp. 234-241.

Gaver, W. (1992). “The Affordances of Media Spaces for Collaboration”, *Proceedings of the ACM 1992 Conference on Computer-Supported Collaborative Work (CSCW '92)*, pp. 17-24.

Gibson, J. J. (1997). *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, MA.

Greenhalgh, C., & Benford, S. (1995). “MASSIVE: A Collaborative Virtual Environment for Teleconferencing”, *ACM Transactions on Computer-Human Interaction*, vol. 2, no. 3, pp. 239-261.

Guiard, Y., Beaudouin-Lafon, M., & Mottet, D. (1999). “Navigation As Multiscale Pointing: Extending Fitts' Model to Very High Precision Tasks”, *Proceedings of the ACM 1999 Conference on Human Factors in Computing Systems (CHI ’99)*, pp. 450-457.

Harrison, S. & Dourish, P. (1996). “Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems”, *Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work (CSCW ‘96)*, pp. 67-76.

Hollan, J., & Stornetta S. (1992). “Beyond Being There”, *Proceedings of the ACM 1992 Conference on Human Factors in Computing Systems (CHI '92)*, pp. 119-125.

Linebarger, J. M. & Kessler, G. D. (2002). "GroupMorph: A Group Collaboration Mode Approach to Shared Virtual Environments for Product Design," *Proceedings of the ACM 2002 Conference on Collaborative Virtual Environments (CVE 2002)*, pp. 145-126.

Lea, R., Honda, Y., & Matsuda, K. (1997). “Virtual Society: Collaboration in 3D Spaces on the Internet”, *Computer Supported Collaborative Work*, vol. 6, no. 2/3, pp. 227-250.

Leigh, J., Johnson, A., & DeFanti, T. (1996). “CALVIN: An Immersimedia Design Environment Utilizing Heterogeneous Perspectives”, *Proceedings of IEEE International Conference on Multimedia Computing and Systems ’96*, pp. 20-23.

Mackinlay, J., Card, S., & Robertson, G. (1990). “Rapid Controlled Movement through a Virtual 3D Workspace”, *Proceedings of the ACM 1990 Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’90*), pp. 171-176.

Mine, M., Brooks, F., Sequin, C. (1997). “Moving Objects In Space: Exploiting Proprioception in Virtual-Environment Interaction”, *Proceedings of the ACM 1997 Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’97*), pp. 19-26.

Morrison, P, & Morrison P. (1982). *Powers of Ten: About the Relative Size of Things in the Universe*, Scientific American Books, Inc., New York.

Norman, D. A. (1990). *The Design of Everyday Things*, Doubleday, New York.

Perlin, K., & Fox D. (1993). “Pad: An Alternative Approach to the Computer Interface”, *Proceedings of the ACM 1993 Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’93*), pp. 57-64.

Pierce, J., Forsberg, A., Conway, M., Hong, S., Zeleznik, R., & Mine, M. (1997). “Image Plane Interaction Techniques in 3D Immersive Environments”, *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, pp. 39-43.

Poupyrev, I., Billinghurst, M., Weghorst, S., & Ichikawa, T. (1996). "The Go-Go Interaction Technique: Non-Linear Mapping for Direct manipulation in VR", *Proceedings of the ACM 1996 Symposium on User Interface Software and Technology (UIST '96)*, pp. 79-80.

Raab, A., & Ruger, M. (1996). "3D-Zoom: Interactive Visualization of Structures and Relations in Complex Graphics", *3D Image Analysis and Synthesis*, pp. 125-132.

Robertson, G., & Mackinlay, J. (1993). "The Document Lens", *Proceedings of the ACM 1993 Symposium on User Interface Software and Technology (UIST '93)*, pp. 101-110.

Sellen, A., & Buxton, B. (1992). "Using Spatial Cues to Improve Videoconferencing", *Proceedings of the ACM 1992 Conference on Human Factors in Computing Systems (CHI '92)*, pp. 651-652.

Singhal, S., & Zyda, M. (1999). *Networked Virtual Environments: Design and Implementation*, ACM Press, New York.

Snowdon, D., & Jää-Aro, K. (1997). "A Subjective Virtual Environment for Collaborative Information Visualization", *Virtual Reality Universe '97*.

Stoakley, R., Conway, M., & Pausch, R. (1995). "Virtual Reality on a WIM: Interactive Worlds in Miniature", *Proceedings of the ACM 1995 Conference on Human Factors in Computing Systems (CHI '95)*, pp. 265-272.

Ware, C., Gobrecht, C., & Paton, M. (1998). "Dynamic Adjustment of Stereo Display Parameters", *IEEE Transactions on Systems, Man and Cybernetics*, vol. 28, no. 1, pp. 56-65.

Waters, R., Anderson, D., Barrus, J., Brogan, D., Casey, M., McKeown, S., Nitta, T., Sterns, I., & Yerazunis, W. (1997). "Diamond Park and SPLINE: Social Virtual Reality with 3D Animation, Spoken Interaction, and Runtime Extendibility", *Presence: Teleoperators and Virtual Environments*, vol. 6, no. 4, pp. 461-481.

Zhang, X., & Furnas, G. (2002). "Social Interactions in Multiscale CVEs", *Proceedings of the ACM 2002 Conference on Collaborative Virtual Environments (CVE 2002)*, pp. 31-38.

Zhang, X., & Furnas, G. (2003). "The Effectiveness of Multiscale Collaboration in Virtual Environments", *Extended Abstracts of the ACM 2003 Conference on Human Factors in Computing Systems (CHI 2003)*, pp. 790-791.

Zhang, X. (2005). "mCVEs: Using Cross-Scale Collaboration To Support User Interaction With Multiscale Structures", *Presence: Teleoperators and Virtual Environments*, 14(1) pp. 31-46.