

A Tangible Evolution: System Architecture and Participatory Design Studies of the Designers' Outpost

Scott Klemmer, Mark Newman, Ryan Farrell, Raecine Meza, James A. Landay

Group for User Interface Research

Computer Science Division

University of California at Berkeley

Berkeley, CA 94720 USA

+1 510 642 4948

srk@cs.berkeley.edu

ABSTRACT

In our previous studies into web design, we found that pens, paper, walls, and tables were often used for explaining, developing, and communicating ideas during the early phases of design. One practice was to produce information architectures using paper on walls. This practice inspired us to work towards The Designers' Outpost, a tangible user interface that combines the affordances of paper and large physical workspaces with the advantages of electronic media to support information design. In this paper, we describe a series of participatory design studies that explored the combination of physical and electronic media in depth. These three studies employed both low- and high-fidelity prototypes. The design teams in our studies encouraged us to support free ink electronic annotations to sitemap pages, versioning of design artifacts, fluid transitions to other tools, and opportunities for collocated and remote collaboration. In parallel, we built a set of prototypes for the underlying vision system. These prototypes led us to difference image-based recognition algorithms and a two-camera infrastructure: a rear-mounted video camera for capturing movement, and a front-mounted high-resolution camera for capturing ink.

Keywords

Tangible Interfaces, Web Design, Sketching, Information Architecture, Computer Vision, Informal Interfaces, CSCW

INTRODUCTION

In our previous studies into web design [18], we found that pens, whiteboards, paper, walls, and tables were the primary tools used for explaining, developing, and communicating ideas during the early phases of design. Later phase design, where detailed page mockups are generated, occurs mostly on the computer. This finding is not surprising, and is consistent with work practice studies across many design and engineering domains [1, 8, 34]. In one common early-phase practice, designers collect ideas about what should be in a web site onto Post-it notes, and arrange them on the wall into categories. This technique, often called affinity diagramming [2], is a form of collaborative sketching used to determine the site structure.

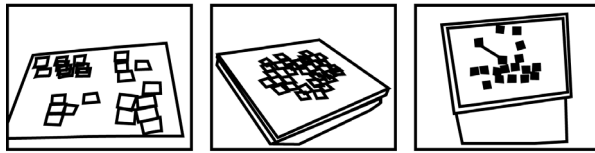
The large workspace offers several clear benefits for the task. In contrast with the heavyweight, formal operations of the computer, it is relatively easy to fill a wall with pieces of paper and move them around to suggest different associations. It permits the representation of large, complex information spaces without the loss of contextual, peripheral information. 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 allowing multiple people to simultaneously view, discuss, and modify the artifact.

However, there are drawbacks to a paper-centric representation. Much of the information exists in the *relationship* between information chunks (Post-it notes). Because structure must be maintained manually, marks the designers make about the data, such as links or groups, often fall out of sync as notes are shifted around. At some point, whether hours after a brainstorming session or months after a project, the paper is removed and the site structure display is lost. The designers in our studies also lamented that versioning is unfeasible in a paper only representation.

This work practice offers few opportunities for remote participants, at a desktop down the hall or in a meeting room across the world. Remote users not only have no way to update the information, they have no way to even access the information. We also found, as others have, the transition from the early paper-centric design stages to the later pixel-centric stages to be highly problematic [10, 33]. Those who worked hard to architect the site feel a sense of lost ownership as the project is "thrown over the wall" to graphic designers and developers. And as the site structure is changed during the development, the early paper artifact drifts further and further out of date.

Supporting Web Practice With The Designers' Outpost

In this paper, we describe a series of participatory design studies (Figure 1) and a series of computer vision studies (Figure 2) to develop The Designers' Outpost, a tangible user interface that combines the affordances of paper and large physical workspaces with the advantages of electronic



First: Paper desk prototype Second: Paper and pixel desk Third: Interactive Wall

Figure 1: The sequence of design studies.

media to support information design for the web. Paper in the physical world becomes an input device for the electronic world. Projectors output electronic information onto surfaces in the physical world. 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 new Post-it notes and organize a site by physically moving Post-it notes around the wall. In addition to easing basic information architecture tasks, our system will support the transition from this early representation to later electronic artifacts, such as a formal sitemap.

Our first design study validated the general approach. It indicated a need to minimize the extra user effort required to use the tool, and encouraged us to allow the interaction to be as freeform as possible. Our next prototype fleshed out the interaction techniques and showed that a drafting desk is too small to support the structures that professional web sites require. While our primary early interest in Outpost was to provide interactive support for information architecture design meetings; designers in our most recent studies found constant interactive feedback distracting. They encouraged us to refocus our interface on: supporting free ink electronic annotations to sitemap pages, versioning of design artifacts, fluid transitions to tools such as DENIM [13] and supporting collocated and remote collaboration. We also found that the system is much more appropriate for information architects than for visual interface designers.

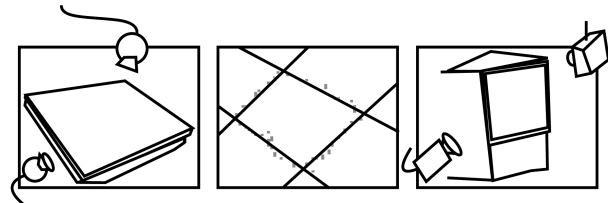
RELATED WORK

Our research is inspired by previous work in two areas – early stage web site design and tangible user interfaces, a technology that seems well suited to this domain. We describe these two areas next.

Web Site Design Practice

The goal of our earlier investigation into web design [18] was to inform the design of systems to better support web design practices. The study consisted of interviews with eleven professional web site designers from five different companies. Each interview consisted of asking the designer to choose a recent project 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 their meaning with respect to the process as a whole.

Three important observations were made during the course of this study. First, designers create many different



First: Difference Image Second: Matlab algorithms Third: Fully Interactive

Figure 2: The sequence of computer vision studies.

representations of a web site. Second, the production and use of these *intermediate artifacts* dominate the day-to-day work practice for most of the design process. Third, we learned that web design is comprised of several sub-specialties, including information architecture [23] and visual design, each of which has its own tools, products, and concerns. We found that information architecture is not well supported by current software tools.

Designers create many different representations when designing a web site. 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 sitemaps, which depict sites as related blocks of labeled information to mockups which depict individual pages in high-fidelity.

Bellotti and Rogers conducted a study on web publishing workflow [1]. They too discovered a tension between paper-based practices and electronic practices. In particular, they found that people were often more comfortable working on paper, but felt that electronic tools were beneficial for stronger communication and awareness among distributed teams. One site director commented “What I would love would be a flat panel I could hang on a wall... For the tacked up paper and string setup we have, a video wall could be really useful, not just for the sake of more expensive equipment, but for working with remote group members, for ease of modification, and for keeping a better record of the evolution of the site.”

DENIM

Based on the results of our studies, we developed DENIM, a sketch-based tool supporting information and navigation design of web sites [13]. DENIM (Figure 3) supports sketching input, allows design at different refinement levels, and unifies the levels through zooming. In particular, DENIM supports visualizations matching the *sitemap*, *storyboard*, and *schematic* representations of a web site. DENIM also allows designers to interact with their site designs through a “run mode,” which displays the sketched pages in a limited functionality “browser” that allows the user to navigate the site by clicking active regions of the sketches and linking to other pages within the site. While DENIM supports authoring site maps, it is best suited for storyboards and page schematics. Outpost targets the artifact generated at the start of the design exploration phase: sitemaps.

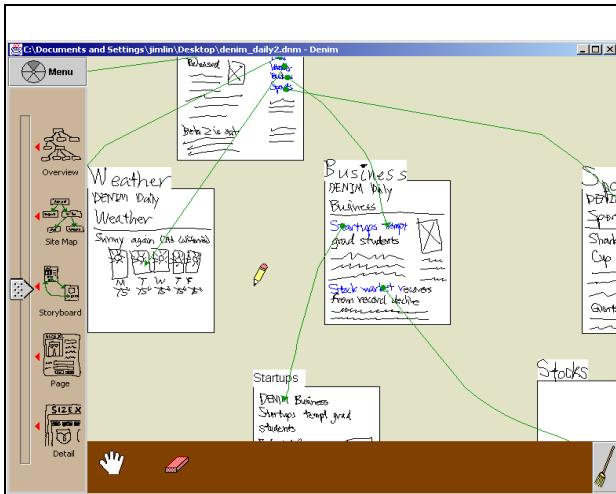


Figure 3. DENIM, shown here in “Storyboard View,” allows designers to design web sites by sketching and integrates sitemap, storyboard, and individual page representations through zooming. Information spaces created in Outpost will be imported into DENIM, serving as baseline sitemaps.

Affinity Diagrams

Our interest in researching computational support for information architecture is motivated by one specific design practice observed during the ethnographic study discussed above. This collaborative 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 are labeled and further grouped into hierarchies of groups. This hierarchical structure serves as a baseline for the structure of the web site. A version of this practice is discussed in detail by Beyer and Holtzblatt in their book *Contextual Design* [2]. They call the technique “Affinity Diagramming.” Usability expert Jakob Nielsen also advocates a version of this method, using index cards to design the information hierarchy [19].

Tangible User Interfaces

In the early 1990s, Pierre Wellner and colleagues produced several prototypes of a DigitalDesk system that used ceiling mounted cameras to track documents and hands on a physical desktop, and a ceiling mounted projector to augment a real desk with electronic information [32, 33]. MIT’s Tangible Media Group later created the metaDESK [27], a digital desk employing tangible interfaces as the controls for and views of a map of the MIT campus. Another MIT project, transBOARD [9], uses tagged whiteboard markers so that remote users can view the whiteboard contents on their desktop screen.

Researchers at Xerox PARC have investigated interactions on walls and electronic whiteboards, including the LiveBoard hardware [6] and interfaces for it [16]. Recently,

they developed the Collaborage system. “A Collaborage is a collaborative collage of physically represented information on a surface that is connected with electronic information, such as a physical In/Out board connected to a people-locator database” [17]. Their system uses computer vision to detect and identify data glyph tagged documents on walls and whiteboards. (A data glyph is a type of 2d barcode.)

Streitz’s group at GMD has developed i-LAND, a system of interaction techniques for working with display surfaces embedded in rooms and furniture [25]. As part of their work, they developed the DynaWall: three adjacent electronic whiteboards that take input via hand gestures. Winograd and Guimbretiere are also working on interaction techniques for large, tiled projector surfaces [35]. Rekimoto and Saitoh [22] developed a system to integrate laptop computers, projected surfaces, and tagged physical objects.

This body of work motivates the concept that for many tasks, computer input is best performed by manipulating physical objects on large surfaces.

THREE DESIGN STUDIES

We approached our research on tools for collaborative information architecture by creating prototypes that enabled user feedback and by designing a computer vision back-end for capture of paper artifacts. In these sections we discuss the Outpost interface and design studies. In later sections, we describe the vision back-end. To explore the viability of combining physical and electronic representations for web site information architecture, we undertook a series of three design studies. We first evaluated the basic concept with a paper prototype study. Next, we built interface mock-ups that envisioned the combination of physical artifact state with interactive feedback. Finally, we created a wall-scale prototype for a set of participatory design sessions with fifteen professional interface designers.

Low-Fidelity Desk: Design Study

We created our initial low-fidelity prototype using cardboard the size of an ITI Digital Desk (41” diagonal), evaluating this paper prototype with two individual participants. The participants wrote on a pad of 3” x 3”

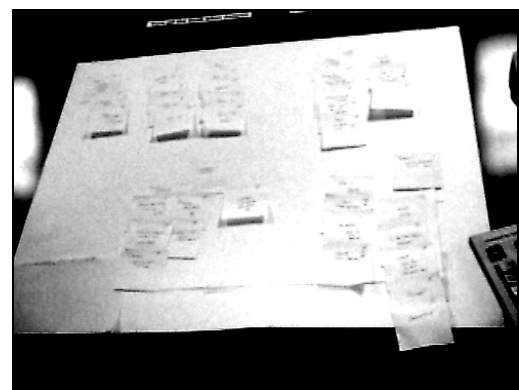


Figure 4. The low-fidelity Designers’ Outpost.



Figure 5. Mock-up of the Designers' Outpost - Collaborating on an information hierarchy with Post-its on a digital desk.

yellow Post-it notes using an inking pen (Figure 4). We gave participants the task of creating the information architecture for a web site about off-campus housing for college students. To start, we handed them six pages of notes from interviews with college students seeking housing. The task included chunking interview information onto Post-its, arranging the Post-its into related groups, and merging two previously saved versions of Post-its into a unified version. A wizard acting as the computer gave verbal feedback about what the computer recognized as groups, which groupings were being selected, and displayed widgets and dialog boxes when appropriate.

We found that participants often forgot required system steps that had no affordance or feedback, such as underlining the note representing a group or pressing an upload button to add a note to the system. This suggested an interface where we automatically recognize as many actions as possible (e.g., a new note should be automatically added when it is placed on the desk). Users were also confused by the three input devices: the inking pen for writing on notes, the virtual stylus for authoring note relationships, and the keyboard for entering version names. As a result, We removed a keyboard from our system design: this simultaneously simplifies the input model and better matches current practice. We also plan to use a single two-ended pen: one end a stylus and the other an inking pen.

Pixel and Paper Mock-up

We took our findings from the paper prototype and created a mock-up of our ideas for combined physical/virtual interaction. The mock-up was built using an ITI Digital Desk, which is a rear-projected surface with the size and slope of an architects drafting table. The interface mock-up showed physical Post-it notes and the corresponding electronic feedback for the core Outpost interactions: authoring groups, authoring links, labeling the groups, and organizing the groups into hierarchies [11] (Figure 5). It

quickly became evident that a digital desk is too small a space for professional web site information architecture; it allows for a maximum of fifty Post-its and two or three users. Information architects often use upwards of two hundred Post-its and four to eight people might be simultaneous participants in design sessions. To build the Designers' Outpost at a full collaborative scale, we moved our design to a SMART Board, a much larger rear-projected surface in the form factor of a whiteboard [24].

Interactive Wall: Interface Prototype

Our low-fidelity and mock-up prototypes informed the design of our interactive prototype. We used this system for a participatory design study with fifteen professional interface designers in five sessions.

The prototype was implemented as a Java application running on a rear-projected 72" diagonal touch-sensitive SMART Board [24] with a 1280x1024 LCD projector. With this prototype, we recognize the location of notes on the board using the board's touch sensor. (The vision-backed system we describe later accurately recognizes location and orientation of notes using computer vision; this system was not robust enough for real use at the time of this design study.) Drawing a line from one note to another note with the board stylus creates a link. The stylus is also used for creating freehand electronic ink on the board. This freehand ink is a feature of the SMART Board and currently is not captured in our system. Virtual ink might be useful for annotating physical notes, for drawing virtual notes, or for freeform scribbling.

Tapping on a note invokes a context menu (Figure 6) that lets users either delete the note or define it as the label note for its group. (In the vision-backed Outpost system described later, removing a note from the board deletes it.) Individual touch-recognized notes are thinly outlined in gray. Notes that are close to each other are deemed to be in a group; these notes are heavily outlined in blue.

Interactive Wall: Design Study

We were ready to bring professional web designers into our lab. We ran five design sessions with between two and five designers per session, for a total of fifteen participants. In four of the design sessions, the designers were colleagues at the same company; the fifth session mixed designers from two companies. Two of the five groups were composed of information architects, two groups were visual designers,



Figure 6. Tapping on a note brings up a context menu.

Role	INFORMATION ARCHITECTS		BOTH	VISUAL DESIGNERS	
Group	A	B	C	D	E
Size	5	3	2	2	3

Figure 7. The five study groups: their size and primary role

and one group had individuals performing both roles (Figure 7). Information architects are mainly concerned with the information and navigation design of a web site. Visual designers typically focus on interaction and graphic design.

Each session lasted roughly two hours (Figure 8). We began the sessions with a high-level overview of the project and a brief demo of the existing prototype. We gave the designers a mock information architecture design task to work out with their team using the prototype. We conversed freely with the designers during the sessions. We asked questions and the designers verbalized what they were doing and offered thoughts on how the tool should work. The design task took 45-60 minutes. This was followed by a fifteen minute demonstration of DENIM and then a 45 minute discussion on Outpost's utility and its relationship with DENIM and their current work practices. They finished with a seventeen question written survey asking about their background and their opinions about the usefulness of Outpost in their work. We videotaped all of the sessions; figures 9 through 11 are stills from this camera. In each study, we had two researchers. One of us was in charge of communication, explaining the system, and facilitating discussions. The other took written notes and was in charge of the video camera.

Interactive Wall: Design Findings

Our findings from this participatory design study offer insight into the designers' collaborative work process and suggest an appropriate interactivity model.

Existing Work Process

Every participant currently works with groups on whiteboards early in the site design process. The information architects all said they currently create sitemaps by placing Post-it notes on the board, while the visual designers talked about sketching page designs directly on the board. Whiteboard meeting capture was highly valued by all five teams. Three of the design teams currently use a digital camera, one uses a whiteboard capture device (the Virtual Ink Mimio [29]), and one assigns a scribe to save information from design meetings. One group even uses an application called Whiteboard Photo [21] to rectify and filter out smudges, dirt, and lighting changes in whiteboard photographs.

Also, every designer said that they currently use either the Visio or Inspiration software packages for creating site maps. Site maps can get quite large; one firm said that two to three hundred nodes is typical.

Overview and Demo	15 minutes
Design Task	45 – 60 minutes
DENIM Demo	15 minutes
Discussion	30 – 45 minutes
Survey	10 minutes

Figure 8. The time breakdown of the design sessions.

The post-test questionnaire included a question "How likely is it that you would integrate Outpost as a regular part of your web site design practice?" The response was a five point Likert scale. A quarter of the participants (four) rated the system the top value, "very likely." Half the participants (eight) rated the system the second value, "somewhat likely." And a quarter of the participants rated the system the fourth value, "somewhat unlikely." No one rated the system the middle value, "neither likely or unlikely" or the bottom value "very unlikely."

Interactive Board Work Process

We observed the groups going through three general phases of design when using the interactive prototype.¹ The designers stated that these same phases were part of their existing practice.

Phase I: Brainstorming

The goal of this phase is to quickly put a large number of concepts on the board. One designer said, "Get all these things on Post-its." The notes simply represent ideas. Sometimes, similar information was placed close together. As described by another designer, "I like being able to just cluster stuff first." Designers did not eliminate ideas or link concepts together into any formal structure at this stage.

Clients, marketing representatives, and engineers often attend key brainstorming meetings. One designer commented that the prototype would be "good for times with the client" because after a meeting they could continue to pare down and hone the artifact without having to start with a new tool from scratch.

The designers were adamant about not wanting any system feedback during this phase. "We didn't do anything here that we couldn't do on a normal whiteboard." One team actually turned off the board. We believe that distracting feedback was the primary reason for three participants to feel it is "somewhat unlikely" that they would regularly use the current prototype in their work.

Phase II: Creating a Top-Level Information Architecture

In this phase, designers migrate from a loose federation of notes on the board to a high level information architecture by clustering related information into groups, pruning unnecessary concepts, and linking notes together.

¹ Two of the groups did not start the third phase during our in lab sessions, but said that is what they would do next given more time.

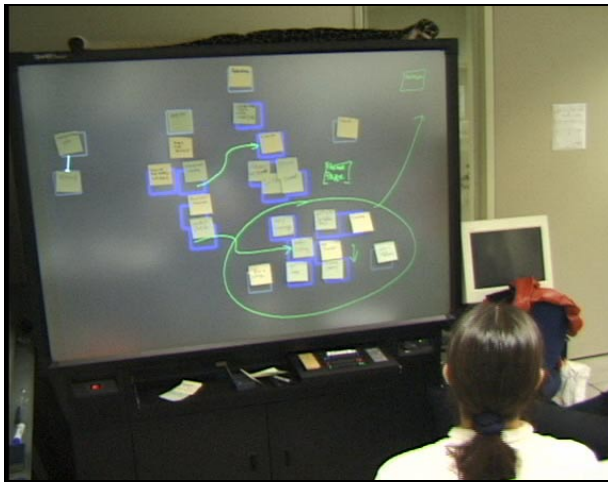


Figure 9. A design team suggests that freehand ink would be useful for both unstructured annotation of the artifact and for performing operations on groups of notes.

The tool support in the interactive prototype was well suited to this phase. This was evident in how fluidly the designers worked in this phase, and by their enthusiastic comments while designing. This was echoed on the post-test questionnaire, where several designers expressed interest in using Outpost for creating top-level information architectures.

Phase III: Drilling Down - Adding Information with Free Ink

The goal here is to “drill down to more detail.” We saw work process differences begin to emerge. The visual designers began to work out basic page designs using empty board space and the board stylus. In contrast, the information architects fully fleshed out the page structure of the site, continuing to add notes.

The key design implication from this phase is the ability to associate freeform ink with individual notes (Figure 9). The visual designers wanted to sketch the design details, and the information architects wanted to add annotations or properties. For example, one information architect said, “I’d like to be able to attach design rationale.” Design rationale is a mechanism for asynchronous communication, embedding in the artifact the motivations for making decisions [14]. The information architects also had a strong desire to use properties for project management. Two groups suggested tagging objects with properties, such as an issue (e.g., will it be possible to get copyright clearance), and later searching for issues across the design.

Overall Process

As reflects their disciplines, the visual designers often talked explicitly about what pages might look like, while the information architecture groups did not discuss visual design at all during the session. Instead, they had active discussions about users and tasks at a more abstract level: “What does the user know here? What is the user trying to do?”

For all the teams, the site representation operated as the central shared artifact for discussion; only about half the time was anyone actually explicitly working at the board. Also, for short periods of time (~1 minute), individuals or subgroups broke off from the main discussion to work on their own. The board remained the anchoring reference point.

We observed two styles of interacting with the board. In the *facilitator* style, one person, usually the senior-most individual, stands at the board (Figure 10). The entire group discusses the site; as the discussion progresses, the facilitator creates notes that synthesize the discussion content. One group also referred to this style as “gate keeping.” This was the primary work practice in three groups, and the groups affirmed that this was their normal work practice.

The second style was *open board*. As with *facilitator*, all group members actively discussed the site. In *open board*, however, there is no central figure; all participants have agency to create notes and directly express their ideas in the artifact. Sometimes everyone was at the board, sometimes just a subgroup (Figure 11). We started the sessions with a single pad of notes and a single marker next to the board. One of the design teams requested one pad and marker per person. This paradigm affords each person their own “input device.” In adding content to the board, information moves from a personal creation space to a shared viewing space.

Several participants commented that they valued simultaneous input with a low-latency response. The SMART Board’s touch sensor only supports one action at a time. Concurrent use of the board has technical design implications for the note sensing technology. This result encourages us to continue building a computer vision system. Vision lends itself both to rich sensing capabilities (note size, color, orientation, and capture of its contents) and simultaneous input.



Figure 10. This is an example of the “facilitator” style.



Figure 11. This is an example of the “open board” style.

Only Information Architects Need Apply

Enthusiasm for the prototype correlated directly with two variables: the percentage of the designer’s work that was web-based, and how much the designer saw their role as an information architect rather than a visual designer. While one visual designer felt, “We don’t really do sitemaps so much. Our interfaces tend to end up with one or two screens,” the information architects saw creating sitemaps as a challenging process of creating the core of a web site. The information architects praised our faithfulness to their current wall-scale work practices, and were enthusiastic about the combined tangible/virtual interaction. However, even the groups that did not value a site map tool were very interested in the wall-scale form factor. After showing them the DENIM prototype, they expressed interest in creating and discussing web sites with DENIM in groups on the SMART Board.

Design Implications

This study underscored several important points about how calm [31] an informal design tool must be; the system feedback should not interrupt the designers flow state.

Smart Yet Silent

We originally felt that one benefit of the prototype was that the system automatically recognized groups based on note proximity and provided visual feedback. However, the designers unanimously felt that automatic grouping by proximity was not useful; they already know where the notes are. The designers offered the more useful alternative of explicitly circling a set of objects to group them.

Furthermore, the group, note outline, and menu feedback was considered distracting. One designer said, “I’m totally disturbed while I’m trying to concentrate on what we are doing. There are too many things flashing.” In hindsight, this result is consistent with the negative user opinion about automatic interpretation and feedback in SILK [12], an informal GUI design tool.

Several participants valued the subtle visual relationships between notes. “Automatically arranging them would take

away from my thinking.” One designer said that she wanted “to work with this before it’s turned on.” Her group wanted to later turn the board on and get the feedback and interaction. This implies that only explicit user actions should cause visible system actions.

In general, interactive features should be available for designers as they move from brainstorming (Phase I) into more explicitly creating a sitemap (Phases II and III). Ideally, interactive functionality should be available but not automatic. When functionality must be automatic, it may make sense to have an explicit “switch” for designers to designate whether they would like an active or passive system.

Sweet Spot on the Tangible/Virtual Spectrum

We have seen informal, wall-scale interfaces that are completely virtual [15], and we have seen wall-scale interfaces that are completely physical [17]. There are appealing aspects to both interaction paradigms (Figure 12); one of our research goals is to leverage the advantages of both.

Fluidity and Physicality

This series of design studies has provided insight into what we believe to be a sweet spot on the tangible/virtual spectrum. Working physically supports collocated collaborative processes. The direct manipulation affordances of physical notes make them easier to see, move, and share.

We reviewed the study videotapes to loosely quantify the pace of interaction. We found that on average, a note was added to the board roughly every 25 seconds. During active periods, a note was added every three to five seconds. Often, there was no explicit interaction for minutes at a time. A good portion of the meetings happened off the board, but referenced the board.

One facilitator began by authoring the sitemap virtually, sketching out square notes and their content. This has the advantage that there is no need to switch between an ink-based pen and a board stylus. A drawback is that the work process was noticeably slower (top speed of one note every seven to ten seconds) because 1) the designer had to create page boundaries rather than using the pre-defined pages torn from a pad, 2) authoring with plastic pens on a plastic surface is awkward for textual input, and 3) the projector ink feedback is much lower resolution than paper. These difficulties negatively impacted the artifact creation process, encouraging minimal input. For example, in one instance, the participant wrote “B” instead of “Business.” Later, he started working physically, and the working pace and artifact quality picked up substantially.

One of our concerns about a tangible/virtual interface was that designers would find it tedious to remove physical objects from the board as they began to flesh out the design in detail (phase III). We asked designers about this issue directly and also watched for frustration with this during the design sessions. Surprisingly, the designers did not

Electronic documents on a virtual desk

Quick to edit, copy, transmit, share, file, and retrieve. Allows keyword searching, spell checking, instant calculations.

Paper documents on a real desk

Three dimensional, universally acceptable, cheap, portable, familiar, high resolution, easier to read. Tactile, can use both hands and fingers to manipulate, and can doodle on with a pencil.

Figure 12. Pierre Wellner's comparison of advantages of electronic and paper documents [22].

mind removing the physical objects; on the contrary, removing them was a natural signifier of a shift in the design process.

At Least a Whiteboard

In our designs, we were careful to preserve many of the successful aspects of working on a traditional whiteboard; the utility of these affordances became apparent in the study. Our prototype permits the representation of large, complex information spaces without the loss of contextual, peripheral information. One designer referred to our interface as “cross-cultural” because engineers, designers, and clients are all comfortable working informally on whiteboards.

Our prototype was not bug free and a touch sensor is a limited input device; sometimes the system provided unexpected feedback or had difficulties redrawing the display. During one session, the system went haywire, and the facilitator jokingly jumped into a boxing stance when she had a tough time “taming the system.” At it’s worst, the prototype was a temperamental whiteboard. “I’m afraid of that thing, that’s why I’m sitting here [on the couch].” This is not the ideal vision of an information appliance [20].

Information appliances should be as easy to learn as physical appliances. When two participants showed up a half hour late, we were pleasantly surprised to see that the participant who was on time was able to quickly bring her colleagues up to speed. After using the tool for only five minutes she was easily able to communicate the conceptual model and the functionality of the prototype.

Extending the Existing Work Process

For a system like Outpost to be successful, it must both support existing work practices and offer designers advantages that only electronic tools can provide. It should offer the annotation features discussed previously, as well as support versioning and a way to transition to later stage tools. Additionally, the system should be flexible enough to allow designers to integrate the physical images they work with every day.

Versioning and Meeting Audio Recording

Our earlier study into web design practice showed that designers desired a way to manage different *versions* of design ideas [18]. Versions play a key role during the design exploration phase. In order to keep track of project milestones and variations, designers are forced to invent ad-hoc methods, usually involving saving multiple versions of files and using complex, cryptic file names to encode the properties of each version. In the physical world, they must either manually photograph, photocopy, or scan an artifact to save the state, or abandon the current state and keep working.

We talked with the design teams about different solutions for versioning support within Outpost. The designers met our suggestion for this potential system functionality with great enthusiasm. Additionally, they felt that an artifact that integrated an audio recording of the design meeting would be helpful. This provides a way of capturing the substantial activity that happens above and beyond directly working with the board. An effective interface should probably associate indexes into audio with board interactions for ease of retrieval; the proper methodology for doing this is still an open research problem.

Several design teams said they would like to replay a design meeting at a faster pace. One designer said that replaying, “gives you an insight into the process you wouldn’t normally get.” The design solution of a shuttle dial, a la video editing, seemed natural to almost everyone. It should be a useful mechanism for browsing changes. We also suggested a *version jar*, a physical object that “contains” a version of an artifact. The jar must be augmented with a system identification mechanism, such as an RFID tag [30]. The designers supported this idea as well, suggesting that it would be useful for explicitly book marking checkpoints and finished designs. Versioning is a central enough issue that having both techniques makes sense. We are currently exploring the possibility of using a property-based document management system [5] as the storage infrastructure for versions.

Transitioning to Other Tools

Every group mentioned that migrating the design artifact to other tools for further refinement would be an essential advantage of the Outpost system. Many of the designers currently photograph meeting whiteboards even though this only produces a static artifact. They were very interested in the prospect of returning to their desk with an interactive site representation that they could continue to work on.

We propose that an appropriate tool for Outpost to transition to is DENIM. DENIM offers the ability to edit the information architecture, specify page level details, and create the navigational structures for a web site. Its interface is oriented towards a single designer working at a PC with a pen interface. Outpost is appropriate for creating sitemaps. DENIM becomes more appropriate when the design team starts to storyboard the specific pages and create schematics. We have begun building a new Outpost

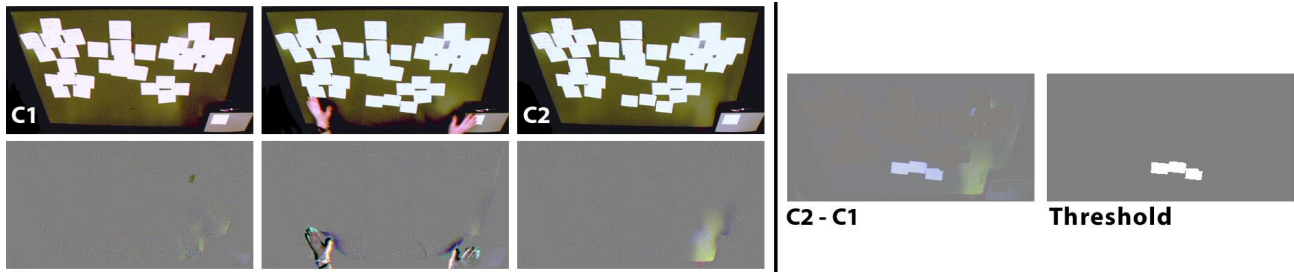


Figure 13: Excerpts from an image sequence from our prototype steady state algorithm. Raw camera frames are shown in the top row, single frame difference image is shown in the bottom row. Raw and thresholded $C2 - C1$ difference at right.

prototype on top of the DENIM code base. In this MVC paradigm, Outpost and DENIM are different interface views and controllers that employ the same model [26]. We plan to add networking socket code that will send design changes to registered Outpost/DENIM clients, enabling real-time and asynchronous participation by remote members of a design team.

Long projects magnify the benefits of having a site map artifact remains in use through the entire design cycle. One design team we spoke with was in the midst of a redesign for a large web site they had originally designed almost a year ago. Outpost has the potential to be a valuable tool for this situation because it lowers the burden of keeping the information architecture updated.

Integrating Physical Images

During the sessions, the design teams often referenced earlier projects, work they'd seen, and personal experiences. As we discussed this referral process with them, we discovered that they often bring in photographs, pages of magazines, and other visual aids to help them design. Outpost should allow designers to stick these visual artifacts on the board, capture, and associate them with pages in the information architecture.

COMPUTER VISION PROTOTYPES

We now briefly present the three computer vision prototypes we have built as a means of illuminating the technology issues involved in building wall-scale tangible interfaces. We use computer vision in our application to locate, capture, and identify the information on Post-it notes (and eventually images) that users place on the board. Computer vision is an appropriate methodology for this task because it can provide automatic untethered and untagged tracking and capture of artifacts users place on the board. Our first prototype was a simple system that computed the difference image between frames, and analyzed this difference image to tell when the board state has changed. In the second prototype, we used Matlab to prototype the full set of algorithms necessary to support the Outpost application. The third prototype was built on top of OpenCV, and implemented the computer vision algorithms at interactive rate, as well as a socket-based network connection for communicating with the Outpost UI.

Difference Image Vision Prototype

The foundation of our recognition system is a technique known as a difference image (Figure 13). Subtracting frame m of a camera stream from frame n results in the difference image $D_{(n-m)}$. D expresses the change in board state between the two points in time.

Difference Image Algorithm and Use

We use difference images to reason about two properties of the board. First, we subtract consecutive frames (the bottom three gray images in Figure 13). The content of a single-frame difference image expresses the activity of the board. If we sum the absolute difference values of all pixels in the image, we can compute an activity metric A . If A is larger than some threshold, then we call the board *active*. Otherwise, we call the board *calm*. In reality, the content of a single frame difference image also results from noise in the camera sensor array and from lighting changes in the world (e.g., someone walks between a light source and the board). We found that in general, mandating that activity be above a threshold was reasonably tolerant to these issues. However, sometimes noise and lighting changes inadvertently triggered activity. In the wall vision prototype described later, we also employ spatial and temporal filtering techniques that help account for camera noise and lighting changes. In this later prototype, we found that this makes activity detection highly accurate.

We also use difference images as an object detection primitive. Our application does not require that we track notes while they are moving, but it does require that the system is aware of a note when it is initially placed down, placed at a new location, or removed. When the board becomes *active*, our system saves the last *calm* frame $C1$. When the board becomes calm again, we capture the new calm frame $C2$. Subtracting $C1$ from $C2$, and thresholding the result tells the system what has changed during that period of activity (the right-hand image in Figure 13). This thresholded difference image becomes the input for our note recognition system.

Design Implications

In building this first prototype, we realized that locating a note is a completely separate problem from identifying a note. Dividing our vision task into two distinct parts (location and ink capture/note identification) enabled us to realize that the system architecture should have two

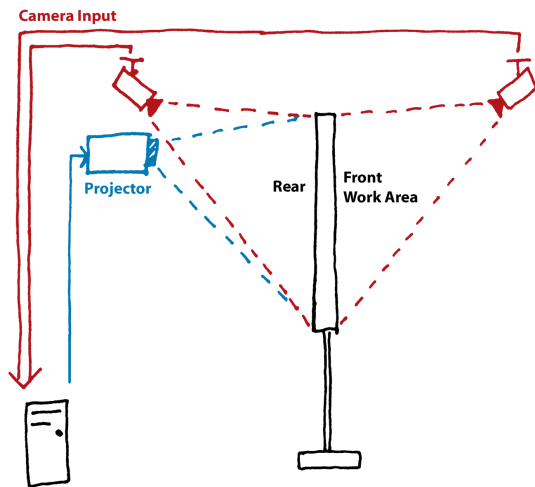


Figure 14. The Outpost system infrastructure.

cameras (Figure 14). In order to obtain an occlusion free view of the board for our difference image algorithm, we followed the metaDESK researchers [27], mounting a video camera behind the board. Interactive frame rates are crucial for this camera. Because the notes are fairly large (three inches square), standard video resolution (640x480) is acceptable for location detection.

Our ink capture task has the opposite set of constraints: we require high resolution for capture but not interactive speeds because the ink capture does not control the board feedback. This suggests a high resolution still camera. Finally, we found that consumer grade web cameras compress images in ways that made computer vision difficult [28].

Matlab Algorithms Prototype

Before building a fully interactive system, we designed and prototyped a set of algorithms in Matlab. Our pipeline performs the following set of operations on $D_{(C2-C1)}$:

- 1) Rectify the perspective camera view of the board plane, bringing the board into a 2D plane using a 3x3 homography map matrix. A homography matrix describes an arbitrary projective transformation. There are more precise algorithms for camera calibration; we chose a homography because it is very fast.
- 2) Threshold the resulting image, producing a three-level image. *Positive pixels* are pixels that have gotten significantly brighter, *neutral pixels* have not changed much, and *negative pixels* have become significantly darker.
- 3) We cluster the pixels using the connected components algorithm, labeling each positive and negative pixel with a cluster ID number. Our system interprets note-sized clusters of positive pixels to be added notes, and note-sized clusters of negative pixels to be subtracted notes.

- 4) We compute the center of mass and the orientation of the note. Inspired by Freeman's work [7], we originally implemented orientation-finding using a second moment algorithm. However, a second moment is not a robust metric for low-resolution captures of objects in the presence of camera noise. The small sample size yields a situation where the maximum second moment is not substantially different from the minimum second moment, and camera and sampling noise combine to make the minimum second moment nearly arbitrary. For this reason, we moved to an expectation-maximization (EM) algorithm [4] that finds the best fitting square on the set of outline pixels of the note. This method is highly robust, even for a highly degraded image outline and small sample size. Theoretically, EM is a more expensive algorithm because it is iterative. In practice, we have found that a very small number of iterations (in our case six) is enough for the solution to comfortably converge.

Interactive Wall Vision Prototype

We combined the system implications from our first vision prototype with the algorithms from our Matlab prototype to produce a wall-scale interactive prototype. This system employs three sensors: a touch sensitive SMART Board, a rear-mounted 640x480 industrial digital video camera, and a front-mounted three megapixel USB still camera (Figure 14) to achieve the multiple person, low-latency input and capture desires that the designers in our study were interested in. This prototype offers an interactive-rate solution for detecting the location of notes. In our current research, we are beginning to address high-resolution capture and unique identification. Our research goal is to leave the representation informal and not semantically analyze ink on the notes, but rather to use the ink as a unique identifier for the note.

Our vision system is written in C++ on top of the Intel OpenCV library [3]. The vision system and the user interface run as separate processes and pass semantic events (e.g., $add[x, y, \theta]$, $remove[x, y, \theta]$) through a socket network connection. Currently both processes run on the same computer, however because of the architecture, it would just as easy to have them be on separate computers. The only reason that we do not do this is network latency between machines.

When the system starts up, it automatically detects the corners of the board. We do this by setting the display entirely white and capturing a frame, setting the display completely black and capturing a frame, computing the thresholded difference image between the two frames, and finding the set of outline pixels. We use EM here as well, finding the best fitting four lines on the set of outline pixels. We use these four corners to automatically compute the homography transform. In this prototype, our camera distortion is small. We originally intended to apply a homography transform to every $C2-C1$ difference image.

We were able to substantially improve the performance of the system by applying the homography transform only to the logical coordinates of the detected note corners. Thus, we only transform four points per note found, as opposed to 640x480 pixels per difference image; a savings of roughly four orders of magnitude.

We also revised our mechanism for finding the location of a note. Because the notes have a sticky stripe across the top, the top edge is flush with the board, straight, and accurate. Sometimes the bottom and side edges curl away from the board, however. (This observation generalizes to images and other paper artifacts taped onto the board.) For this reason, we compute a four-line EM on the note outline, and use the top line to compute orientation and location.

For the most part, this prototype has been very successful. We are currently having difficulties with process scheduling on a one processor machine under Windows 98. Often, either the vision or the interface process will be given use of the processor for extended periods of time. Because both need to run interactively in order to achieve interactive performance, we are moving the system to a two-processor machine with Windows 2000; we believe this will resolve the difficulty.

Additionally, the current vision prototype occasionally makes recognition mistakes. We have found the errors fall into three categories:

- 1) Missed actions: A note was added or subtracted, but the system did not recognize this.
- 2) False positives: Nothing actually happened, but the system issues an action. (False positives are usually triggered by motion of a user's shadow.)
- 3) Location and orientation misreporting: the system properly detects an event, but it is reported at a slightly incorrect location and orientation. Two factors contribute to these errors: 1) The recognition system can treat slightly the wrong set of pixels as the object pixels, 2) The homography transform is not completely accurate.

We are solving these issues in two ways. First, we believe we can improve the vision algorithms. In this arena, we believe that we can better use shape to detect changes. For example, we may be able to assume that a cluster with three straight sides and an oddly shaped fourth side is really a note with a user's hand on the fourth side. (Currently, we may discount this cluster because it is too large to be a note.) We also may benefit from using edge detection as opposed to simple thresholding. And finally, we are considering replacing the homography with a more accurate (albeit slower) representation of camera calibration.

Perhaps more importantly, these errors remind us that intelligent interfaces make mistakes. This has encouraged us to rely more on "simpler" sensing technologies, namely

the board's touch sensor. For example, we have revised our system to first recognize an add note action with the touch sensor. A simple down/up touch on empty board-space adds a candidate note to the model at an approximate location. The model then queries the vision system, asking if any notes were added in the region of the touch. The vision system responds either yes or no. If a note is added, the vision also responds over the socket connection with the precise location and orientation information it found.

CONCLUSION

We propose The Designer's Outpost as a tangible interface for collaborative web site information design. Its 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 a sequence of three design studies, culminating in a study of fifteen professional designers in five design teams working with an interactive prototype. These studies show that the interaction techniques for working with electronic whiteboards should be calm and that there is substantial merit in a system that is simultaneously tangible and virtual. The designers we spoke with encouraged our pursuit of a versioning system for information architectures, and were enthusiastic about fluid transition to tools such as DENIM.

We have also described three vision studies that yield an interactive rate system for finding notes on a large surface. In using two cameras, we are able to separate note detection (which must be interactive) from high resolution capture and identification (which can be slower than interactive).

We are now revising the user interface to provide for more calm interaction and to include facilities for versioning and remote collaboration. We are improving the vision back-end so that the detection is more robust and so the system can uniquely identify notes and capture their ink. Our current goal is to build an Outpost prototype that is robust enough for us to deploy at a design firm for an extended period of time.

ACKNOWLEDGMENTS

Thanks to John Canny for everything from giving us lab space to computer vision help to soldering advice; an HCI renaissance man. Thanks to Marti Hearst for helping to conceive the original concept for The Designers' Outpost. Thanks to Jimmy Lin and Jason Hong for Java programming advice. Thanks to Iwan Ulrich for 1394 driver code. Thanks to Mark Bilezikjian for help running the wall study. Thanks to Gerelee Goltsev, Bjorn Liljequist, Shilpa Sood, and Michael Toomim for early design ideas on properties. Thanks to Smart Technologies for the donation of a SMART Board. Thanks to the support of Trevor Pering, Mark Holler, and Gary Bradski at Intel. This work is supported by a grant from the National Science Foundation. Finally, a big shout out to all the designers who participated in our design studies.

REFERENCES

1. Bellotti, V. and Y. Rogers. From Web Press to Web Pressure: Multimedia Representations and Multimedia Publishing. In Proceedings of *Human Factors in Computing Systems: CHI '97*. Atlanta, GA. pp. 279-286 1997.
2. Beyer, H. and K. Holtzblatt, *Contextual Design: Defining Customer-Centered Systems*. San Francisco: Morgan Kaufmann, 1998.
3. Bradski, G., Open Source Computer Vision Library. 2000.
4. Dempster, A.P., N.M. Laird, and D.B. Rubin, Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 1977. **39**.
5. Dourish, P., W.K. Edwards, A. LaMarca, and M. Salisbury, Presto: an experimental architecture for fluid interactive document spaces. *ACM Transactions on Computer-Human Interaction*, 1999. **6**(2): p. 133-161.
6. Elrod, S., *et al.* Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations and Remote Collaboration. In Proceedings of *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*. pp. 599-607 1992.
7. Freeman, W.T., *et al.*, Computer vision for interactive computer graphics. *IEEE Computer Graphics and Applications*, 1998. **18**(3): p. 42-53.
8. Hynes, C.M. and G.M. Olsen. Quick but not so dirty web design: Applying empirical conceptual clustering techniques to organise hypertext content. In Proceedings of *DIS: Symposium on Designing Interactive Systems 1997*: ACM Press. pp. 159-162 1997.
9. Ishii, H. and B. Ullmer. Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In Proceedings of *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems*. pp. 234-241 1997.
10. Johnson, W., H. Jellinek, J. Leigh Klotz, R. Rao, and S. Card. Bridging the Paper and Electronic Worlds: The Paper User Interface. In Proceedings of *Proceedings of ACM INTERCHI'93 Conference on Human Factors in Computing Systems*. pp. 507-512 1993.
11. Klemmer, S.R., M.W. Newman, and R. Sapien. The Designer's Outpost: A Task-Centered Tangible Interface for Web Site Information Design. In Proceedings of *CHI: Human factors in computing systems 2000*: ACM Press. pp. 333-334 2000.
12. Landay, J.A., *Interactive Sketching for the Early Stages of User Interface Design*, Unpublished Ph.D., Carnegie Mellon University, Pittsburgh, PA, 1996.
<http://www.cs.berkeley.edu/~landay/research/publications/Thesis.pdf>
13. Lin, J., M.W. Newman, J.I. Hong, and J.A. Landay, DENIM: Finding a tighter fit between tools and practice for web site design. *CHI Letters: Human Factors in Computing Systems, CHI '2000*, 2000. **2**(1): p. 510-517.
14. MacLean, A., R.M. Young, and T.P. Moran, Design rationale: the argument behind the artifact. *Proceedings of the SIGCHI conference*, 1989: p. 247-252.
15. Moran, T.P., P. Chiu, W. van Melle, and G. Kurtenbach, Implicit Structures for Pen-Based Systems within a Freeform Interaction Paradigm, in *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*. p. 487-494, 1995.
16. Moran, T.P., W.v. Melle, and P. Chiu. Tailorable Domain Objects as Meeting Tools for an Electronic Whiteboard. In Proceedings of *Conference on Computer Supported Collaborative Work (CSCW) 1998*: ACM Press. pp. 295-304 1998.
17. Moran, T.P., E. Saund, W.v. Melle, A. Gujar, K. Fishkin, and B. Harrison. Design and Technology for Collaborage: Collaborative Collages of Information on Physical Walls. In Proceedings of *Proceedings of Symposium on User Interface Software and Technology 1999*: ACM Press. pp. 197-206, November 1999.
18. Newman, M.W. and J.A. Landay. Sitemaps, Storyboards, and Specifications: A Sketch of Web Site Design Practice. In Proceedings of *Designing Interactive Systems: DIS 2000*. New York, NY. pp. 263-274, August 17-19 2000.
19. Nielsen, J., Test it! Jakob Nielsen's UI Tips. 1998.
<http://www.builder.com/Graphics/UserInterface/ss01a.html>
20. Norman, D.A., *The Invisible Computer : Why Good Products Can Fail, the Personal Computer Is So Complex, and Information Appliances Are the Solution*, 1998.
21. PolyVision, Whiteboard Photo, PolyVision Corporation. <http://www.pixid.com>
22. Rekimoto, J. and M. Saitoh. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In Proceedings of *CHI: Conference Proceedings on Human*

- factors in computing systems 1999*: ACM Press. pp. 378-385 1999.
23. Rosenfeld, L. and P. Morville, *Information Architecture for the World Wide Web*. Sebastopol, CA: O'Reilly, 1998.
 24. Smart, Rear Projection Smart Board, Smart Technologies Inc.
<http://www.smarttech.com/rearprojection/>
 25. Streitz, N.A., J. Geissler, T. Holmer, and S. Konomi. i-LAND: An interactive Landscape for Creativity and Innovation. In Proceedings of *CHI: Conference Proceedings on Human Factors in Computing Systems 1999*: ACM Press. pp. 120-127 1999.
 26. Tesler, L., The Smalltalk Environment, *Byte Magazine*, vol. 6(8): pp. 90-147, 1981.
 27. Ullmer, B. and H. Ishii. The metaDESK: Models and Prototypes for Tangible User Interfaces. In Proceedings of *Proceedings of Symposium on User Interface Software and Technology 1997*: ACM Press. pp. 223-232, October 1997 1997.
 28. Ulrich, I., and Nourbakhsh, I., Firewire Untethered: High-Quality Images for Notebook Computers. *Advanced Imaging Magazine*, 2000: p. 69-70.
 29. Virtual Ink, *Mimio*. Virtual Ink.
<http://www.mimio.com/>
 30. Want, R., K.P. Fishkin, A. Gujar, and B.L. Harrison. Bridging Physical and Virtual Worlds With Electronic Tags. In Proceedings of *CHI: Conference Proceedings on Human factors in computing systems 1999*: ACM Press. pp. 370-377 1999.
 31. Weiser, M. and J.S. Brown, The coming age of calm technology. *PowerGrid Journal*, 1996. **1.01**.
 32. Wellner, P. The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display. In Proceedings of *Proceedings of the ACM Symposium on User Interface Software and Technology*. pp. 27-33 1991.
 33. Wellner, P., Interacting with Paper on the DigitalDesk, *Communications of the ACM*, vol. 36(7): pp. 87-96, 1993.
 34. Whittaker, S. and H. Schwarz, Back to the Future: Pen and Paper Technology Supports Complex Group Coordination, in *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*. p. 495-502, 1995.
 35. Winograd, T. and F. Guimbretiere. Visual Instruments for an Interactive Mural. In Proceedings of *CHI: Conference Proceedings on Human factors in computing systems 1999*: ACM Press. pp. 234-235 1999.