Chapter 7
Orienting in Virtual Perceptual Fields

7.2 Visual Momentum Across Displays in Virtual Perceptual Fields

I claim that designers of computer based information systems create a virtual perceptual field. It is easy for designers to create virtual perceptual fields where the observer must function without the assistance of orienting perceptual functions because of the keyhole property and because the computer affords designers freedom from the constraints acting on the referent real world objects (Woods, 1995a). When dealing with virtual perceptual fields in the computer medium, the burden is on the designer to explicitly build in mechanisms to support the operation of the orienting perceptual systems and the functions that they perform in a fully articulated cognitive system adapted to a changing environment. The designer of a virtual perceptual field develops the characteristics that allow or that undermine the role of orienting perceptual functions.

There are many techniques that can be used to support the coordination between orienting perceptual systems and focal attention in a virtual world (e.g., Woods, 1984). The remainder of this chapter describes several techniques organized by the heuristic of "visual momentum across displays." This is just one approach to aiding navigation. Some of the techniques suggested by this point of view are similar to or overlap techniques that have been developed by others (e.g., Furnas, 1986; Henderson and Card, 1986; Lamping, Rao and Pirolli, 1995; Kahn, 1995). In the case of other innovations, Visual Momentum may provide a more general and abstract basis to understand why these systems aid navigation.

Visual Momentum

Given that meaningful activities often involve moving across displays, designers should be able to aid the transition from one view into the artificial data field to another. The concept of Visual Momentum is borrowed from perception and cinematography (Hochberg, 1986) and refers to the impact of a transition from one view to another (a cut in cinematography) on the cognitive processes of the observer, in particular on the observer's ability to extract task-relevant information. As applied to HCI, "the amount of visual momentum supported by a display system is inversely proportional to the mental effort required to place a new display into the context of the total data base and the user's information needs. When visual momentum is high, there is an impetus or continuity across successive views that supports the rapid comprehension of data following the transition to a new display" (Woods, 1984, p. 231).

At one end of the dimension lies poor transitions which consist of (a) total replacement of one view for another and (b) the absence of any visible cues to the virtual field of possible views. When Visual Momentum is low, each "glance" into the artificial data field is independent of previous glances so that the observer must reorient from scratch to each new view as it is called into the limited viewport.

At the other end of the dimension the observer works within a conceptual space in which individual views are grounded. A conceptual space depicts relationships in a frame of reference (Woods, 1995a; 1996). In between lies a variety of techniques for building a sense of a conceptual space analogous to a physical space so that orienting and moving about the virtual perceptual field can employ the same perceptual and cognitive processes that allow us to fluently explore and reorient to new events and changing views in naturally occurring physical spaces. Some of the techniques to increase Visual Momentum are longshots, landmarks, content-laden cues to structure, spatial dedication, coordinating what can be seen in parallel and what in series as a function of task demands (Henderson and Card's "rooms"), center-surround, side effect views, cues to status. All of these techniques and many others (e.g., trails, bookmarks, 3D spatial metaphors) function as visible cues to the structure of the space of possibilities and cues to the status of those different parts of the domain represented by the artificial data field behind the keyhole. Designers can orchestrate these kinds of techniques to create what the user experiences as a tangible conceptual space to support effective workspace coordination.

The Longshot -- Showing the Big Picture

It is commonly accepted that an overview display will support coordination or navigation across the many views available within the virtual data space. In Cases 4 and 5 designers provided what they thought were summary displays, but these displays did not support practitioners and were rejected or little used. Case 4 illustrates that simply changing the relative density of data visible as the field of coverage of the underlying topology expands or contracts does not necessarily support effective across display transitions. Case 4 also shows that simply aggregating details into a display of generic status filters out too much information for it to support operators as an overview of system state and how state is changing relative to goals. What do designers need to do to create effective overviews?

Woods (1984) referred to summary displays that actually serve as effective orientation and navigation aids as longshot displays. In cinema, a longshot is an establishing view that shows relationships between characters and summarizes relevant information. It keeps the viewer involved in the flow of the plot by allowing him/her to step back from the details, discover why these details are important, how they relate to previous views and to establish a frame of reference to help the observer comprehend upcoming views.

The successful longshot in computer based information systems serves the same purposes. In essence it is a kind of global map (Billingsley, 1982; Kahn, 1995; Kahn, 1996). It helps practitioners step back from the details of the monitored process to assess overall system status. It helps them decide where to look next within the system. It helps them relate the view currently under examination to previous views and to integrate new views into their assessment of process state as they are called into a visible viewport.

Three functions contribute to the effectiveness of a longshot -- the status summary function (cues to status), the orienting function (cues to structure and mapping the structure to domain semantics), and the movement function (direct manipulation).

<u>The Status Summary Function: Status at a Glance</u>
Longshots containing status summary information allow users, in a mentally economical way, to step back and assess their overall situation with respect to the underlying process, device, or activity they are engaged in. By providing task relevant status information, longshots are <u>content-laden</u>. By showing the status of the process or activity behind the computer interface, longshots can help users decide where to look next in the virtual field. The concept of content-laden navigation aids which provide cues to status as well as to structure will echo throughout several of the techniques to enhance Visual Momentum across displays.

In a physically distributed control center (an open workspace), an experienced operator can stand at the back of the room and gain enough information to describe the status of the system, i.e., the current state, the current epoch or phase of operation, the direction things are headed in (e.g., deteriorating or recovering), the stance of actors towards the system (such as routine operations or a tense critical period). If a longshot includes the status summary function, it will externalize pertinent summary information found within the display structure, and thus allow an operator the opportunity to step back and assess the status of the monitored process and to quickly see how the system is behaving. The longshot provides cues to status across the entire system so that practitioners can remain in tune with changing conditions and the "big picture" while they are focused on a detailed task or part of the underlying process. Cues to status are an important technique to enable practitioners to easily "check read" or peripherally pick up what might be interesting changes in other parts of the underlying process that should guide a shift in attentional focus (Woods, 1995b).

Status summary information can also support tasks in domains that are more self-paced. Tasks in these domains often include updating and maintaining information within the display structure. For example, some spreadsheet users update lists of information or financial figures each month. A longshot could help spreadsheet users by providing status summary information about these updating activities (i.e., which areas of the sheet have been updated, and which ones have not; Watts, 1994).

For a longshot to support the status summary function, it must include the following attributes:

1. The summary information must be distilled.
The point of a summary is to distill the relevant information that characterizes the situation as whole in a concise, recognizable form. The relevant information must be represented in a way so that observers can size up the state of affairs at a glance. One indication of a concise distillation of relevant factors is that it is informative even if it is shrunk to a relatively small size. Summaries in which information is solely conveyed through the digital display of elemental data are rarely adequate for this function. Similarly, designers can easily over-summarize and provide too little data to be of value to experienced practitioners (e.g., group alarms).

2. Information in the summary must be abstracted.
To broaden the view of the system, information in the summary display should be abstracted to a higher level of information than that of raw data and other details about the system. It is important to note that abstracted information is not simply a lack of detail, but rather an integration of details that informs the observer about higher level questions (Vicente and Rasmussen, 1992). It involves collecting information from various areas of the system that speak to broader issues about routine and exceptional conditions. To do this, the designer may need to transform lower level data, integrate these data and contrast them to related values. Examples of abstracted information include answers to questions like: what mode is the system working in? is the system functioning normally, or is there a malfunction in one of the subsystems? what activities are taking place at this time?

3. The longshot must include information about change and sequence.
If practitioners are assessing the status of a system, they must be able to recognize patterns of change within that system. In addition to information about the current behaviors and states of the system, information about what has happened recently, and what kinds of trends may be developing often contributes to the overall assessment of the status of the system. For example, in event-driven worlds faults often result in a cascade of disturbances. A content-laden longshot will help the practitioners keep track of the big picture of disturbance evolution while they are pursuing diagnosis and response.

4. The longshot must show information that is relevant to the viewer's context.
The information in the longshot view should answer questions and provide information that makes sense to a practitioner in his or her task context. Examples

of potentially relevant information include the activities that are currently ongoing and helping practitioners see if events are developing in accordance with their expectations. The longshot provides the larger context about the semantics of the field of practice in which one examines different and more focused views.

5. The longshot should support "check reading."
An abstracted and distilled overview should help users pick out what conditions or changes are potentially interesting quickly and in a mentally economical way given the current context and ongoing lines of reasoning.

One technique we have used for fulfilling the status summary function is to use a reduced scale version of a pattern-based or emergent property display as a longshot (e.g., Watson, Eastman, Woods, 1990; Ranson and Woods, 1996). The patterns in this type of integrated display are much larger (lower spatial frequency) than the elements from which they emerge. As a result, the patterns still stand out at a reduced scale while details become obscured (essentially the size reduction acts as a high spatial frequency filter). The patterns are effectively a dynamic summary of the status of the portion of the underlying process or device they cover. The patterns support the "check reading" activity that allows practitioners to peripherally pick up what might be interesting changes while they are engaged in other lines of reasoning (Woods, 1995b). Note that users have to be familiar with the fully detailed version of the pattern based display for the reduced scale version to work successfully as a longshot.

<u>The Orienting Function</u>
The orienting function of a longshot helps operators orient to where they are (the currently visible views) relative to the set of views that they could examine in this context. The orienting function helps them comprehend cuts from one view to another. Longshots contain map-like characteristics which show users where they are located in relation to the important parts or landmarks within the virtual perceptual field. The map can serve as a representational framework for capturing what options are relevant to the current situation, which options have been recently selected/inspected (a trail), and support user browsing through potentially relevant views. The latter raises the question of how a longshot can be developed to invite exploration (Norman, 1988).

The structure of the overview display should reflect the structure of the views within the workspace as they represent the semantics of the field of practice. This type of conceptual map makes movement in the large network of possible displays the equivalent of moving in the semantics of the domain from the point of view of a practitioner. When a summary supports both this function and provides information about significant status and change in the underlying process or device, (the system status function), the longshot helps the operators formulate relevant questions and helps them decide where to direct their attention next.

Supporting the orienting function helps produce interface transparency where the practitioners can concentrate on their activities and goals in their field of practice instead of being focused on the interface control mechanisms themselves.

The following principles should be followed to support the orienting function of a longshot:

1. The longshot must be coordinated with other views.
The longshot needs to be coordinated with the other more local kinds of views available to users and the reverse as well. A longshot serves as a visual representation of the virtual data space otherwise hidden behind the narrow keyhole. This means the longshot must capture the structure of the virtual data space (cues to structure). When it fulfills this purpose, the longshot can serve as a reminder of what types of local views are available and which of these views are relevant to the user's task context. The goal of coordinating the different views is to enhance the user's ability to extract and integrate information across cuts or transitions from one view to another (the fundamental definition of high visual momentum).

2. The overview display should include relevant frames of reference.
Coordination between longshot and other views depends on choosing, coordinating or creating frames of reference related to the practitioner's tasks in the field of practice. A frame of reference specifies relationships between parts and part-whole relationships. Depicting these relationships for a field of practice creates the longshot. The longshot defines and makes apparent the frames of reference within which other views exist and can be integrated. The other views have a meaning as a portion, neighborhood or perspective of the larger frame of reference defined by the longshot. The portion currently visible becomes a place (where the user is "located") within the larger space defined by the longshot. However, system designers sometimes create "overview" displays that do not show the operators where they are in relation to where they could be. Problems range from showing a list of menu options without indicating where the operator is located within the display structure, to showing a broad view of the structure of the monitored process without indication of how the process is broken up into displayable chunks (Woods, 1991).

3. The overview display must always be available in parallel with other views.
Experience from several projects (Woods, 1984; Bolt, 1984; Case 4) all point out that a longshot is effective only if it is available physically in parallel with other types of views. The point of a longshot is to help the practitioner know where to look next in the virtual field. If the longshot is not constantly available, then the practitioner has to decide when it is appropriate to consult it with the associated cognitive burdens. If other constraints make it impossible to constantly show the longshot, relax those constraints. An effective longshot is one fundamental kind of view which should be available in parallel in a well coordinated workspace.

4. A longshot combines information about the state and behavior of the monitored process with indication of the different types of available views relevant to the context.

When the representation of the types of available views is combined with the information about the monitored process, the overview display shows where operators could shift their attention within the display structure by showing what is happening in other areas. Therefore, a longshot integrates abstracted information about the status of the system while indicating the set of options or other views that are of potential relevance (Mitchell & Saisi, 1987). In effect, a longshot serves as a map of possibilities annotated with distilled data about the state of the process.

The Movement Function

In addition to showing overall system status and giving clues for where to look next within the system, the overview display should indicate how an operator can move to an area of interest within the display structure. The longshot functions as a map of other display possibilities in part. The map should make it apparent how one calls up or moves to or navigates to other views. Norman (1988) refers to this problem as the gulf of execution.

As a map of possibilities, a longshot provides a view of multiple options in parallel and shows the relationships between these options at multiple levels. This has led Billingsley (1982) and Woods (1984) to emphasize the need for maps or map-like structures to aid movement through a virtual space over a series of menu choices. Too often, systems of menus are designed in such a way as to force users through a long series of menu options which provide few options in parallel and which produce multiple across-display transitions without providing any content to the observer (see many Web sites). The longshot can support the movement function by allowing users to select from it the views for display in other viewports. Kahn (1995, 1996) shows a variety of ways that maps can be used to support movement within a Web site.

But a longshot can be more than a structured menu with multiple options in parallel. Since the longshot is a map that organizes parts and relationships between parts for the domain of interest, it is one of the prerequisites for providing a direct manipulation interaction (Hutchins, Hollan and Norman, 1986). Designers can use this map to allow users to directly specify where they want to go or how to instruct the system without thinking in detail about the interface mechanisms. Many of the interface and navigation problems in Example 2, the infusion device for home care of pre-term labor, discussed at the opening of this chapter disappear once a meaningful frame of reference for user activities is found (Obradovich and Woods, 1996). In this particular case, user activities are all about different rates of infusion over different time intervals and different bolus sizes at different points in time. In other words, what is informative in this field of practice are dose-time relationships. Since none of this is visible in the infusion device, the designers are forced to create awkward, arbitrary, and cognitively demanding sequences of displays and interaction to program the device. Not surprisingly, the result is classic HCI deficiencies in the device producing typical user problems and errors (Norman, 1988).

The solution to these problems is to provide a dose-time structure as the basic frame of reference. Users will be unable to see all of the relevant dose-time picture in detail at one time because of multiple therapy plans and because the relevant information may stretch out into the future or back into the past. But a longshot of the dose-time frame of reference can support navigation across more detailed views, and it can support direct manipulation mechanisms for instructing the system about desired therapy plans. This radical restructuring of the basic concept behind the interface, based on the semantics of the field of practice, would completely restructure the nature of the interaction. Users would no longer have to learn or remember the syntax of the interface; instead, the interface would match the semantics of their activity--setting up, modifying and monitoring dose-time relationships.

Longshots contain two kinds of information: information about the underlying process or field of practice (e.g., state information) and information about the virtual data space (e.g., what views are available). It is important to perceptually distinguish these two different categories of information. Case 1, earlier in this chapter, is one example where the visual indications of where one can click to open new views/windows are incomplete and ambiguous. Similar confusions about which visual marks are "clickable" and which are not have been found in usability tests of Web sites (Nielsen, 1996).

Status summary information, combined with the multiple parallel options and structured relationships between parts, then can act like a preview function to help the observer to quickly find and focus in on what is interesting given their context and goals.

Longshots allow practitioners to step back from the details of a part of the underlying process or one subtask to see important trends and activities about the process as a whole. This role means that longshots can be an important element in creating open workspaces that support cooperative activity if the longshot is available in common to the entire work team. Just because a display is available to all through a large screen display or other shared display mechanism is not sufficient for it to function as a longshot. The criteria in this section need to be met for the overview to contribute effectively to the shared workspace.

When an overview display provides the three functions discussed above--(1) status summary information, (2) visible cues to the semantic structure of the virtual field, and (3) allows practitioners direct mechanisms to shift their "gaze" within this space--then the overview functions as a longshot. The WWW is currently has many sites which lack effective longshots and demonstrate the cognitive and performance consequences as users are forced through multiple steps blindly trying to get to information relevant for their purposes. Because effective navigation is central for users to be able to take advantage of the Web's power, we have seen a wave of innovation in a number of techniques, most notably techniques for mapping web sites. Kahn (1995, 1996) provides many examples, principles and a method to help designers create effective longshots and increase visual momentum across displays for Web sites and other human-computer interfaces.

Landmarks

Another way to support navigation in the computer medium is to include landmarks in the artificial data field (Woods, 1984). Hochberg and Gellman (1977) define landmarks as "features that are visible at a distance that provide information about location and orientation." In the computer medium, landmarks are features in the interface that are visible at a glance and provide information about location and orientation.

Passini (1992) notes that landmarks play an important role in navigation in physical space. For example, the Space Needle, which can be seen from a distance from most areas of downtown Seattle, can be used to keep track of where people are in relation to where they are going. It provides information about both the relative distance they must travel to reach their destination, as well as the direction they must travel to get there. Once people become familiar with the area around the Space Needle, the structure can also serve as a reminder of the buildings and areas that surround it. In this case, the Space Needle serves as a summary of the surrounding area, as well as a frame of reference to help people establish their location relative to their destination.

Landmarks support repair processes for people who have become disoriented. If people lose their sense of where they are located, they can look for familiar landmarks to help them re-orient to their surroundings. For example, if people decide to take a new route to work, and the street takes unexpected turns, they can look for familiar landmarks in the distance to help them figure out where they are located, and which direction they are going. Landmarks also support repair processes in a virtual environment. If computer users attempt to choose a display from a menu, and mistakenly choose the wrong view, they can use the landmarks in that view to figure out where they are located within the virtual space.

Content free landmarks are structuring cues; content-laden landmarks provide perceptual cues to structure and signify something about the content of that area (e.g., the topic or the summary status). Both types play a role in building a virtual perceptual space.

Watts (1994) followed up her field work on the navigation strategies adopted by extended spreadsheet users (example 4 in the introduction to this chapter) with a simulation study that examined the role of landmarks and longshots as well as other cues to structure in supporting navigation. In this study, experienced users of large extended spreadsheets were given data related in a large set of interconnected data tables and asked to carry out various tasks modeled on the tasks user perform in actual organizations. She varied the cues to support navigation such as landmarks, longshots and others (e.g., spatial dedication). The results showed that landmarks are important cues for aiding navigation. Participants more directly targeted relevant data with landmarks than without. In fact, they considered landmarks so important that five of the seven participants when working with an extended spreadsheet without landmarks stopped performing the task and actually added landmarks to the extended spreadsheets before continuing with their tasks. Both content-laden and content-free were important to the participants in the study.

Spatial dedication

Information is spatially dedicated when it appears in one fixed physical location. The classic exemplar is the physically distributed control center where each data channel occurs in a fixed location and, ideally, functionally related data channels are located in one region of the physical space.

The fixed spatial structure of data serves as a memory aid for users who are familiar with that structure. While the entire field of data is not visible in parallel, the fixed spatial structure allows users to travel directly to desired data channels or topics (e.g., the electrical portion of the plant is always just to the right of the steam portion or the feedwater pumps are below and to the right of the boiler). For example, if extended spreadsheet users always keep their raw data tables in the top left region of the extended spreadsheet, and their calculations in the bottom right area, they will always know where to find those kinds of information.

In a control center where data is physically distributed and spatially dedicated and operator engaged in one activity can still see or pick up changes in other parts of the control center. Since data channels are in fixed locations and groupings, if they notice a perceptually distinct change in one location or area, as experienced operators in that space, they know something about what has changed without directing focal attention to that specific data channel or set of data channels. Similarly, glancing at an area of the control center provides a quick check read of what is going on in that region of the monitored process.

The same issues apply to computer-based systems and virtual perceptual fields. We consistently observe users spontaneously try to introduce some degree of spatial dedication in many different kinds of settings (where they place the same kind of information in one consistent location) to help themselves overcome navigation burdens. For example, in Case 1 (anesthesiologists in cardio-thoracic surgery; Cook and Woods, 1996) and Case 3 (space mission operations) practitioners invested effort during low tempo periods to set up their computer-based workspace as a fixed tableau of views tailored as best they could to the demands of the tasks they perform at higher tempo, more critical periods. In other words, they invested effort to avoid or minimize the need to navigate and transition across displays during these potentially busy periods. Another source of information about navigation needs comes from cases where practitioner's attempts to avoid navigation break down. When circumstances conspire to force practitioners outside of their preferred configuration of views, they become caught up in the interface itself and recognize themselves that this activity is a diversion from their responsibilities.

These cases of practitioners expending effort to avoid navigation burdens come from event-driven, high consequence domains. However, even in apparently self-paced domains the pace of work varies and there are periods of relatively high demands. Watts (1994) also observed spreadsheet users invest effort to organize a dedicated information space to avoid navigation burdens. This kind of user tailoring was limited and brittle in this case as in the medical and space examples because the computer-based system was not explicitly designed to support this kind of user behavior.

These cases illustrate the repeated observation that users seem to prefer a spatially dedicated representation even if it is crude and deficient in other ways over keyholed computer systems despite apparent flexibility to call up many different features, displays, and options. Why? We think this phenomenon occurs because low visual momentum computer systems create new burdens for already loaded users. If practitioners used the computer-based system in the ways that designers envision, then they will be forced to interact with and devote limited attentional resources to the interface at times when their attention should be most focused on their job. Practitioners avoid this situation by throwing away flexibility and using characteristics of the computer to create their own spatially dedicated set of views of the underlying process or device.

"Rooms": What can be seen in parallel and what in series

The fragmentation of information into small chunks connected in large networks inevitably leads to the dissociation of related data. Data that must be seen together to support user tasks are often physically located in separate areas of the interface. One problem that can result is what Henderson and Card (1986) called "thrashing." If users must serially view each piece of information related to their task, they are forced to thrash back and forth between the displays to meet their task goals.

One way to reduce the problems caused by the dissociation of related data is by allowing users to see related views in parallel. This capability expands the narrow keyhole provided by the CRT to coordinate inter-related views. For example, the solution to many of the navigation problems in Case 1 (the operating room information system) is not to be found in more efficient navigation mechanisms. Rather the source of the problem is that practitioners using this interface cannot work in parallel on two kinds of activities (Cook and Woods, 1996). In this context there are two fundamental kinds of information and functions in a patient monitoring system: (a) monitoring patient vital signs which is always an important task and (b) various kinds of special asynchronous operations such as the sequence of tasks involved in measuring and computing the patient's cardiac output (a third category is interface manipulations). In the system studied, these two kinds of functions compete for the same limited real estate (opening up the window to support measuring and calculating cardiac output obscures a significant part of the patient's vital signs). But one kind of information (patient vital signs) needs to be continuously available, while the other is needed only on particular occasions. Thus, from a workspace coordination level of analysis, two parallel viewports are needed--one sufficient to show patient vital signs continuously and the other to handle the different kinds of asynchronous capabilities.

Another example comes from a hypertext system that was used to prototype an electronic version of the multiple paper volumes of emergency procedures used to guide operator actions in nuclear power plant accidents (Elm and Woods, 1985). The volume of procedural guidance is quite large, and the material is organized at multiple levels in a way designed to help operators apply the plans that best fit the actual problems in the plant. Because problems can evolve and change and because of the possibility of complicated diagnostic situations, the procedure system was designed with double checks and branches which were implemented in paper through foldout pages, cautions at the beginning of a task, and kick out steps to other parts of the procedure system. For example, while working through a particular procedure, various notes, cautions, and double checks would be relevant the entire time operators carried out this task (e.g., if a parameter exceeds a limit, stop doing this procedure and switch to another). This material was included in the paper system as a foldout page so that operators could see the double checks at the same time as they followed the specific steps in the procedure. The paper procedural system also contained background material about particular steps or sections of the procedure.

Converting the paper procedures to an electronic format appeared to offer a variety of advantages which could be obtained simply by converting the current paper system into the format of this hypertext system. The hypertext system was organized as a hierarchy of frames. The network consisted of two kinds of frames. Menu frames that contained a few options to select different parts of the procedure space (e.g., procedures for different kinds of faults and procedures to ensure that basic safety concerns were being addressed successfully). Content frames contained actual procedural guidance. Users could only see one frame at a time, either one content frame or one menu frame. Because of the small size of the frames, each contained only a few (one to three typically) statements from the paper procedure documents (e.g., steps, cautions, notes). As a result, finding task relevant procedural material could involve moving through as many as 4 menu frames before arriving at any procedural guidance about how to handle the evolving plant problems. Examining other material related to the current Content frame on view would involve calling up one to several menu frames before arriving at the related material. Each transition was an act of total replacement, and users could only look at one frame at a time.

Not surprisingly, usability tests in a simulator showed that different kinds of users (the procedure developers, experienced operators and experts in the hypertext system) all got lost in the network of displays in the sense that they could not follow the procedures in synchrony with the evolution of the incident in the power plant.

The navigation problems arose in part because users could not see related material in parallel (other problems existed as well, such as no longshot was available). Some of this was domain specific (e.g., seeing a procedure and the relevant foldout page in parallel) and other instances were generic (users were unable to see options and content in parallel). As in the operating room case, the navigation problems could only be addressed by changing the workspace coordination to allow users to see inter-related views in parallel. Figure 7.7 illustrates the workspace redesign.

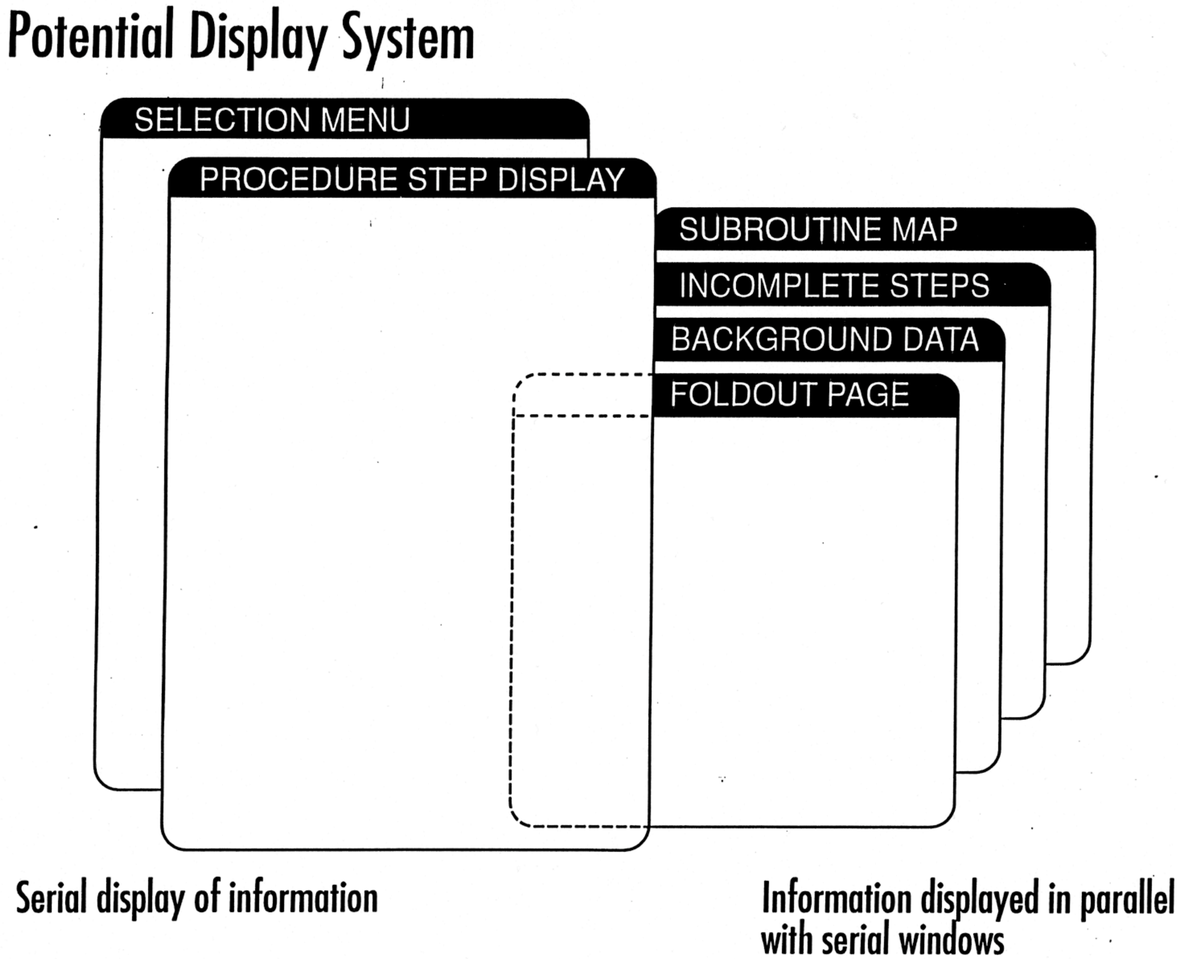

Figure 7.7. The workspace redesign for a computer-based version of the emergency procedures in a nuclear power plant showing what kinds of views are available and which can be seen in parallel and which in series (from Elm and Woods, 1985).

Watt's 1994 field study of users of extended spreadsheets also showed that users will invest time and effort arranging their workspace to be able to see inter-related views in parallel. For example, study participants consolidated related information so they would not have to interrupt their tasks to attend to the interface management task of navigating between distant areas of the extended spreadsheet.

Designing a coordinated workspace is about deciding what kinds of views are available and which can be seen in parallel and which can be seen serially. This is a shift from simply building up a collection of single views to creating a "task-tunable" workspace of multiple views or perspectives (e.g., Card, Robertson and York, 1996).

1. Developers should study/analyze what views need to be seen in parallel (this is conditional on what kinds of views are being considered or designed, and it requires understanding user activities during coherent tasks). Henderson and Card (1986) suggest empirically tracking what views users call up together to identify sets of views that need to be coordinated. Methods for cognitive task analysis that analytically map meaningful domain contexts identify views that are inter-related and need to be seen in parallel (e.g., Mitchell and Saisi, 1987; Woods and Hollnagel, 1987).

2. Explicit representation of the workspace in terms of the kinds of views and their inter-relations is a prerequisite for the design of a coordinated workspace (Malin et al., 1991). Figures 7.2 and 7.7 are examples of explicit representations of the workspace (other examples include state-transition diagrams and innovative display ideas such as the document lens, Robertson and Mackinlay, 1993, or the information grid, Card et al., 1991). Developers very often leave this information out and represent the design only in terms of the baseline views that appear by default in the standard viewports. When developing or evaluating a computer-based display system, we have found it necessary to characterize prototype systems in terms of a coordinated set of multiple views. Explicit, thoughtful design at this level of analysis of a computer-based display system cannot go forward without some way to represent the coordination of multiple views. The simple act of trying to represent the workspace often reveals many potential user problems, initiating rounds of design and evaluation that can lead to new innovations.

3. In general, users are likely to need to see specific kinds of views in parallel, e.g.,

- content and options need to be seen in parallel;
- longshot and local views need to be seen in parallel.

Other techniques are based in part on providing inter-related views in parallel (e.g., the center-surround technique and side effects views discussed in the next sections).

4. Designers should make provisions for users to be able to compose, save, and manipulate sets of views as a coherent unit, e.g., Henderson and Card's Rooms concept.

Center-Surround

In a natural perceptual field, people are able to switch attention and re-orient to interesting information in their environment (Rabbitt, 1984). The coordination of human focal attention and orienting perceptual functions such as peripheral vision supports the process of knowing where to look when. Orienting perceptual functions provide information about broad patterns in the surrounding environment and pick up changes that might warrant a shift of attention away from the current focus. Woods (1984) suggested that designers of virtual perceptual fields can model their systems on this characteristic of the human perceptual system--a center-surround technique. To support knowing where to look when, designers should surround a highly detailed central view coordinated with lower resolution, i.e., more distilled, views of physically or functionally related data. Lamping et al. (1995) have called this technique "focus plus context." Furnas (1986) proposed supporting navigation in a similar way by using an optical analogy, the fisheye lens.[1] All of these labels have in common the technique of balancing a high resolution detailed view with summary views of related parts of the virtual perceptual field.

The simplest way to develop a surround is to show nearby areas of the artificial data space in less detail. If the artificial data space is organized around discrete chunks, overlapping the chunks can help the viewer integrate the individual pieces into a complete picture of the overall space (as atlases show portions of adjacent regions around the region of interest). For example, the redesigned electronic procedure display system for nuclear power plants, instead of showing a few procedural steps in complete detail, showed one step in detail embedded in the context of previous steps and upcoming steps shown in less detail (Elm and Woods, 1985).

Another way to develop a surrounding context is to use locality maps, that is, a map of the neighborhood surrounding the current display of interest. One question becomes what criteria should be used to define nearby related material. In many designs, developers use the physical topology of the underlying device or process to define nearby material. Users can then pan and zoom to move the viewport across this large topology (or to move the topology relative to the viewport).

More powerful from the point of view of aiding practitioner information extraction and performance is representing functionally inter-related data in the surround. In this technique, the surround represents distant information and displays within the virtual space that are functionally related to the high-resolution view on the screen. This technique is model-based as it presupposes a model of relationships between different types of views and information about the domain in question.

For example, Woods and Hollnagel (1987) used one cognitive task analysis technique to map meaningful contexts for process control domains (see also Mitchell and Saisi, 1987). As a result, they were able to develop an information system concept where the computer system automatically displayed lower resolution views of contextualy related topics or areas (e.g., other related goals, mechanisms and requirements). When a practitioner selects a topic or area of interest for display in the high resolution viewport, he or she has defined a focus of attention. A

---

[1] The fisheye lens, if taken literally, has the drawback that the surround area is optically distorted. The extant working model for a focus plus context technique is the center-surround organization of the visual field that supports the coordination between orienting perceptual functions and our focus of attention.

functional model of the domain plus information on current status allows one to know what other topics are relevant to the primary view of interest. Based on this model, summaries about the status of these contextually relevant topics can appear in the surround automatically. These functionally related summaries in the surround can vary with context, i.e., the type and level of information displayed could depend on the state of the underlying system or on the state of the problem solving process.

The surround should cue the observer to related views within the display space and should provide status information as well. Just like the orienting function of peripheral vision, status information can help users decide when to change their focus of attention. The status summary function of the surround means that (as for longshots) the data displayed should be distilled and abstracted, reveal activities and change, and should support check reading.

Designers have to decide what kind of information at what level of summary and abstraction is appropriate for a surround. What is nearby needs to be thought of as a semantic property of the conceptual space if one is to use the center-surround technique to help observers know where to look when (linking domain semantics to the structure of the virtual perceptual field). In many domains this is complicated by the fact that there are multiple semantic relationships between topics, areas and data.

Recently, there have been a large number of innovations which are based, in part, on a center-surround concept (Mackinlay, Robertson and Card, 1991; Mitta and Grunning, 1993; Robertson and Mackinlay, 1993; Sarkar and Brown, 1994; Rao and Card, 1994; Bartram et al., 1995; Mackinlay, Robertson and Deline, 1996; Greenberg, 1996). These innovations go by varied names depending on their individual history of development but they all depend on the basic concept of providing a higher resolution focused view suurounded by contextually relevant information and views.

Side Effect Views

A specific kind of surround view is side effect views (Woods and Hollnagel, 1987). Side effect views provide users with information about distant areas of the information space that might be affected by actions they are taking or activities going on within their high-resolution central view. The goal is to give users a global picture of the state changes, both main effects as well as side effects, that occur as a result of actions taken.

Designers may want to consider side effect views whenever a domain has multiple interconnections between systems and functions so that actions can have multiple effects--the intended or main effects as well as other unintended "side" effects. In this kind of situation, it is relatively easy for practitioners to err by missing side effects of their plans and activities. Examples of domains with multiple interconnections are common in process control. For example, in nuclear power plants changing net water inflow to the reactor has multiple effects--coolant inventory changes, but reactivity (the nuclear reaction) can change as well because boron is dissolved into the coolant water to act as a moderator of the nuclear reaction. Increasing or decreasing net water inflow affects coolant level and boron concentration.

To reduce errors involving missed side effects, one nuclear control room computer-based display system included a specific side effects window (Woods and Hollnagel, 1987). If an operator is interested in one function (e.g., calls up displays on coolant inventory as the primary focus), the system also provides a summary of the status of other functions which could be influenced by changes in the primary area of interest in a side effects window (e.g., the reactivity/boron concentration function).

These kinds of complex interconnections happen in many other domains as well. For example, extended spreadsheets couple together multiple data tables. Values from one table may be used in calculations in other data tables. If a user of an extended spreadsheet changes the value of a cell on the screen, it may impact calculations and values in other data tables which are not visible because the extended spreadsheet is much larger than the available keyhole.

A side effect view can indicate whether a change in the table on a display causes numbers in distant areas of the extended spreadsheet to change in important ways (Watts, 1994). In this study, side effect views were proposed to explicitly represent the interconnections within and between spreadsheets, and to provide immediate feedback about how changes in one spreadsheet cell affects other cells and data tables. Side effect views can help users notice errors quickly, and can help users find the cause of errors by showing which cells contribute to an erroneous cell. Note that for a side effect view to provide immediate feedback about user actions, the view must be available in parallel with the active area of the spreadsheet. Also, the feedback must provide information that is relevant to user tasks and goals. For example, if users have the goal of keeping specific cells in their spreadsheet positive (i.e. if users do not want the cell representing a project budget total to be less than zero), a side effect view can aid this goal by showing when a formula result goes negative.

Cues to Status

Cues to status are an important technique to enable practitioners to easily "check read" or peripherally pick up what might be interesting changes in other parts of the underlying process that should guide a shift in attentional focus (Woods, 1995b). Many of the techniques discussed involve providing cues to status as well as to structure--content-laden landmarks, center-surround, side effects views.

For cues to status to be effective, users need to be able to "size up at a glance" the status or activities going on in that part of the artificial data space. This means that the cues should meet the criteria for the status summary function:

- a distillation of what is important about that aspect of the field of practice,
- an abstracted representation that goes beyond simply making raw data available,
- representations that support quick check reads and pattern recognition so that practitioners can pick up potentially interesting changes without disrupting their ongoing line of reasoning.

Conceptual Spaces And Workspace Coordination

Workspace coordination is as fundamental a part of the design of a computer-based information system or device as the design of individual displays. This level of analysis and design is important because users can easily get lost in a large network of displays and options hidden behind the narrow keyhole of the computer screen. But practitioners, as responsible agents in a field of practice, cannot afford the cost of getting lost in the interface. Instead, we observe practitioners adapt the system and their behavior to avoid being burdened by navigating in the interface itself when they should be focused on their job. They avoid navigating through too many displays when the job is resource-limited, fast paced, event-driven or critical to their goals.

Simply providing context free aids that increase the efficiency of moving from one display to another is not enough to help practitioners perform coherent tasks across displays. While we are struck by the competency people exhibit in navigating physically distributed, spatially organized information spaces, simply reproducing the superficial properties of a physical space also is not enough to aid practitioners. We do not see all parts of a large physical space at once and in parallel. The competency we exhibit in physical spaces for focusing in on potentially relevant data comes from the coordination of orienting perceptual functions and focal attention. By supporting these mechanisms in a virtual perceptual field designers can break down the keyhole of the computer medium.

While the computer medium fundamentally presents the observer with a keyhole, the computer also provides designers with the power to develop new techniques to enhance practitioners ability to process data across displays as they perform coherent tasks. The problems associated with keyholed computer-based systems have spurred developers to use this power and innovate many different kinds of techniques for workspace coordination.

We think that techniques that contribute to breaking down the keyhole are based on a few basic concepts that should guide developers as they think about how to design a coordinated virtual workspace.

1. Build a conceptual space. Workspace coordination is about deciding what kinds of views should be available to practitioners and how to organize those views into a coherent conceptual space. A conceptual space depicts relationships in a frame of reference.

2. Link structure to semantics. Building a conceptual space is intimately concerned with understanding the nature of practice in some domain of activity:

- what are coherent activities and sequences,
- how are activities interleaved given resources and demands,
- how do activities ebb and flow regularly over different task epochs or irregularly driven by events,
- how do new events and changes require practitioners to shift focus,
- what data needs to be seen in parallel in different contexts?

This is the kind of information that is needed to develop what kinds of views should be available, how to coordinate different views in parallel or serially in pace with the interleaved and changing tasks of practitioners.

3. Provide visible cues to structure. Many of the techniques to aid navigation are methods for making the underlying structure visible to users. But the second principle points out that simply adopting a spatial metaphor may not be enough -- the representation may look like a physical space or we may adopt a fixed layout to pan and zoom across. The visible structure needs to capture the important semantics of the field of practice (e.g., Vora, Helander and Shalin, 1994).

4. Provide cues to status at a glance. A recurring theme in many of the techniques discussed specifically in this chapter is adding cues to status. In order to re-orient attention the observer needs to be able to peripherally pick up some information about what is going on (activities, changes, inter-relationships) in other parts of the artificial data field. The goal is to allow observers who are focused on one view or activity to "check read" or to pick up what might be interesting changes in other parts of the virtual field that should guide a shift in attentional focus.