Costs, Clocks, and Habits: Design Issues for Elegant Peripheral Awareness Displays

Carl Gutwin
Department of Computer Science, University of Saskatchewan
Saskatoon, Saskatchewan, Canada S7N 5A9
carl.gutwin@usask.ca

INTRODUCTION

A number of peripheral awareness displays have been invented for the computer desktop (e.g. IM clients, Tickertape [2], Sideshow [1]). These displays often provide valuable information, but they still have something of an uneasy relationship with the desktop. Users are often reluctant to give up screen space that they perceive as too valuable for mere awareness (e.g. [1]), and as a result, the awareness information is often difficult to retrieve. In this short paper, I consider the idea of information costs in the design of peripheral awareness displays. Information cost appears to be a fairly straightforward design factor, but the example of a simple wall clock shows that its relationship to peripheral awareness is subtle. Using the example of the clock, I consider how other factors interact with information costs to produce both elegance and peripherality.

INFORMATION COST

Awareness is (at least partly) about information which has to be gathered, and that gathering can take greater or lesser effort depending on the representation and location of the information. I think that the notion of information cost is an important one if we want to get these designs right – so, rather than thinking about displays being in peripheral vision (or peripheral whatever), what we may want is information sources that have reasonable costs. The cost can be split up into two categories: gather cost (the effort required to obtain the information) and interpret cost (once you have the information, how much effort is required to integrate it with your existing knowledge). Figure 1 maps these two costs onto Neisser's perception-action cycle [4], which is a good model for the maintenance of awareness (e.g. [3]).

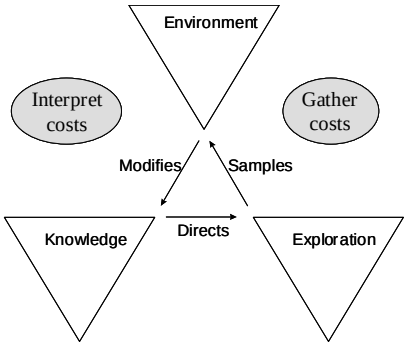

Figure 1. Information costs in the perception-action cycle

Both of these types can be broken into (arbitrary) levels. For example, I can think of five general levels of gather cost, based on the idea that if you're paying attention to something then the cost to gather information about it is basically zero, and if you have to go looking for it, the costs are high. So, the information could:

1. occupy your full attention (lowest cost)
2. occupy part of your attention
3. be out of focus but visible
4. be hidden but nearby
5. be in an unfamiliar location (highest cost)

It's obvious from this list that although we may want to reduce information costs in peripheral awareness displays, we don't want to reduce them too much because they will start to intrude on focused attention. So, there's a reasonable region that might range from "something more than 2" to "something less than 5." (We might keep the full scale around, however, since people do move awareness things into focus and I can imagine that displays could shift up and down).

For interpretation costs, there are costs that I would call decoding costs, and then integration costs. For decoding, the only effort levels that I can think of are perceptual vs. cognitive, although there are likely more. On this scale, perceptual representations are better (and get higher elegance marks), although some kinds of information are not well suited to this approach. For integrating the decoded information into what you already know, the issue is whether the information in is a form that is compatible with what you've already stored away. Rather than levels here, this forms a continuum from perfectly compatible to completely incompatible. This might argue for customizability, so that if the user wants to think about some piece of information as a pie chart, they can see it as a pie chart.

So, as an example, how does a 'network traffic bubbler' ambient display fit into these scales? It's quite far out of attentional focus (probably sitting in the corner somewhere), so if you are sitting at your workstation, the costs of gathering information from the bubbler are actually fairly high. This points out, however, a key thing about displays situated in an environment: that because people are mobile, and because they tend to look where they're going, you can place a display in a likely location and actually get it into focused (or at least shared) attention.

So, gather costs can become low, but update frequency is then based on the viewer's mobility.

In terms of interpretation costs, the bubbler is a perceptual and qualitative display, so it will be easy to get the information (at least, some/lots more/less type of information) without having to think too hard about it. The bubble-as-traffic metaphor would be compatible with mental models that are similar (e.g. network as pipe and flow); but perhaps there are other models that would make it more difficult to interpret the meaning of the display.

This also raises the issue of whether the display answers the questions that people want to answer with the information, but this is just the same old representation problem that infovis always has to deal with. An interesting question raised by the representation issue, however: perhaps peripheral awareness displays should ONLY attempt to fill information needs that are amenable to perceptual ambient displays, where you get fuzzy general state rather than exact state.

In the next section I'll look at another example, this one from the everyday world, and see how information costs play out with awareness displays that everyone knows about.

INFORMATION COSTS AND ELEGANCE IN CLOCKS

There are a lot of good examples of elegant peripheral awareness displays in the real world. I'd like to look more closely at clocks, since I think that they provide about the same amount and kind of information that might be in a computational awareness display, and because we might get some agreement that they are also a good example of elegance.

Clocks are interesting in terms of information cost, because even though they are about as canonical an example of a peripheral awareness display as you can get (at least with person-made displays), their information costs are higher than might be suggested by their ubiquity.

Clocks are never located in the normal focus of attention, and are only rarely even in the actual periphery. Wristwatches are perhaps closest, but even these are not visible without several motions. If you don't wear a watch, gather costs for clocks go from the non-trivial (turning around to look at the wall) to the substantial (walking to another room that has a clock). However, the costs for clocks rarely go up to number 5 on the scale above, because we generally know where the clocks are. At the other end of the scale, clocks do occasionally intrude into attention (alarms, hourly chimes).

In terms of interpretation costs, clocks are again not at the low-effort end of the scale. Although analog clocks allow perceptual processing, there is still a strong cognitive component to them; plus, they are both highly abstract and complex in terms of representation (consider how long it takes kids to learn to tell time). Of course, these difficulties are long past, and so for most people analog clocks don't present much of a decoding cost. Digital clocks of course ignore all of this perceptual learning and require cognitive effort. Also, clocks aren't particularly ambient, in that you can't get a general sense of the time just by being nearby; you actually have to attend to them to get much out of them (although there are lots of ambient time displays, such as shadows on the floor).

I would argue that information costs for clocks are reasonable, but not particularly low. So why do clocks work so well as peripheral awareness displays?

I think there are two main reasons: one is that effort within a certain range isn't that important after all, and the second is that clocks don't move around. I'll begin with the latter, since the first reason kind of falls out of the second.

Clocks don't move (I mean the clock itself, not the hands), and the fact that they don't move makes it easy for people to remember where they are. Furthermore, we don't put things in front of clocks, so knowing where a clock is also means being able to use it just by looking in the right direction.

Why are these important? Stable location and trivial visibility are two crucial requirements for forming a habit: if a display is always there and always visible, it is easy to start looking at it out of habit. As a result, we have become a society of clock watchers. There is an old trick about asking someone the time just after they've looked at their watch – chances are they won't know because they will have looked purely out of habit. This may be good or bad, but the important thing for awareness displays is that it manages to occur at all.

Once you have a habit, then the effort starts to become less important – and as long as it stays within the range of what can be done relatively easily and automatically, then people will likely not even notice. With less perceived information costs, the display becomes less obvious and (maybe) more elegant.

I think that this is a big issue for computational awareness displays. If a display is consistently in the same place and consistently visible, then I would be willing to bet that (a) it fades into the background (background is elegant), and (b) people start to look at it without thinking about it (automatic processing is elegant).

Notice though that approximately zero computational awareness displays are fixed in location and always visible. If you don't agree, think of your favourite awareness display: where is it now? is it always in the same place? is it always visible? do you ever open a document in front of it? does a screen saver ever blank the screen? do you ever shut the machine off? In short, it's not the same as a clock.

MAKING AWARENESS DISPLAYS MORE LIKE CLOCKS

So, what would it take to make a computational awareness display as elegant as a clock? Assuming that you've already done a good job of providing useful information and that your data representation is appropriate (i.e. do good visualization design), I think there are three things that need to be done:

1. *Situate the display in the environment*. As discussed above, the display should be fixed to a location. However, part of our ability to remember where a clock is located relates to the differentiability and consistency of the environment surrounding the clock. So, the display should be in the same place and also in recognizable surroundings. This implies that displays should not be just windows on monitors, they should be dedicated displays that sit outside the context of a desktop workstation.
2. *Never move it, change the source, or turn it off*. These 'nevers' imply that the display has to stay where it is forever, not just for a week or for a project; it has to show the same source forever (no borrowing the display to look up movie listings); and it has to be robust enough to run all the time (forever). Are we ready to commit enough resources to making our awareness displays like this? Do we really care that much about traffic conditions?
3. *No more than medium information cost*. People can handle medium costs, but not much beyond that. The parts of #2 above follow from this, but there are other things implied as well: the display should be in an accessible location and should be easy to see without any interaction required (e.g. no click to view – looking is low-effort but manipulating is high-effort).

There are other advantages to making clocks. There are lots of clocks that are pleasant to look at (aesthetics are elegant), and even sound nice. Clocks even provide a notification function through their alarms, so there is a possibility that the awareness display could do a bit more than just give out the information (I'm not sure whether the idea of a clock means that displays could have unlimited functionality, though).

WHERE I'M GOING WITH THIS: PROJECTWATCHER

ProjectWatcher is a system that shows the state of, and changes to, a CVS repository (see Figure 2). ProjectWatcher combines information about the files in a repository with information about the people working on the project. It can be used as a peripheral awareness display that shows the state of the project or changes to the repository. It can also be used to answer awareness related questions about other people (such as who is working on what, or who has worked on a particular file recently).

Each file in the repository is represented as a simple block, and blocks are ordered by the date of the file's creation. Blocks are coloured according to which developer on the project last worked on the file. Outline highlights are added to blocks in situations where the file is different between the local version and the repository.

At present Projectwatcher is a normal window-based application that sits at the edge of the computer screen. However, starting soon there will be a clock-version of Projectwatcher that will sit on my desk and keep me up to date on one and exactly one ongoing software project (a reimplementation of Spacewar, for what it's worth). It will be built using a PocketPC glued into a nice wooden frame (this makes for an expensive clock, but hopefully costs for building dedicated displays will come down eventually).

Figure 2. Mockup of the clock version of Projectwatcher

CONCLUSION

A few points from all this might be worth remembering. First, information cost is a useful (but not always telling) design factor in peripheral displays. Second, displays don't have to be all that close by to support peripheral awareness – people are willing to look around if it's a habit, and there are other factors that make moderate information costs bearable. Third, the example of the clock is a possible way forward for elegant awareness displays – situated, fixed location, visible, consistent, always there. Finally, going this far towards turning an information source into an everyday object highlights the question of whether the information is valuable enough in the first place to go to all that trouble, and raises the question of whether we can learn from the clock example even within a desktop environment.

SOFTWARE AVAILABILITY

ProjectWatcher (the desktop version) is available from the Saskatchewan Interaction Lab web site: hci.usask.ca.

ACKNOWLEDGMENTS

This research is supported by IBM and NSERC.

REFERENCES

1. JJ Cadiz, Gina Danielle Venolia, Gavin Jancke, and Anoop Gupta, Designing and Deploying an Information Awareness Interface, *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work.*
2. Geraldine Fitzpatrick, Simon Kaplan, Tim Mansfield, David Arnold, Bill Segall. Supporting Public Availability and Accessibility with Elvin: Experiences and Reflections. Journal of Computer Supported Cooperative Work, 2002 v.11 n.3/4 p.447-474.
3. Carl Gutwin and Saul Greenberg, A Descriptive Framework of Workspace Awareness for Real-Time Groupware, *Journal of Computer Supported Cooperative Work*, 2002 v.11 n.3/4 p.411-446.
4. Neisser, U., *Cognition and Reality*, W.H. Freeman, San Fransisco, 1976.