The Designer's Outpost:
A Task-Centered Tangible Interface for Web Site Information Design

Scott Klemmer and Mark Newman
December 8, 1999

> "To enter [Jackson] Pollock's studio is to enter another world, a place where the intensity of the artist's mind and feelings are given full play."
> - Robert Coodnough, *Art News*, May 1951

Abstract

The Designer's Outpost is a tangible user interface that combines the affordances of paper and large physical workspaces with the advantages of electronic media to support information design for the web. Based on an earlier ethnographic study, we have analyzed web site design practice and developed a system to support the practices used by designers during the early phases of information design.

The Outpost system is implemented using a digital desk: a large (33"x27") display screen with the form factor and slope of an architect's drafting table. Designers interact with the system by writing on Post-It Notes and arranging them on the desk in related groups. The system tracks the designers' manipulation of Post-It Notes through digital cameras and displays feedback about the structures that it infers. Designers can also use a special pen directly on the desk surface to perform other kinds of interactions with the system, such as linking chunks of information and specifying group labels.

Web Site Design Practice

In earlier work, one of the authors conducted an ethnographic study into web design. The purpose of the study was to learn about web site design in order to develop systems to better support design practice. The study consisted of interviews with eleven professional web site designers from five different companies. All of the designers except one were interviewed in their work environments. These designers represented a range of different experience levels, backgrounds, and primary responsibilities with respect to web design within their own firms. Each interview consisted of asking the designer to choose a recent project that was completed or nearly completed, and walk the interviewer through the entire project, explaining what happened at each phase. The designer was asked to show examples of artifacts that he or she produced during each phase and explain the meaning of the document with respect to the process as a whole. At the end of some of the interviews, the designer was asked to give copies of the documents discussed during the interview to the interviewer for the interviewer's reference. In this way, examples of design process artifacts were collected from four designers.

Three important observations were made during the course of this study. First, designers create many different representations of a web site during the process of designing a web site. Second the production and use of these *intermediate artifacts* dominate the day-to-day work practice during most of the design process. Third, it was learned that web design is made up of several sub-specialties, including information architecture and graphic design, each of which has its own tools, products, and concerns.

Designers create many different representations of a web site during the process of designing a web site. This observation is similar to what was reported by Sumner and Stolze in their studies of speech application designers[Sumner97]. Examples of pervasive and significant intermediate artifacts include sitemaps, storyboards, schematics, and mockups. These representations depict the web site at varying levels of detail, from site maps, which depict sites as related blocks of labeled information to mockups which depict individual pages in high-fidelity, specifying exact colors, typefaces, layout elements, and graphics. Brief descriptions of these common representations are given in Figures 1-4 below.

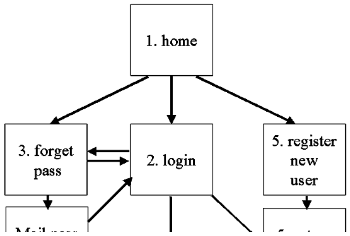

Figure 1: Sitemaps are high-level representations of the entire web site. Pages or sub-sites are represented as labeled blocks of information.

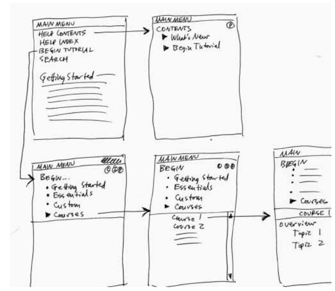

Figure 2: Storyboards are representations of interaction sequences. Minimal page-level detail is shown.

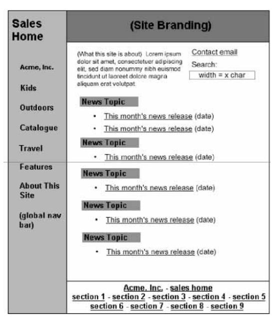

Figure 3: Schematics show individual pages at a medium level of detail. They depict the *information* on a particular page, but do not specify in detail how that information should be presented.

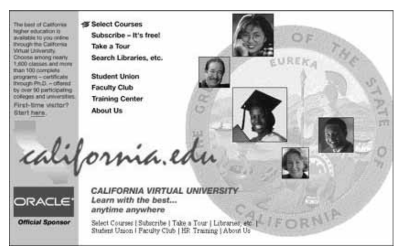

Figure 4: Mockups are high-fidelity representations of individual web pages. They depict *exactly* how the page should appear, including detail about graphics, colors, and typefaces.

These representations, and the intermediate artifacts that embody them, are extremely important to the design process. They provide structure to the design process by articulating a set of short-term goals, the satisfaction of which drives day-to-day design work. These artifacts also structure the communication that takes place throughout the design process. They are used to communicate with clients, developers (the people who will realize the design as a finished web site), and with other designers. They are also used by designers to improve their own understanding and ideas, in other words, to communicate with themselves.

Finally, a division of labor exists within the field of web design, and not just along the more obvious lines of designer-client, designer-developer or designer-project manager. Within the scope of "design" itself, there exist roles such as "information architect" and "graphic designer." In some firms, these roles are divided among multiple practitioners, whereas designers at other firms take on these roles at different junctures in the course of a design project. Even in the latter case, where individual designers take on multiple roles throughout a project, the existence of different roles with different concerns is consciously discussed, and it is clear to these designers (at least in retrospect) when they are performing the different roles.

One of our general research interests is in supporting informal interactions such as those often employed by designers during early phases of design, when the focus is on high-level structure and the space of design possibilities is being explored. In keeping with this interest, we have focused our attention on the artifacts and work roles that focus on high-level representations. These artifacts and roles were observed to be most central to design activity during the early phases of the design process. As was mentioned earlier, the purpose of the investigation into web site design was to guide the development of systems to better support web design. One such system, called DENIM, has already been built and is described in [Lin00]. In the remainder of this paper, we describe a system called Outpost that is currently under development. Outpost also supports web site design by focusing on specific practices employed by certain designers when performing early-phase information design.

Information Design Using Post-It Notes

Outpost is motivated by one specific design practice that was observed during the ethnographic study discussed above. This practice consists of arranging Post-It Notes on a large surface such as a wall, table, or desk in order to explore the information structure of a web site. Designers write chunks of information on Post-It Notes and stick them to the wall. They then move the notes into spatially proximate groups representing categories of related information. Groups thus constructed are labeled and further grouped into hierarchies of labeled group. The structure of hierarchical groups is then used as a baseline for the structure of the web site.

A version of this practice is described by Beyer and Holtzblatt in their book Contextual Design [Beyer98]. They call the technique "Affinity Diagramming" and specify certain details such as the colors that group labels should be (blue for first-level groups, pink for second-level groups, and green for the highest level groups, also called "areas of concern"). Jakob Nielsen also advocates a version of this method using index cards to design hierarchy [Nielsen99].

The large workspace used by this work practice provides several clear advantages with respect to the task. For one, the large space permits the representation of large, complex information space without the loss of contextual, peripheral information that would accompany a mapping the information space onto a conventional graphical user interface. Such an interface would likely map the information space to a virtual space with some small portion visible on the monitor at any given time. The large physical space also provides an immersive environment in which the designer can move around the information space. Collaboration is aided both by the persistence of the artifact, which supports asynchronous collaboration and constant awareness of the state of the project, as well as by the greater-than-human-sized space which allows multiple people to simultaneously view, discuss, and modify the artifact. Finally, the creation and exploration of the information models uses cheap, readily available materials (Post-It Notes, walls, pens) and requires very little investment in time and effort to make use of the tools.

By digitally enhancing a physical workspace in which designers can carry out the information design practice described above, we hope to overcome some of the drawbacks of the strictly paper-and-pen approach. Electronic media are easy to replicate and distribute. The immersive, persistent aspect of the information representation is advantageous to design team members who are fortunate enough to be collocated with each other and with the space where the representation is maintained. Those who are distant from the space, whether permanently or temporarily, have no access to the representation. A computerized representation is much easier to export for viewing at remote locations.

Using an enhanced display as the background surface for the information space representation allows the possibility of additional feedback and interactivity with the representations. Making the surface itself interactive allows the possibility of additional expressive capabilities, such as the ability to draw or make meaningful gestures directly on the surface. Finally, space is at a premium in many organizations, and it is often not possible to permanently dedicate a significant amount of physical space to a specific project. Capturing information space representations into a digital format allows temporary configurations to be made persistent and remain interactive after the physical space is no longer available to maintain the representation. Furthermore, this persistence can be achieved with no need for someone to manually translate the physical representation to a digital one.

Figure 5: A designer sitting in front of a Post-It Note covered wall. Post-It Notes represent chunks of information and are arranged spatially into groups of related information. (from *Contextual Design*, Holtzblatt and Beyer)

The Outpost Solution

The goal of the Designer’s Outpost is to provide a tangible user interface that combines the affordances of paper and large physical workspaces with the advantages of electronic media to support information design for the web. In particular, it is designed to support the practice described in the previous section.

The current implementation of the Designer's Outpost is designed to run on a digitally enhanced drafting table-like surface (referred to as the "desk" from now on). The user interacts with the system by writing on Post-It Notes and arranging them on the desk in related groups. The user can also use a special pen directly on the desk surface to perform other kinds of interactions with the system, such as linking chunks and specifying group labels.

The basic operations supported by Outpost are writing, grouping, labeling, and linking. Creating hierarchical groups is also supported. Ink is written on Post-It Notes using a Cross iPen tablet and a compatible inking Cross pen. The Post-It Notes are then transferred to the desk surface where they can be grouped. The system assumes that information chunks (represented as Post-Its) that are near each other are meant to be in the same group. Visual feedback is provided to indicate the system's interpretation of the group boundaries and to enhance the sense of group cohesion.

One Post-It Note from a group can be selected as the label for that group. We intend to use the desk pen for specifying the group label, either through a gesture made on or near the target Post-It Note, or through a contextual menu that can be activated by the pen in the vicinity of the target. Other approaches that we considered included special marks made on the Post-It Note and the use of different Post-It Notes for labels. Both of these solutions were discarded, however, because we did not wish to constrain the user to have to decide before adding a Post-It to the desk whether or not that Post-It should serve as a group label. Labels are displayed by the system with distinguishing highlights and borders.

UI ENVISIONMENT

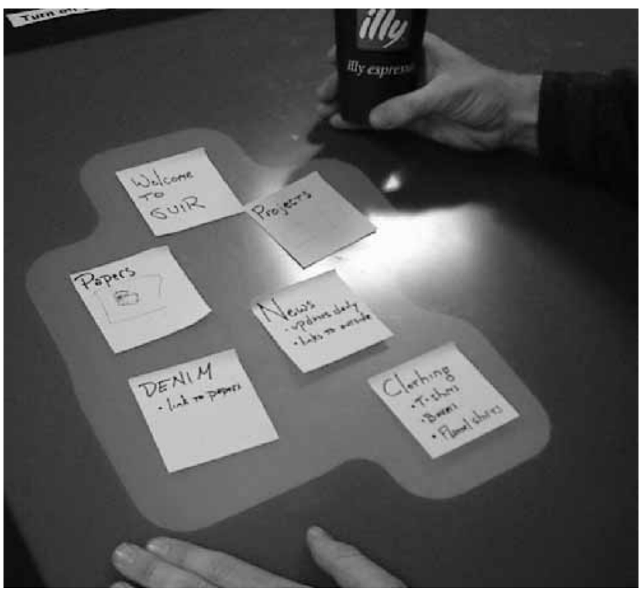

Figure 6: Groups are formed by placing Post-It notes near one another. The system provides feedback about the groups it has recognized by outlining the group.

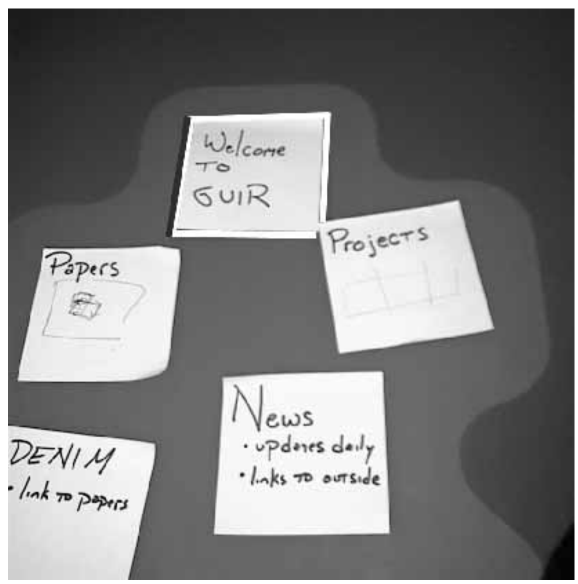

Figure 7: Groups may have labels. Labels are highlighted by the system using a clearly distinguished border color.

The ability to link chunks and groups of information is an feature provided by Outpost that is not readily provided by the existing practice of grouping Post-Its on a wall. Some variants of the existing practice do exist that allow links between different pieces of information to be indicated. Such variants place the Post-Its on large sheets of paper instead of directly on the wall, or they place the

Post-Its on a whiteboard or other erasable surface. When links are drawn on the wall, however, maintaining links between items that may be moving around is tedious if not impossible. Outpost provides links as a way to enrich the vocabulary available to designers for describing the relationships between objects in the structure. Groups express only one kind of relationship between chunks of information, namely *similarity* along some unspecified dimension. Links can be used to impose additional structure on the information by describing other relationships, such as "reachability" (you can get there from here) or information flow (information produced here is used to generate the information there).

Groups can be further grouped into larger groups. This multi-level grouping allows the construction of hierarchies. In order to explore different arrangements at a level above a simple group, is it probably infeasible to use proximity to specify second- and third-level groups. Using proximity would require that entire groups be moved in order to reassign them to new higher-level groups. Since this would involve moving potentially dozens of Post-It Notes, we do not think that proximity should be used for creating higher level groups. We intend to use desk pen operations, possibly in conjunction with links, to specify higher-level groups.

Finally, hierarchies created with Outpost will be available for import into DENIM to serve as baseline site maps for further exploration and refinement of the site. Chunks of information will be regarded as DENIM web pages, and the Outpost structures will be represented as hierarchical site maps in DENIM, with labels appearing in a higher level of the hierarchy than the group members and all group members linked (accessible) from the page representing the group label.

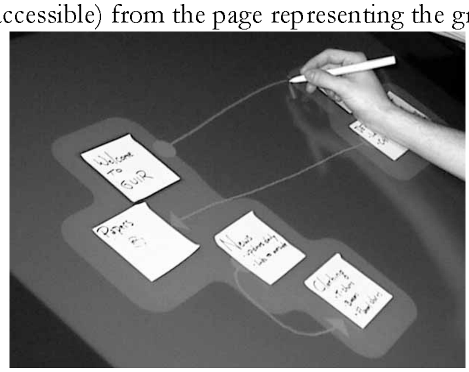

Figure 8: Links provide another layer of expressive capability in Outpost. Other types of relationships beyond simple grouping can be expressed using links.

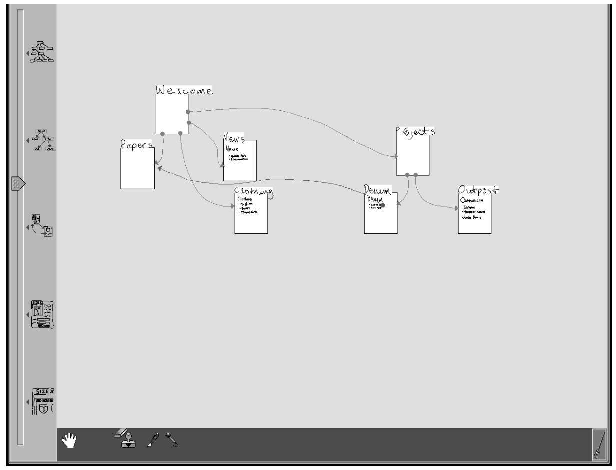

Figure 9: Information hierarchies created in Outpost can be imported into DENIM and serve as baseline sitemaps

Early Work, Low-Fi Prototypes, and User Testing

Early design work, detailed task analysis, and user testing was performed by Raecine Sapien in Summer 1999 [Sapien99]. In addition to suggesting several design directions which we have incorporated and expanded upon in the current design, Raecine created a low-fi prototype using poster board and Post-It Notes to explore how users would interact with a system like Outpost. She conducted a series of user-tests with the low-fi prototype and produced a medium-fi interactive prototype in Java based on the observations from the first user tests. This second prototype used only "virtual" Post-It Notes and explored how users would interact with the grouping mechanisms and visual feedback about the group structure. Our current work is motivated and inspired by the prototypes created by Raecine and the user feedback she collected about those prototypes.

Related Work

"Consider a future device for individual use, which is a sort of mechanized private file and library. … It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works."[Bush45] In many ways, the field of digital information architecture was born with Vannevar Bush's 1945 Atlantic Monthly article *As We May Think*. The motivations for the Designer's Outpost are similar to Bush's Memex. Our work shares with Bush's vision desk-centric space for intuitive data entry and manipulation. Informal interfaces support natural, ambiguous forms of human-computer interaction [Hearst98]. Interaction modes that informal interfaces might support include speaking, writing, gesturing, and sketching. These forms of communication have existed as part of the graphic design process long before the advent of computers. While details of these communication practices must be tailored to somehow express computational tasks, a main thrust of this work is to build the system around actual design practice. For this reason, we see the most important body of related work to be the work that designers perform every day.

Desk interfaces have a rich legacy as well. At MIT Lincoln Lab in 1963 Ivan Sutherland demonstrated Sketch Pad, a pen-based interface for drawing 3D models.[Sutherland63] Inspired by this seminal work, in 1995 Zeleznik et al developed Sketch,[Zeleznik96] a desk and pen-based gestural interface for 3D model construction. Like the Outpost, Sketch uses a rear projected architectural drafting table. Ishii et al have set as their mission to "employ the physical word itself as an interface," using computationally enhanced physical props. A recent project, the Luminous Room,[Underkoffler99] is "an interior architecture space whose surfaces have been made capable both of displaying visual information and performing visual capture." This two-way information stream is of particular relevance to our work. Ishii et al assert the success of an ongoing project, the metaDESK:[Ullmer97] It "has a certain legibility of interface in that its affordances suggest and support user's natural expectations from the device." While the technology of the Tangible Media Group is absolutely amazing, the success of this work is crippled by the fact the group practices technology-centered design. In the group's twenty-eight publications, only one formal user study was conducted. [Underkoffler99b] (A few papers included anecdotes from lab visitors and conference attendees.) We have practiced task-centered design, grounded in our ethnographic studies of the web design process.[Newman99]

Moran, and colleagues at Xerox PARC have been working on interfaces for Tivoli, an interactive whiteboard.[RonbyPedersen93] Discussing an internal ethnography of PARC meeting practice, Moran notes that "Meetings are usually part of a larger work process in which there is a knowledge base of materials that is used between meetings. During a given meeting, only a relevant 'working set' of materials is needed to effectively conduct the meeting. … Pen-based computational systems that allow scribbling and gesturing on wall-size displays can support a whiteboard metaphor for working meetings."[Moran98] We draw from Moran's research on spatially organizing material over large surfaces. This fall, Moran et al presented the Collaborage system; a camera captured whiteboard.

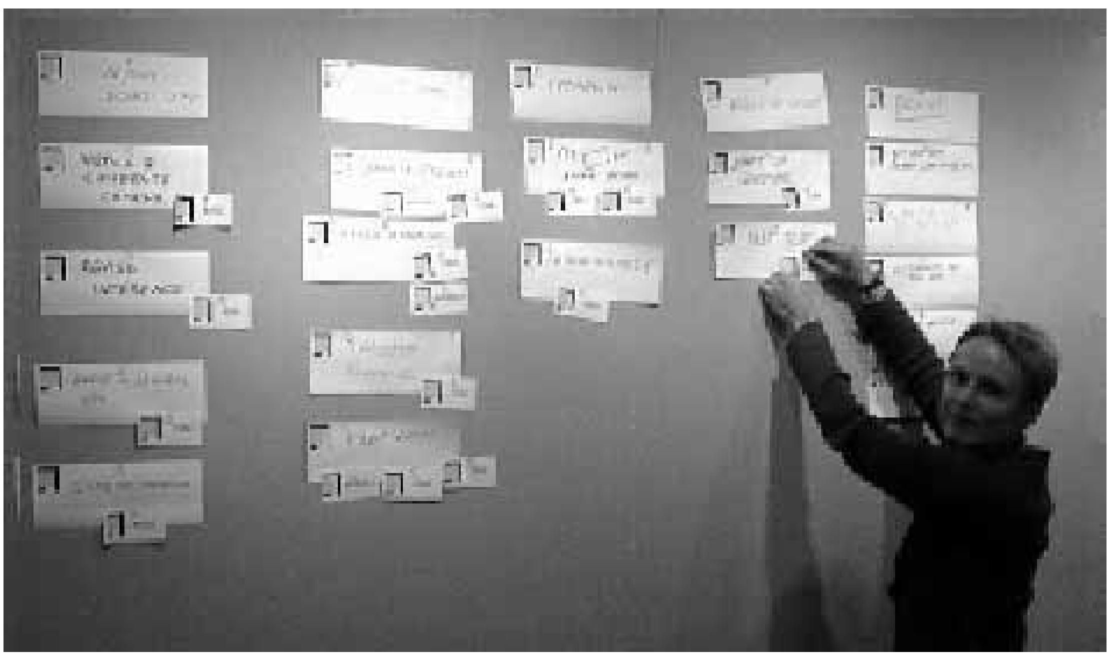

Web site design teams make use of information architects to help organize the hierarchy of the site map, labeling, and searching systems. Often these architects are more comfortable with sketching than coding. Informal interfaces such as SILK and DENIM combine the fluidity of sketching with the appeal of computer based design tools, namely that the final product is interactive and can be easily modified.[Landay95] Still, there is a haptic aspect to paper-based design that is missing from even informal screen-based design tools. Wellner notes that "Consequently, we have two desks: one for 'paper pushing' and one for 'pixel pushing.' " Wellner lists the advantages of "pixel pushing" electronic documents as: "Quick to edit, copy, transmit, share, file, and retrieve. Allows keyword searching, spell checking, instant calculations." And the advantages of paper documents as "Three dimensional, universally accepted, cheap, portable, familiar, high resolution, easy to read, tactile, can use both hands and fingers to manipulate, and can doodle on with a pencil."[Wellner93] The goal of the Outpost is to have the best of both worlds; in our system paper pushing is pixel pushing. A user has the same fundamental capabilities in the Outpost system as in a non-computational paper based system: she can create new "pages" by writing on a new Post-it note and organize a site by physically moving Post-it notes around the desk. The system also captures text, and infers structure about the site; we discuss this further in a later section.

System hardware

The key hardware in the system is comprised of a Windows 98 PC connected to an ITI Digital Desk, an inking pen, and a tablet. The desk is large enough for users to comfortably work with up to around eighty Post-it notes[Sapien99]. We use a Cross iPen tablet and a compatible inking Cross pen. To create a new page, users write on a post-it placed on the tablet. The tablet has the advantage of much better resolution (critical for ink capture) and portability. This is clearly non-ideal interaction. In a few years, when higher resolution desks are available, this will no longer be necessary. The tablet, the desk, and the Post-its are all capable of pen input. It would benefit the user if one pen worked with all elements of the system. The difficulty is that the desk has a non-inking active pen; and it is undesirable to leave a physical ink trail when gesturing on the desk or tablet screen. One potential solution may be hard pencil lead, which leaves no marks on the smooth plastic desk and screen surfaces. Replacing the desk's plastic tip with a pencil lead would work for the desk + screen case. In the desk + tablet case, one could tape a non-inking Cross pen to the side of the desk pen, and slightly raised. While I believe the desk pen only transmits marks when the tip is depressed, the Wacom tablet is aware of the pen any time the tip is near the tablet surface. The tablet's model of pen position is translated form the user's model of pen position, but if we fashion the "doublepen" in such a way that it encourages one holding style, the system can automatically compensate because translation is deterministic.

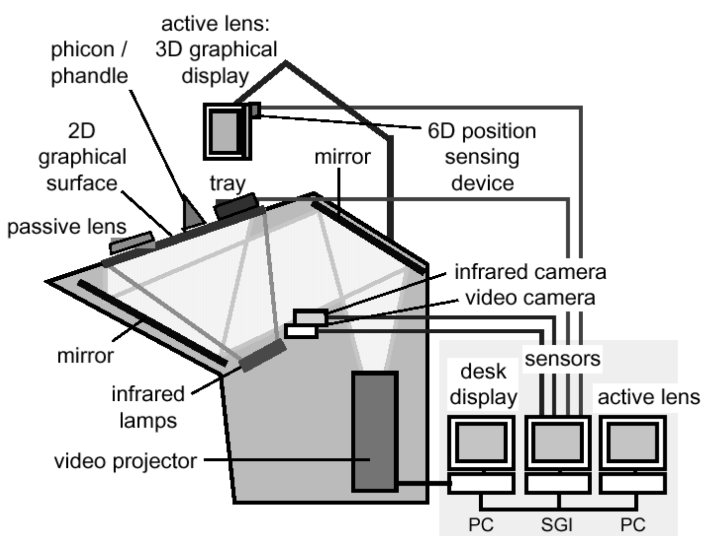

Figure 3: The Tangible Media Group's metaDESK

Camera hardware

"The main difficulty in implementing a system to support computer-based interaction with paper documents is to relate a selected location on the paper to the text at that position." [Wellner p. 589] The Wellner digital desk "uses supplementary cameras zoomed into narrow fields of view" to obtain images for “limited” OCR. The Marcel project is more rational; it uses “low-resolution video images, corrects them for skew, and correlated them with document images obtained from a printer or high-resolution scanner.”[Newman92] The drawbacks are 1) printing and scanning are slow processes, and 2) the user’s body often occludes the above mounted camera’s view. We adopt a similar method, sans these 2 two drawbacks. We use pair of 640x480 capture resolution, 24-bit color, USB cameras for image input. The Microsoft Research Vision SDK handles the details of low level image capture. One camera is mounted inside the desk, and is able to track the perimeters of Post-its without the difficulty of being occluded by the user. (We borrow this idea from the Ishii metaDESK.[Ullmer97]) The second camera is ceiling mounted. Like Marcel, our ceiling camera captures a low-resolution version of the document, which is rotation and translation normalized for comparison. Also like Marcel, this capture is not how the system stores note data. We capture our high-resolution version for storage as the user writes the note on the desk or tablet; eliminating the need to scan.

“We are especially interested in using [computationally] passive, minimally tagged physical objects as TUI controls whenever possible. ” [Ullmer97] Our paper interface uses off the shelf Post-its; completely passive and untagged. Long term, we may be interested in position tracking privileged physical objects for use as phicons; here we plan to follow the Ishii philosophy of minimal tagging. Because these objects are “included” with the system and not brought in externally, it makes sense to have some tagging when it can offer a performance gain.

The Vision Pipeline

At system start-up there is a one-time camera calibration sequence. Our goal is to track solid colored rectangular objects: to know their location, and to know when they are added or removed. In the first step of the calibration, the user places a large sample of the target color on part of the screen. The user then moves the object to a different part of the screen. By computing the difference of these two images, we have constructed the three basic input signals. They are: object add, object remove, and no change. We use an iterative algorithm to partition these signals. We then use the signal partitioning to perform a three-value threshold at runtime. Second, we let the camera capture ten frames of no activity. Our cameras have a lot of noise; we use these ten frames to parameterize our notion of a frame that has no activity. Finally, we calibrate the size of the capture object. We ask the user to place one example object of the appropriate size onto the desk. We capture a frame, and then ask the user to remove it, and capture a second frame. We difference these two images, and perform a three-value threshold. Using the connected components algorithm, we can find the size of the added object. We take a percentage of this (currently 80%) as the minimum allowable size of a valid object signal. In the future, we could also infer that signals much larger in mass than the example signal are a set of multiple overlapping Post-its.

With this calibration information, our object recognition task is fully parameterized. For each captured frame n, we compute the pixelwise difference of n and n-1. If the single difference image contains activity above the noise threshold, we claim the system is in a state of flux. (We save the last calm image c1.) We wait until the system returns to a state of calm, and examine the new clam state image c2. We perform an “uberdiff” of c2 – c1. We then threshold and perform connected components on the uberdiff. The output of this is the set of objects that have changed. (The thresholded image difference subtracts out environment and the portion of the state that is unchanged.) We also know which of these objects are remove objects and which are add objects. We can then pass this knowledge about location and state to our state maintenance system for semantic identification.

Future Work – Camera Calibration

In order to provide correlated feedback about semantic structure of object relationships, we need to know where in desk space the tracked objects are. This is the canonical augmented reality problem; we need to correlate camera pixels with desk pixels. In our case, the correlation problem is constructing a projective mapping from desk coordinates to camera coordinates. Our application has some properties that make calibration difficult, and others that simplify it. One difficulty is that the optics of our camera system are poor; the deviation from a theoretical camera model is probably significant. Our mapping must also be real time, and cannot be a significant computational portion of the entire vision process. Our task is eased because we know the shape of all target objects ahead of time. Post-its are small rectangles. Two corner points uniquely determine the three DOF (two translation plus one rotation) of a rectangular object in a plane; and four corner points enable error minimization. The desk is a large rectangle. Calibration only needs to happen once, when the system is started, so it's okay if the process is lengthy. Finally, the desk is capable of broadcasting images, enabling a feedback loop. Four points uniquely define a projective warp; we use the four corner points of the desk screen. Regardless of the specifics of the calibration algorithm, the camera to desk lookup table needs to be computed twice: once for the top camera and once for the bottom camera. Both lookup tables can be built simultaneously using the same display sequence of structured light data.

The bottom camera is responsible for the location of Post-its; the top camera is responsible for ink recognition. Ink recognition is a slow task. Luckily, e only need to perform ink recognition when a new Post-it appears on the desk. While a Post-it remains on the desk, even as it is moved around on the desk, the system uses frame-to-frame coherence algorithms to continue associating Post-it data with the physical Post-it object. The metaDESK project (using the same desk hardware as this work) used a below screen camera and found that objects disappeared from view once they were more than 10cm above the table. Above this height we start with a fresh slate.

Future Directions

When writing on the desk, the system knows the location of all post-its; it would be beneficial for users id we allowed for adding new ink to existing Post-its. On the tablet, we need more information. The metaDESK has objects that are stored in tagged trays. "By recording when these (labeled) objects leave the tray and applying this label to the next unlabeled object in the infrared vision scene, we are able to successfully identify and track physical objects on the desk." In general, this heuristic seems too error prone for our interests. For the particular case of pulling a Post-it off the desk and placing it on the tablet, it may work, but it still seems fragile. An alternate option is to require that the user place the tablet on the desk prior to Post-it entry. Now the Post-it in question is under surveillance from the above mounted camera and we can notify the system it is being updated.

After placing Post-its on the desk for the first time, users are free to keep or discard the physical Post-it. Because web site design takes place over an extended period of time, users may find it convenient to throw away the mass of Post-its, or they may lose them. The virtual Post-it stays. Users may write on the virtual Post-it on the desk. (And if we have a screen tablet to use as a high resolution Toolglass, they may also place the tablet on the desk and write.) One of our motivations for this work was the excellent haptic feel of Post-its; how do we retain this after they are discarded. The answer is generic Post-it "frames:" a paper Post-it laminated to a metal sheet with the center cut out such that only the border remains. Placing a frame onto the desk "attaches" the two. Users can then move the physical frame around to update the arrangement in the virtual space. An alternative to the frame interaction would be to use PDA's. Placing a PDA on the desk surfaces "picks up" a virtual post-it. It can be carried around and text can be added to it. When the PDA is placed back on the desk surface, the note is "dropped."

Sets of Post-its can be grouped together. Sapien's user testing provides valuable feedback here as well; an intuitive grouping rule is that any Post-its that are in physically touching are in the same group. Another structural element we support are links. If a user draws a line on the desk from one Post-it to another, we display the line trace (with an arrowhead drawn at the end) and semantically add a link from the originating page to the destination page.

In the design process, it is important to have meta-data about the design artifact. We support two types of meta-data: color and annotations. In our system we choose to have the system not infer structure from Post-it color. Designers are free however, to use any color Post-its they choose, and the virtual Post-its retain the color of the physical Post-its. Designers may want to use color to represent potential problem pages, correlate color with which team member is working on the project, or any other information they choose. Designers are free to change the color of the virtual Post-it at any time, using a color marking menu. Annotations are semantically interpreted however. 3M makes arrow shaped Post-its; our system recognizes these Post-its (and the writing on them) as representing information about a page or area. Overlapping an annotation with a page semantically attaches the two. While both Post-its remain on the desk, the two move independently. The two move in unison when only one Post-it is on the desk, or if a frame is used for manipulation. Our last bit of meta-data is timing information. The system representation of the Post-it has a creation time stamp, along with a time stamp for the end of every move and a time stamp whenever the physical Post-it disappears from or appears on the desk.

Conclusion

The Designer's Outpost is a task-centered tangible interface for web site information design. It's functions are informed by observations of real web site design practice and we believe it provides many of the affordances of current tools and practices while offering the advantages of electronic media.

We have described the user interaction with Outpost as well as the hardware and software infrastructure. To date, we have assembled the hardware and completed a substantial portion of the software infrastructure. In the near future, we intend to finish the implementation of the design described here, and validate Outpost's usefulness with practicing web designers.

Bibliography

[Beyer98] H. Beyer and K. Holtzblatt. *Contextual Design*. Morgan Kaufmann Publishers, Inc., San Francisco, 1998.

[Bier94] E.A. Bier, M.C. Stone, et al. Toolglass and Magic Lenses: The See-Through Interface. In *Proceedings of SIGGRAPH 1994*. pp. 445-452, ACM Press, August 1994

[Bush45] V. Bush. As We May Think. *The Atlantic Monthly,* July 1945.

[Gross96] M.D. Gross and E.Y.L. Go. Ambiguous Intentions. *Proceedings of the ACM Symposium on User Interface Software and Technology 1996*, pp. 183-192, ACM Press, New York, 1996.

[Hearst98] M. Hearst, M.D. Gross, J.A. Landay, and T.F. Stahovich. Sketching Intelligent Systems. *IEEE Intelligent Systems, vol. 13*, pp. 10-19, 1998.

[Landay95] J.A. Landay and B.A. Myers. SILK: Sketching Interfaces Like Krazy. In *CHI: Conference Proceedings on Human factors in computing systems 1995*, pp. 43-50, ACM Press, 1995.

[Landay96] J.A. Landay, *Interactive Sketching for the Early Stages of User Interface Design*, Ph.D. dissertation, pp. 242, CMU-CS-96-201, Carnegie Mellon University Computer Science Dept., Pittsburgh, 1996.

[Lin00] J. Lin, M. Newman, J. Hong, and J. A. Landay. Denim: Finding a Tighter Fit with Web Design Practice. In *CHI 2000*, To Appear.

[Moran98] T.P. Moran, W. van Melle, and P. Chiu. Tailorable Domain Objects as Meeting Tools for an Electronic Whiteboard. In *CSCW '98. Proceedings of the ACM 1998 conference on Computer Supported Collaborative Work*. pp. 295-304, ACM Press, 1998.

[Moran99] T.P. Moran, E. Saund, W. van Melle, A. Gujar, K. Fishkin, and B, Harrison. Design and Technology for Colllaborage: Collaborative Collages of Information on Physical Walls. In *Proceedings of Symposium on User Interface Software and Technology 1999*. Pp. 197-206, ACM Press, November 1999.

[Nielsen99] J. Nielsen. Test it! Jakob Nielsen's UI Tips. In *C|net's Builder.com*. http://www.builder.com/Graphics/UserInterface/ss01a.html

[Newman99] M. Newman and J.A. Landay, *Sitemaps, Storyboards, and Specifications: A Sketch of Web Site Design Practice as Manifested Through Artifacts*. Technical Report UCB//CSD-99-1062, University of California, Berkeley, Computer Science Division, Berkeley, CA, September 1999.

[Newman92] W. Newman and P. Wellner. A Desk Supporting Computer-based Interaction with Paper Documents. In *CHI: Conference Proceedings on Human factors in computing systems 1992*, pp. 587-592, ACM Press, 1992.

[Numinous99] Numinous Technologies is no longer, http://news.cnet.com/news/0-1003-200-340452.html

[Rettig94] M. Rettig. Prototyping for Tiny Fingers. *Communications of the ACM, vol. 37*, pp. 21-27, April 1994.

[RonbyPedersen93] E. RønbyPedersen, K. McCall, et al. Tivoli an electronic whiteboard for informal workgroup meetings. In *CHI: Conference Proceedings on Human factors in computing systems 1993*, pp. 391-398, ACM Press, 1993

[Rosenfeld95] L. Rosenfeld and P. Morville. *Information Architecture for the World Wide Web*. O'Reilly & Associates, Inc. Sebastopol, CA, 1998.

[Sapein99] R. Sapien, M. Newman, and J.A. Landay. Combining Tangible and Informal Interfaces for Early Web Site Design. Unpublished work.

[Stollnitz96] E.J. Stollnitz, T.D. DeRose, and D.H. Salesin. *Wavelets for Computer Graphics: Theory and Applications*. Morgan Kaufmann, San Francisco, 1996.

[Sumner97] T. Sumner and M. Stolze. Evolution, Not Revolution: Participatory Design in the Toolbelt Era. In *Computers and Design in Context*, M. Kyng and L. Matthiassen, eds. MIT Press. Cambridge, MA, 1997.

[Sutherland63] I.E. Sutherland. *Sketchpad: A Man-Machine Graphical Communication System*. MIT Lincoln Laboratory Technical Report #296, January 1963.

[Ullmer97] B. Ullmer and H. Ishii. The metaDESK: Models and Prototypes for Tangible User Interfaces. In *Proceedings of Symposium on User Interface Software and Technology 1997*. Pp. 223-232, ACM Press, October 1997.

[Underkoffler99] J. Underkoffler, B. Ullmer, and H. Ishii. Emancipated Pixels: Real-World Graphics in the Luminous Room. In *Proceedings of SIGGRAPH 1999*. pp. 385-392, ACM Press, August 1999

[Underkoffler99b] J. Underkoffler and H. Ishii. Urp: A Luminous-Tangible Workbench for Urban Planning and Design. In *CHI: Proceedings of Conference on Human Factors in Computing Systems 1999*, pp. 386-393, ACM Press, May 1999.

[Wellner91] P. Wellner. The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display. In *Proceedings of Symposium on User Interface Software and Technology 1991*. pp. 27-33, ACM Press, 1991.

[Wellner95] P. Wellner. Interacting with Paper on the DigitalDesk. In *Communications of the ACM Vol. 36 No. 7*, pp. 87-96, ACM Press, July 1993.

[Wong92] Y. Y. Wong. Rough and Ready Prototypes: Lessons from Graphic Design. In *CHI 1992*, pp. 83-84, ACM Press, New York, 1992.

[Zeleznik96] R.C. Zeleznik, K. Herndon, and J.F. Hughes. Sketch: An Interface for Sketching 3D Scenes. In *Proceedings of SIGGRAPH 1996*, ACM Press, New York, August 1996.