Papers

Beyond the Plane: Spatial Hypertext in a Virtual Reality World

Rosemary Michelle Simpson
Computer Science Department
Brown University
rms@cs.brown.edu

Scenario Component - Rosemary's Environment

Rosemary, a hypertext researcher, is working with a group of colleagues who are at globally spread out locations. She wears a lightweight head mounted display and uses her feet, knees, voice, and hands to direct the system and for telecommunications with her workgroup.

At the moment she is alone in her workspace, rapidly creating and organizing her information farm in preparation for a meeting. The first thing we notice about her workstation is the total absence of any keyboard, mouse, or computer, as we know it. The screen she uses is very large. It looks, in fact, like a large flat drawing board set at a 45 degree angle. At her feet are a set of pedals, not unlike those used by organists, except that these pedals don't seem to be attached to anything.

When you look at the video 'drawing board' you notice that in addition to the lack of keyboard there is a total lack of menus or windows on the screen. What you see is a series of 3D objects that are moving around without any apparent movement of her hands. Objects come into view, rotate, interact with other objects, and change their appearance and behavior as they interact with the environment. In short, the display seems to be alive.

You notice that the she periodically moves her head, appears to be listening, speaks, touches the screen in various places, occasionally making drawing gestures that are followed by the appearance of an object that then joins the rest of the menagerie. The appearance of the screen at times shimmers with the Star Trek effect as objects come into view and fade out of view. Other times an object moves, showing motion blur as it does, but then appears to be two transparent versions of itself with one version gradually fading away.

What you notice as you watch her work is that she is using her voice, moving her head, gesturing with her hands, and sometimes touching the screen. Her feet move to touch the different pedals, and occasionally she uses her knee to move a previously unnoticed knee pedal . Her whole body seems to be involved in interacting with this moving display. At one point she instantiates a keyboard, apparently by use of a foot pedal and a touch on the video drawing board where she wants it.

Part of the head mounted display is a microphone and an ear piece. With this she issues voice commands to the system and hears responses from the system about how long it's going to take for a particular action to occur.

Rapid gestural motions with both hands bring a 3D room into view, which attaches itself to a house in the cityscape already present on the screen. She speaks into the microphone and a graphical display, which we learn is a module interconnection tracker, appears with one of the nodes highlighted. When she gestures, it rotates, bringing other parts of the graph into view. Then, when she touches first the highlighted node and then another part on the structure, the node detaches from its original location and reappears in the new location. At the same time, the room also detaches itself and reappears, attached to a different house several streets over.

This room is her library and she has just positioned it within the library of CMU, whose physical location is thousands of miles away but whose virtual location is represented at the street where she just put her library. In fact the cityscape we see is her own personal working information farm, in which each location can have associated sounds, animations, and interactive simulations. By placing her library within the model of the CMU library she has made a link between her local information sources and the multimedia database at CMU. This link will later be available to her colleagues when they meet.

She hears a ring through the ear piece and Sonja, the language designer who is in Germany appears in her space - Sonja has her own room to which she has added touches that are her local versions of resources, that modify the basic structure which Rosemary has established for her guest rooms.

At one point in the discussion Rosemary gets up and moves around the room, touching various points in the room that have a small metallic dot on them This free movement is, first of all, possible because she is not connected by wires to her screen. In fact, the entire room is her workstation and is interconnected through the metallic dots.

Author's Note

This is a piece of the scenario for a paper-in-progress, "Beyond the Plane: Spatial Hypertext in a Virtual Reality World".

Everything in the scenario, which will present several very different environments, each of which maps to the cognitive style of the person using it and all of which interact with each other, must be based on a reasonable extrapolation of research currently underway or planned for the near future. This is science fiction that is a real possibility for the next ten years.

The system detects her position within the volume of space through sensors in her clothing and through the places she touches. This orientation modifies the representation of the information she's working with that she sees through her headpiece.

The system is aware of the environment at each point in space. Since the entire room is mapped within the system and since the middle of the room is where her central organizing model is located, all she has to do is to position herself there and ask the system for the model to be displayed in her head piece. She also can hear recordings that relate to the information. and can use the pointer in her hand to grab and move the contents of the model around. If she chooses, she can have her information farm/cityscape map to the entire volume of the room, so that she can move around in it and work with it without having to be at her video board.

At one point in her moving around she realizes that she needs information to answer a question asked by her systems co-worker who is working on some scaling problems having to with the distributed systems architecture project. The information farm that she's working with consists of many thousands of items. At any given moment she may need to have a sense of the farm as a whole or of some particular part, and this 'part' may in fact consists of dynamically generated collection in a dynamically generated structure. At this moment she wants to find and work with all of the items that have to do with operating systems language design, and interface design, so she describes what she wants to the system, asking that an agent be set up to find these elements. Since she knows she's going to need it in the future, she asks that this description be saved as a permanent agent and also asks that it be entered in the written record that will be produced as a result of the upcoming meeting. As she's describing what she wants the agent to be looking for, she remembers some documents in another library in Toronto Canada, so she asks that a system-active agent be set to search for those documents, bring them into her farm, and link them to the existing information on the subject.

Each of the dots placed around the room has multiple purposes. One function of these dots is to provide a spatial orientation within the three dimensional space of the room. Other functions of those dots are to provide an extension of that part of the room into virtual space, where virtual space means whatever it has been defined to mean for that location. Virtual spaces at different locations may be constrained by each other so that changes in one can affect the other. In addition, virtual spaces may be wormholes, or ports, into the virtual spaces of colleagues.

Each space can be overloaded to have several different representations depending on the overall context of the system at the time that Rosemary activates the dot, or on how she describes that dot to the system. All the dots have a representation in the cityscape on the virtual drawing board.

Notice that there are several different kinds of agents in this system. There's the information farm agent, who goes out and collects requested information, arranging it in a dynamically-determined structure. Information farm agents can also provide different kinds of views over the farm on command. In addition, there are itenerant agents that go out across the Web to gather information from designated places or to search for resources and bring them back into the farm. These itenerant agents communicate with the information farm agents, sometimes negotiating with them about how to store and structure the results. Each agent can simply be a transient computational entity, instantiated for a one-time task, or can be a complex avatar with a persistent and evolving context that modifies and is modified by current requests. The avatars assume personalities congruent with their tasks and can be given visual and behavioral attributes that reflect their behavior.

At one point Rosemary, who is designing the hypertext system of the future, wonders what Vannevar Bush would think about some of the problems she's investigating and so she asks the multimedia database at CMU to build her an interview with Bush. The base for the interview is created from records of what Bush had said and written about the subject over the years, as well as from commentaries from trusted sources about his work. Once that base is built, she then interviews Dr. Bush. He points her to some of his management writings for insights on how to resolve the social issues involved in handling interactive telecollaborative hypertext systems. A related footnote link points her to the talk Tim Berners-Lee gave at the Brown-MIT Memex Symposium in October, 1995 where he described social plumbing.

Next, Rosemary wants to talk with Bob, the database person on her team, about issues having to do with linking the volatile and heterogeneous collection of materials from the information farm into the global database. She also wants to talk about how the database is going to interact with the new distributed operating system/language environment they are working on. Now, she knows that Bob would be asleep at this time since he is in

Hawaii, so she calls up the 3D email avatar/agent of Bob that she has in her system and asks if there has been any new information added while she was asleep. She requests that he scan over the collected set of information from Bob and see if there is any theme that might be related to the question she's asking. What she is asking this avatar to do is, in effect, to come up with a new set of information by organizing all the data pertaining to this subject and Bob's contribution to it.

What she finds out raises some other questions that she then communicates with Bob's home avatar. Even though Bob is asleep, his home avatar can access those things to which Rosemary has access privileges. In fact, there has been some new material gathered but not yet communicated that does answer the question but raises a third possibility in her mind and so she stores it with both of Bob's avatars, including a link to the meeting agenda.

As Rosemary is walking around the room and thinking about the upcoming meeting, she puts on some music. It's the Bach Musical Offering and its structure makes her think about relationships among some of the problems she and her team have been working with. So she asks the system to record the particular part of the music that she's interested in and link each voice in the fugue to a particular sub- problem description. She decides to watch a visual representation of the sub-problem interactions as the music plays. To make this easier to view, she selects object representations for each sub-problem , links them to particular voices in the music, and allows them to modify themselves relative to the music - size represents loudness and spatial relationships represent musical relationships.

Then she starts dancing...

She decides to enter and modify what is happening so she assigns one of the sub-problems to herself. Then she starts dancing, listening to the change in music as she interacts with and modifies the objects.

She then goes to a painting by Kandinsky which she realizes will provide, in combination with the music, a map into the information structures that she's evolving. One of the insights that the Kandinsky painting evokes is that the information structures contain significant discontinuities along multiple overlapping and quite different continuums. Some of these structures exist in different times and, while they never touch directly, they indirectly influence each other.

Since the members of her team have different mind styles and one of them is blind, she needs to map her insights into a variety of modalities into order to present them in a way that each team member can understand and thus contribute to.

The blind team member is a gifted linguist, singer, and dancer who understands information kinesthetically. He uses sonar-type devices to perform elaborate dances that integrate and express information structure relationships, and then translates that understanding into precise verbal descriptions. His system detects his location in space as he is dancing and creates a node-link structure that represents the dance. The nodes encapsulate text, sound, and spatial coordinates and the links trace the trajectories. The nodes and links are time-typed to show the relationships in time and space among the nodes.

The resulting hypertext can be played as an animation. In addition, the dancer can play it back, create new versions branching out from particular nodes, and change the nodes, links and time relationships. These versions can be played simultaneously, as in HyperCafe, or selected separately. In addition, other types of information may be attached to the nodes, such as videos, links to other information sources, comments, etc. The overall effect can be similar to the songlines created by Australian aborigines to express their understanding of space, time, and events.

Next, as she is trying to express an idea that she's not been able to bring into symbolic space, Rosemary asks the system to turn on the body-locating and recording mechanism and starts to dance again. The system records the trajectories of her motions, the speed, and emotional state changes (through sensors attached to her clothing). She calls for particular pieces of music which are linked to the place in the dance where she was when she asked for that music. At other times she touches books and pieces of art. The system records which book and which part of the artwork she touches, linking that information into the dance structure. At one point she stops the dance and asks the system to replay parts of it. She wants to look and see how the various components relate. She makes links between later and earlier parts of the dance, as well as to the music. This results in multiple music links to various parts of the dance as well as links between the art and sculpture in the room to the music and dance structure.

She then calls up other dances - great dances of the past and other cultures - Isadora Duncan, Bali, African forest dances - and links them in, sometimes splicing them into her

dance, sometimes having them proceed in parallel, sometime requesting that an indeterminacy function be attached to vary how the link is made.

During the meeting, the various group members come into the public space, taking seats on the couches. Some of them prefer not to be visually present and are represented by a variety of avatars, sometimes only as voices. As they discuss the issues on the agenda, the 3D information farm/cityscape is shown. They walk around the cityscape and look at the various parts of the structures. At some point they request various perspective changes - fisheye views, perspective wall, cone tree, etc. At other points they simply walk through the city, asking for different representations of the concepts, data, structures, and behaviors. Each person can ask the system to map their preferred mind style representation into the information.

To be continued...

Officers Reports, continued from page 4

Is there a life for Hypertext after the Web?

What are the implications of all this for the future of SIGLINK? Hypertext is not a closed, well-defined, domain. It typically is a trans-disciplinary domain. It ranges from computer science to human sciences and to art. The specificity of people working in Hypertext is a wide mental bandwidth which enables them to master techniques that are to be integrated for developing hypertext. The success of hypertext relies on the ability to integrate human computer interaction, multimedia, collaborative work, electronic documents, typography, graphics, rhetoric, networking, databases, task modeling, information retrieval, artificial intelligence, multi-agent approaches, object orientation, software engineering, and so on. The key words are openness and tolerance. We cannot live alone. We do need the Web and, of course, we hope the Web does need our know-how. Our strength and future is in mastering the integration of the widest set of techniques into systems which are agents to help reading, not in promoting a single and unified solution to build a hypertext.

The key feature of such agents, that make them different from other ones and which is responsible for their "Hypertext" label, relies on the interaction style: the computer is no longer visible from the reader side. A hypertext is an interface which should completely hide the computer. The reader "interacts" only with "responsive documents". I would conclude with a single thought: if the computer is still obvious in an application, it is probably not yet a true hypertext application! So, we still have a long way to go.

Standards, continued from page 5

The SGML-ERB and SGML-WG now turn to link semantics and stylesheets, and hope to achieve a similar synergy. DSSSL is known to provide the stylistic power needed, and so far looks amenable to simple subsetting such as worked for SGML. Readers who wish to participate should apply to chairman Jon Bosak, `jon.bosak@sun.com`, explaining their qualifications.

HyTime is still undergoing its "Technical Corrigendum" process, though I am happy to report that a draft of the final amended text has been distributed. It benefits greatly from DSSSL's formalisation of documents' tree structures (since it does infra-document locations largely by tree-node positions).

It seems to me that a key area for further development involves developing a taxonomy of link semantics (both rhetorical and structural). At one end are links that trigger arbitrary programs (useful and necessary, but hardly amenable to analysis); at the other are systems with "click to move to card X" as the only semantic; and a whole continuum between. With documents larger than one screen and with multi-window, interfaces, a whole range of hypermedia interaction issues come to the fore: how should go-back lists work when many panes are present, and the user can change context with merely a glance? How should we handle links whose scopes partly overlap? What does it mean to link to a discontiguous target, especially when editing happens? These are only a few of the most obvious questions that are becoming daily experiences, and need first analysis, and eventually, once better understood, some standard terminology and models.