

Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use

Jonathan Grudin
Microsoft Research
One Microsoft Way
Redmond, WA 98052-6399 USA
+1 703 936 0784
jgrudin@microsoft.com

ABSTRACT

Software today does not help us partition our digital worlds effectively. We must organize them ourselves. This field study of users of multiple monitors examines how people with a lot of display space arrange information. Second monitors are generally used for secondary activities related to principal tasks, for peripheral awareness of information that is not the main focus, and for easy access to resources. A second monitor improves efficiency in ways that are difficult to measure yet can have substantial subjective benefit. The study concludes with illustrations of shortcomings of today's systems and applications: The way we work could be improved at relatively low cost.

Keywords

Awareness, multiple monitors, displays

INTRODUCTION

As more of our activities are supported digitally, we use more applications and devices at work, in the home, and on the move. The seductive goal of getting information to and from "anyone, anywhere, at any time" is reflected in the CHI 2001 theme.

Yet information overload is a major concern: we cannot handle "everything, everywhere, all the time." We need to partition our digital worlds. Software today does not provide much help. With greater understanding of how people work, it could do better.

A single monitor on a computer is like a one-room house or a village with a single general store: everything must go into it. This is the dominant mode of computer use. Much effort has gone into helping us organize this space, with windows, frames, icons, menu bars, task bars, and so forth.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCHI'01, March 31-April 4, 2001, Seattle, WA, USA.
Copyright 2001 ACM 1-58113-327-8/01/0003...\$5.00.

With the spread of networks and the Web, the amount of accessible information is rapidly overloading displays, which have not grown at the same pace. This has contributed to the already substantial efforts to make more efficient use of monitor space, described in the next section.

Additional devices – laptops, home computers, personal digital assistants (PDAs) – can help us selectively route information. But often a goal is to access information from any location – multiple views into the same space. Then, as with a desktop computer, we need to focus on managing information. Just as houses have several rooms and towns have many stores, our digital worlds can benefit from partitioning and from interfaces that support it sensibly.

A natural way to partition is to take advantage of peripheral awareness. Widely-available monitors are not large enough to extend far into our peripheral view, and software has difficulty sensing where on a monitor our attention is focused. Like a one-room house, single monitors do not provide structural support for partitioning space.

Multiple monitors for a desktop computer are now relatively inexpensive. But software supports them poorly, in part because multiple monitor use is relatively new and we do not understand how people make use of them.

Early in the era of graphical user interfaces, Cypher [2] envisioned software environments that support multiple activity structures – virtual "places" that one moves among, each supporting one activity. Henderson and Card's Rooms system [6] was a software instantiation of virtual places. Software that allows switching among multiple screens or desktops on one monitor realized some of these concepts.

Multiple displays are a feature of control rooms that oversee complex operations, such as air traffic control and factory production. Each display is usually devoted to a single view, possibly fed by a different system. Together they serve a single real-time activity. The circumstances of control room activity differ from those of PC use, but they raise related issues of visual and attentional awareness.

Harrison and Dourish [5] "spend some time thinking about how people turn a space into a place" in the context of

shared media spaces, from text-based MUDs and MOOs to video conferencing. They argue that a *space* is an opportunity, whereas a *place* is imbued with meaning. In their case, this is social meaning, rooted in community practice. But individuals also transform space into place for personal purposes.

This paper brings together these threads, suggesting that the increase in available information has amplified our need to partition our digital worlds into different places, and that multiple monitors can be used to do this in a way that conveniently correlates physical spaces with virtual places.

After motivating the field study of multiple monitor use, I describe a range of behaviors and provide suggestions for improving software interfaces to a technology that could spread rapidly.

THE WORLD THROUGH A 17-INCH WINDOW

Most of us view our digital worlds through a monitor. A large monitor a few feet away covers about 10% of our visual field, the area we can focus on by moving only our eyes. It covers about 1% of the area we can see with a quick neck movement.

You would probably find it frustrating to navigate and understand your physical surroundings if your view was restricted to 1% to 10% of its normal range. With a monitor, you view a growing digital world without the eye and neck movements that are central to vision in the physical world. You make little use of peripheral vision.

The size of a standard monitor has increased slowly, not keeping pace with Moore's Law. However, our digital world – the amount of potentially useful information that is accessible from our drives, over intranets, via the Internet, and on the Web – has grown much faster than Moore's Law. Information related to our tasks; information about the status, activities, and work products of team members; organizational information; reports from professional and educational organizations; commercial and leisure opportunities; news, sports, and financial bulletins; messages from friends and family...

Some information we want to see as soon as it is available. Other information we prefer to have accumulated and sent periodically. Some things should not interrupt us, yet ideally be available at a glance when we have time for them. Other information we prefer to get only upon request.

Work proceeds on several fronts to address the challenges that result from these diverse requirements. First I review software efforts, then those that require enhanced hardware.

Enhancements to windows and GUI features. Windows were an early and successful effort to partition virtual space. A window is generally limited by being restricted to one application, but effort has gone into stretching window capabilities. A window can show two objects of the same type, two views of a single object (e.g. an outline alongside a page view), or can spawn a second window. Fisheye

views [3] access more information by emphasizing that which is particularly salient. Windows can be minimized. To avoid being buried, indicators are squeezed around display borders and windows can block "full screen" expansion of other windows. "Alpha-blending" creates translucent windows that float over other windows, simple key sequences bring different windows or desktops into view, and some menus or task bars disappear until hovered over. Research includes 3D interfaces [9], zooming [1], and "stationary scrolling"[10] of entire displays.

Better awareness and notification indicators. This subset of display enhancements is a major focus of CSCW research and product development. It draws on features described above (such as alpha-blending) and includes sounds and pop-ups that remain or fade away. In one such study, Gutwin et al. [4] tested six awareness indicators and identified two that performed best.

Much larger high-resolution monitors. Wall-sized flat-panel displays are on the horizon and can be simulated now, as in the GMD DynaWall composed of three SmartBoards [11]. For some time, this solution will be very expensive.

Head-mounted displays and other immersive environments. Head-mounted displays may achieve a reasonable cost and be useful in some places, but they are inherently single-user and inhibit interaction. Navigating the virtual space they create requires mastering new input devices and adjusting to the mismatch between real and perceived movement.

Ambient information. As demonstrated by Ishii and his colleagues [7], external events of interest can be signaled by digitally controlled environmental lighting, sounds, air motion, and even smells.

Ubiquitous computing, information appliances. Instead of a single computer and display, we can have multiple artifacts with specialized displays and other output devices [8]. This will come at its own pace and price. In a sense, information appliances are an extreme case of multiple monitors. The infrastructure and design to make them practical are major undertakings. In contrast, the infrastructure for multiple monitors is already in place.

Dual or multiple monitors. For a decade, Macintoshes have allowed a person to use two monitors on one machine as a single surface, dragging objects from one to the other. Windows 98 and 2000 also support multiple monitors. Multimonitor use, requiring a second monitor and video card, has a noteworthy advantage in the quest for display real estate: It is available and relatively inexpensive. Two standard monitors cost much less than one very large monitor, and prices are falling.

Dual monitors have drawbacks. They take up more desk space. They are not supported by old operating systems or by all computers – few laptops can easily handle two-monitor configurations, for example. Another problem – or opportunity – is that even recent systems and applications do not make effective use of multiple monitors.



Figure 1. A 2-monitor configuration. Note bezels interrupting the display surface.

One constraint is that displays do not connect seamlessly: a few inches of metal or plastic separate the surfaces even when the monitors are side by side, as in Figure 1. A window opened across two monitors is interrupted. On first consideration this is purely a hindrance, but it provides some benefits. It forces people to think about the positioning and size of screen objects. Without this constraint, a “maximize window” operation or casual decisions often govern placement. Another benefit is that when an object is on one display or the other, the system or application can more easily deduce where a user’s attention is directed and the software can behave more appropriately. These points are returned to in the discussion.

Despite the drawbacks, multiple monitors can address the explosion in interesting, accessible information. They are invariably popular with those who try them. One reviewer of an early version of this paper wrote that she did not have enough desk space for two monitors; within two months, she and others in her lab had acquired a second monitor.

This study addresses the following questions: How do multimonitor users make use of the extra space? Is the second (or third) monitor an extension of their workspace, or do they use it for different purposes? How can software exploit this capability more effectively?

FIELD STUDY BACKGROUND

Dual or multiple monitors (“multimon”) usually refers to a single processor with multiple video cards and monitors in simultaneous use. The system treats the independent displays as a single continuously addressable space, so objects can be dragged from one monitor to another or can straddle monitors.

This was the arrangement at 13 of the 18 sites studied. My focus is on the behavior of people working with interrelated displays, so a few variants were included that also involve simultaneous use of multiple displays.

Method

Participants were recruited primarily through two channels. Microsoft has a “multimon” email distribution list. A query to this list yielded 7 volunteers. Another 7 people already using multiple monitors in their work were recruited from outside the company by recruiters whose job is to find participants for usability tests. Another 4 interviews were obtained opportunistically; two were recruited by word of mouth and two interviews that focused on PDA use were extended to cover the use of multiple displays. The number of interviews was sufficient to reveal patterns as well as diversity in use.

The participants were relatively experienced computer users, as is usually true of leading-edge technology adopters. Inexperienced users may have fewer online activities and could benefit less from dual monitors – until they become more experienced.

Apart from the PDA users, all participants were interviewed in their workplaces. They did not work during the interview, but windows open at the outset usually provided a snapshot of their work. Interviews, which were recorded, were scheduled for an hour and usually ran over. Offices or workspaces were photographed and some screen captures were collected. Participants received a gratuity.

Much of the structured interview consisted of exploring their configuration and reviewing their online activities. The hardware configuration was discussed. Then the distribution of objects and activities on their monitors was explored, beginning with the visible objects (windows, icons, task bar items, and so forth), continuing to cover buried objects, and concluding with activities and software objects not presently visible. They were queried on their use of communication and notification tools, web browsers, and so forth if these topics did not arise spontaneously.

Participants

The nine external participants included four CAD designers, two managers of small businesses who worked primarily from home offices, a 911 operator, and people working in field sales support and product marketing.

The nine in-house participants included four developers, three testers, one usability engineer, and one program manager. Two were actively involved with developing or testing multimonitor use at the time of the interview. All were technically expert, knowledgeable users.

Twelve participants had used multiple displays for one to three years; two had experience of four months or less.

MULTIPLE MONITOR CONFIGURATIONS

Physical layout

Nine of the 18 participants had one processor and two stationary desktop monitors. Three differed only in having a third monitor, one of which had been acquired just prior to my visit. Of these 12, six had monitors that varied in size. Monitor sizes ranged from 15” (for the smaller of a pair only) to 21” with 17” and 19” the most common.

One person had managed to configure a laptop to behave as a second monitor, varying from the previous participants only in that one monitor was mobile and smaller.

The five remaining participants had more radical variations. Two had hand-held computers (PDAs) that were connected and constantly synchronizing with their desktop system when they were in their offices. They could not drag a window from one display to another, but they used the PDA display as an extension of the primary monitor. A home office worker had a second monitor that displayed the same image as the first, but with a different resolution. Often used for graphics, one monitor displayed the big picture, the other showed a subset in greater detail. The usability engineer had two PCs and two keyboards and used portals on each PC to drag objects (but not open windows) back and forth. The 911 operator had three workspaces and two monitors; on one monitor a single keystroke switched in a second display. Thus, she combined the dedicated displays of traditional control rooms with “desktop-switching.” (The CAD designers had system software that provided the ability to change ‘desktops’ *as well as* dual monitors; switching desktops changed both monitor displays.)

Frequently monitors were not placed as close together as possible, sometimes even farther apart than in Figure 2.



Figure 2. A 3-monitor configuration. Displays are not pushed as close together as possible to minimize gaps.

Logical organization matches physical layout

All participants with full-sized monitors arranged them horizontally and configured them in the same logical relationship. (The software permits one to place monitors in any relationship; for example, monitors that are physically side by side can be arranged logically as one above the other, in which case dragging something down to the bottom of one causes it to appear at the top of the monitor next to it.)

OBSERVATIONS

1. A second monitor is not treated as ‘additional space.’

Participants reported rarely straddling a single window across two monitors. One expressed surprise that he did not. This is consistent with the physical gaps that are often left between monitors.

Even CAD designers, who often work with large, complex graphical objects, reported rarely extending a window across two monitors. Instead, they typically maximize a design to fill one monitor entirely, leaving the other monitor free for other uses. One said “AutoCAD is on the right-hand side, that’s my primary (monitor). This is more my miscellaneous, everything, on the left-hand side.” ‘Miscellaneous’ meant desktop icons and shortcuts, email, ICQ, AOL buddy list, a CD player, Groupwise Notify, a calculator, a game, and a small AutoCAD window to examine details not visible in his primary monitor view.

A strong demonstration that multiple monitors can be more about partitioning than about increasing space is provided by the two participants who dock their constantly synchronizing palmtop computers next to their desktop monitors. One keeps his calendar visible on the palmtop, the other keeps email visible. The increase in space provided by the palmtop display is not significant and there is no information on the palmtop that is not available to the desktop computer. The value is in having instant access to a resource in a known location in peripheral vision.

A 21” monitor has twice the display space of a 15” monitor. However, users of a 21” monitor are not forced to consider how to segment their digital world to the extent that dual monitor users are.

When a user has monitors of different sizes, this becomes another consideration in designating specific objects to specific monitors – spreadsheets or other potentially large applications are typically directed to the larger.

2. Distribution of tasks among monitors.

People generally indicated that one monitor was used for their “main” or “primary” task – writing code, working on a CAD design, preparing a document, slide presentation, or graphic image, and so forth. This is the task to which most keyboard and mouse activity is devoted over time. The separation of tasks was quite conscious: several respondents said “I use the second monitor for (other tasks).”

Sometimes the second monitor is used as an equal partner in the primary task, notably when material from one object is being copied into another. But people reported almost always using it for secondary or alternative purposes. These purposes are principally of three kinds (with some overlap):

Secondary tasks in direct support of the primary task.

“With this I can very easily glance and read and see what’s going on.”

– Developer referring to second monitor used for compile and debug tools, with source code on the primary monitor.

A developer or tester working with source code on the primary monitor often devotes the second monitor to windows (at times more than 5) used to identify variables appearing in the code, report on compilation tasks, provide debugging readouts, and so forth. Once dragged to the second monitor, these windows are used more for reading than for interaction.

Some CAD designers have a variation on this practice. Rather than extend a model to cover two monitors, they maximize the drawing space on the primary monitor by dragging toolbars to the other monitor. This includes basic toolbars that they would keep open when using a single monitor, but it also includes advanced toolbars that they would not sacrifice display space for on a single monitor. One had 11 toolbars open when interviewed; he could easily find and invoke advanced functions. Another window open on his secondary monitor provided a second view into the object on his primary monitor; he used it to locate parts that could be “reused” or copied into the primary window.

‘Live’ information channels from the outside.

“They pop up, they go on top, but they don’t take focus so that I’ll keep working but I’ll see that it’s there...”

– Tester, about reminders that appear on secondary monitor.

Communication channels to people and events are often placed on a secondary monitor. These include email; instant messaging, ICQ and Groupwise Notify (linking participants to colleagues, friends, and/or family); and news alerts such as news headlines, specialized news filters, sports bulletins, stock quotes from different external sources, often set to be updated at intervals such as every half hour or twice a day.

Participants frequently mentioned that interrupts from ‘push’ applications are much less disruptive when they are confined to a second monitor, with updates still available at a glance. Software cannot easily deduce where we are working and when we can be interrupted; multiple monitors allows users more control of interrupts by positioning them.

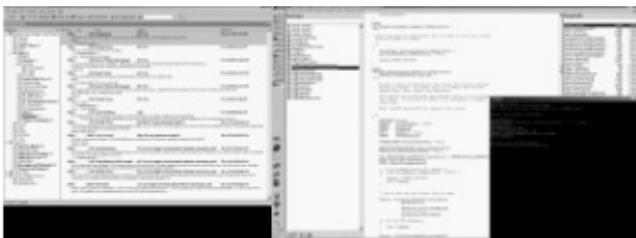


Figure 3. A bitmap from a 2-monitor configuration.

Figure 3 shows one arrangement (screen capture software does not explicitly identify monitor boundaries). The smaller monitor on the left is devoted to email, the monitor on the right holds source code and programming windows.

Using a second monitor as an information channel can be an alternative to or in conjunction with using it to support the primary task. Communication windows can be buried under or tiled with code tracking windows, for example.

One developer who was on numerous project-related distribution lists filtered incoming email into folders. His secondary monitor was used primarily to display his email folder hierarchy. The count of unread messages in each folder is indicated parenthetically in blue next to the folder name. He reported that frequently during the day he glanced over to check these counts for unusual activity, interrupting his work to open email only when he noticed a frenzy of activity. Other users also monitored email folders.

This use supports communication that is important but not critical enough to merit interrupting the principal activity. Burying or minimizing these windows on a single monitor could have a similar effect, but now they can view the windows with a glance, without lifting their hands. When communication does become the primary task, some participants said that they move their email window to their primary monitor to “do email.” Others shifted their attention to the monitor on which the communication application resided and worked there.

Personal resources

“If I ever watch media events, NetShow, I’ll usually have it in the far right-hand corner (of the monitor on the right), because to me this is the farthest stretches of my desktop.”

Common uses of the secondary monitor were for background resources such as To Do lists, contact lists, calendars, and audio controls for playing digital music.

Web browsers were frequently exiled to the secondary monitor. Depending on the browser contents, it could fall into any of the categories listed above in this section.

3. When a button click is too much.

“It’s not intrusive in the sense that I have to click a button to see my calendar.”

– Field sales support engineer

“Alt-tabbing around, I got sick and tired of it.”

– Developer explaining why he had added a third monitor.

Given the ease of minimizing and restoring windows, why bother with a second monitor? Repeatedly, people indicated that they considered it a relief not to have to use buttons, “escaping from the need to Alt-Tab.” The ability to avoid a few keystrokes is welcomed with great subjective enthusiasm, although it might be difficult to objectively measure an efficiency gain. Several factors are implicated.

One is that sometimes it *does* require an effort to get to a buried window and back. Some participants noted that with many windows minimized, “Alt-Tabbing” to get someplace and back is tedious.

Finding features buried in menus is even worse. Opening toolbars through the menu system to find a rarely-used advanced function is much more time-consuming than visually scanning the full set of functions on a second monitor, and the latter approach requires only a single mouse movement and click to invoke the action.

When a second monitor displays processes that support the primary task, it reduces the cognitive load by allowing rapid glances to check on information; for example between a bird's-eye view of an entire graphical object and a close-up of a detail. One respondent explained "I can just read here and turn my head and go to the line, as opposed to reading the number, Alt-tabbing and then maybe forgetting the number or getting distracted."

When accessing communications from outside, there is convenience and an "out of sight, out of mind" element. A glance at a second monitor to see who has sent email is less disruptive than moving a mouse to open and then close an email window. In contrast, windows that are minimized or buried are easily neglected. Informal awareness systems that rely on desktop video, for example, suffer when the large windows that they require are buried and forgotten. Efforts to address this through notification sounds or animations can be effective but distracting.

The significance of the ability to cleanly shift attention was underlined by this observation: Although people did not often straddle windows across monitors, they usually avoided cluttering the right side of the left monitor or left side of the right, making an effort to keep icons, menu bars, audio controls, and other small objects on the left of the left monitor or the right side of the right, far in the periphery. In this way, there is little distraction when they glance between the principal windows on each monitor.

4. People like multiple monitors.

"I would not go back." "(Without it) I feel very limited."
"It has spoiled me." "A single-monitor system with Alt-Tab, it's not as easy for me to go back to a single-monitor system."

These comments reflect the overwhelming consensus of those interviewed. Some were considering adding a second monitor to their home computers and one had done so. A CAD designer who was initially skeptical and the last in his shop to acquire a second monitor is now enthusiastic and "can't imagine" not using it: "You end up with a lot of just, minimization... a lot of extra work."

What is required to provide dual monitors? Functioning, used monitors are inexpensive. Many organizations have unused 15" monitors in abundance. The principal requirement is a second video card.

5. Applications do not make good use of multiple monitors.

"You intuitively think a window should show up... it shows up on a different monitor."

"Most apps will start on the primary monitor even if they were located on another monitor."

One participant described his ultimate in frustration. During breaks he plays a computer game in which one amasses points by executing a variety of maneuvers to commit mayhem. To access unfamiliar advanced features, he must pause the game, bring up the Help display, then close Help

and resume the game. As he does this, his second monitor is blank. Games, he noted, should detect a second monitor and make Help and resource status constantly visible there.

Multimonitor development has focused on getting the system display software and the application program interface to work, with little attention to the human computer interface or intelligent exploitation by system or applications. Even a multimonitor installation program was careless: Clicking on a menu item on one monitor caused a dialogue box to appear on the far corner of the other monitor, where it easily went unnoticed.

Multiple monitor display space is handled by a single coordinate space. Many applications cannot handle the negative numbers that are assigned to a monitor to the left of the primary monitor. "This is a demonstration of what most apps do wrong, how they're not multimon aware," explained one developer, demonstrating an object that popped up a dialogue box far from the focus of attention.

One user prefers to keep desktop icons on the far right of his rightmost monitor, but when he invokes the 'arrange icons' utility, it moves them to the left. So he must move them back each time.

Systems can support multimonitor use better and serve as an example, but applications will make the real difference. Most are given no thought as to how they appear on dual monitors. Dialogue boxes and other items are placed randomly or inconsistently. There is no memory for where interface elements were placed when an application is reopened. Several informants laboriously and regularly rearrange multiple windows or drag toolbars into place.

Progress has been slow, primarily focusing on avoiding the most obvious clumsiness. The next step is to exploit multiple monitors positively.

The PowerPoint application illustrates both the potential for intelligent use of multimon and the need to better understand its use. If one edits slides on the secondary display (as defined by the system during installation) and opens 'Slide Show View,' the slide show appears on the primary monitor. One can continue editing and see the changes reflected in the slide show, or advance the slides.

This terrific capability even pre-dates system support for dual monitors. Unfortunately, it is not discovered by some (perhaps most) multimonitor users because they edit slides on the primary monitor. When the Slide Show is opened in that situation it overlays the normal edit view, leaving the other monitor untouched. This great feature is rendered less effective by not being matched to the pattern of use.

I was told of a CAD programming application with a similar nice: one can program on one monitor and see what is happening in AutoCAD on the other.

"Thank God they finally made it multimon aware."

– Developer describing an internally developed utility.

THE BENEFITS OF ARBITRARY DIVISION OF SPACE

Most users initially consider a second monitor to be inferior to a twice-as-large monitor. I began using dual monitors myself only because of the substantial cost difference. In retrospect, it seems clear that a single large space is not always preferable. Being able to park objects out in the periphery is an advantage.

Consider, as an analogy, space in a house. People generally value large rooms – and they value more rooms. A house with one large bedroom is not the same as a house with two bedrooms of moderate size. In the two-bedroom house, the second room is used for different purposes – perhaps as a guest room and office. One *could* use the master bedroom for these purposes, but we usually don't, even if it is large. The wall makes a difference.

Similarly, with multiple monitors, the division into two spaces can facilitate diversity in use. One *could* extend a single window to cover two or three monitors, but people usually don't.

Consider the difference between a single large meeting room and several small breakout rooms. Each arrangement can be valuable, but different tasks are optimal in one or the other. As with monitors, rooms of equal size may be arbitrarily assigned different functions, whereas rooms of different size may be assigned functions in part based on the requirements of different tasks.

Very large displays will find significant uses when they can be afforded. These are likely to include shared physical spaces where large, manually updated status boards or heavily used bulletin boards are found today. History suggests that further growth in display size will also be welcome on the desktop, particularly as flat-panel displays come down in price.

Yet a place will remain for the arbitrary division of space provided by multiple monitors. Many people prefer a house with several rooms – accepting the arbitrary constraints – rather than one large room in which things can be freely arranged. The users with their miniscule palmtop displays running alongside their large PC monitors made this point: The benefit was not the small increment of space, it was the partition. Space with a dedicated purpose, always accessible with a glance.

OPPORTUNITIES FOR DESIGN

“It would be really neat if they could realize that maybe the Help should show up on another monitor.”

– Participant discussing an office software application

“It would be a dawdle, just give me a spec.”

– A developer (not one of the 18 interview participants) asked the difficulty of coding application enhancements for dual monitor users

Doubling (or tripling) display space will not solve our information management problem, but it will help. As noted above, not all information is equal – we want some things

brought to our attention, others we prefer to keep in the periphery. Multiple monitors do this naturally, segmenting display space in a manner that is clear to the human and also to the software... if the software takes notice.

Although most multimonitor users dedicate one monitor to the focal task and use the second for peripheral awareness and less frequently used resources, they get little assistance from applications in arranging things. More often than not, applications make awkward use of the capability.

Consider the CAD designers who open toolbar menus and drag them to the second monitor for easy reference and access. At a bare minimum, the system should remember placement so people do not have to repeatedly organize themselves. It would be even better for the application to propose a logical arrangement of all toolbar functions on the second monitor.

Every designer should consider the multimonitor scenario, however briefly, during the design process. When two design alternatives are roughly comparable, one may lend itself to more graceful use with a second monitor. Opportunities to exploit a second monitor are often evident, as in making status information visible for fast-paced games. Help windows could in general open on a second monitor. Almost any application has some information that would usefully be placed in the periphery. The developer quoted above was commenting that given the current system support for multimonitor, it is relatively easy to code a multiple monitor solution that is clearly specified.

Of course there is a down side. If our primary applications colonize the second monitor, we reduce its usefulness for notification and awareness of external events, or for easily accessing personal resources. Thought must be given to the division of features between multiple monitors. If applications cannot resist using all available space, we may need three monitors.

The opportunity exists to design suites of awareness and notification features that draw on the full range of communication and agent software that a person selects. Space could be segmented according to events originating with team members, organizational changes, news and alerts arriving over the Internet, and communications from friends and family.

Why isn't more attention paid to dual monitor use? Multimonitor use has been too rare to be a priority for most developers. But as some people begin to exploit the capability and as monitor costs continue to drop, use could snowball. When they see benefits, more people will obtain a second monitor, making it more worthwhile to design for multiple monitors, leading in turn to yet more adoption, and so on... Intelligent design can get ahead of this curve and increase the acceleration along it.

CONCLUSION

In hindsight, these findings may seem to be “common sense,” but relatively few people demonstrate awareness of

them. For example, one very large corporation recently had a “Single Pane Of Glass” campaign to limit employees to one computer and one monitor per office. The resulting mountain of extra monitors were offered for sale at \$25 apiece. Also, as I have noted, almost no designers, developers, testers, or usability engineers (including some who use multiple monitors themselves) consider multiple monitor scenarios for the applications they are building.

This could change.

“The other parts of the company are seeing the CAD department and how productive they are with the dual monitor thing... The more times you have to flip, and flip from one screen to the next or open and close sessions, you lose your train of thought... and they're starting to see that, and then they're having a change of attitude.”

ACKNOWLEDGMENTS

Thanks to Saul Greenberg, Gayna Williams, Nelle Steele, Shari Schneider, Anoop Gupta, Scott LeeTiernan, JJ Cadiz, Leysia Palen, and Kevin Schofield for suggestions and comments. Thanks also to Nancy Bell and to Victoria Grady and Jeffrey Carr of the Usability Support group for identifying and recruiting participants, and to the participants for their time and attention.

REFERENCES

1. Bederson, B.B. and Hollan, J.D. (1994). Pad++: A zooming graphical interface for exploring alternate interface physics. *Proc. UIST'94*, 17-18.
2. Cypher, A. (1986). The structure of users' activities. In D. Norman and S. Draper (Eds.), *User Centered System Design* (pp. 243-263). Erlbaum.
3. Furnas, G. (1986). Generalized fisheye views. *Proc. CHI'86*, 16-23.
4. Gutwin, C., Roseman, M. and Greenberg, S. (1996). A usability study of awareness widgets in a shared workspace groupware system. *Proc. CSCW'96*, 258-267.
5. Harrison, S. and Dourish, P. (1996). Re-Placing space: The roles of place and space in collaborative systems. *Proc. CSCW'96*, 67-76.
6. Henderson, D.A. and Card, S. (1986). Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, 5, 3, 211-243.
7. Ishii, H. & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *Proc. CHI'97*, 234-241.
8. Norman, D.A. (1998). *Invisible Computer*. MIT Press.
9. Robertson, G.G., Czerwinski, M., Larson, K., Robbins, D.C., Thiel, D. and van Dantzich, M. (1998). Data Mountain: Using spatial memory for document management. *Proc. UIST'98*, 153-162.
10. Smith, R.B. and Taivalsaari, A. (1999). Generalized and stationery Scrolling. *Proc. UIST'99*, 1-9.
11. Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P. & Steinmetz, R. (1999). i-LAND: An interactive landscape for creativity and innovation. *Proc. CHI'99*, 120-127.