The Designers' Outpost: A Tangible Interface for Collaborative Web Site Design

Scott R. Klemmer, Mark W. Newman, Ryan Farrell, Mark Bilezikjian, 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. These wall-scale paper-based design practices inspired 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. With Outpost, users collaboratively author web site information architectures on an electronic whiteboard using physical media (Post-it notes and images), structuring and annotating that information with electronic pens. This interaction is enabled by a touch-sensitive SMART Board augmented with a robust computer vision system, employing a rear-mounted video camera for capturing movement and a front-mounted high-resolution camera for capturing ink. We conducted a participatory design study with fifteen professional web designers. The study validated that Outpost supports information architecture work practice, and led to our adding support for fluid transitions to other tools.

Keywords

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

INTRODUCTION

Our previous studies into web design [17] 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 consistent with work practice studies across many design and engineering domains [1, 8, 24].

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. We have built a tool, The Designers' Outpost, that supports this practice. It combines the advantages of both paper and electronic media. A video of the Outpost system is available on the web at http://guir.berkeley.edu/outpost/.

Current Physical Practice: Benefits and Drawbacks

The large workspace of a wall or whiteboard offers several clear benefits for collaborative design tasks. Large workspaces permit the representation of large, complex information spaces without the loss of contextual, peripheral information. 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. Paper and walls "make information, any kind of information, tangible, easy to manipulate, and easy to organize" [19]. Collaboration is aided both by the persistence of the design 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. Covi *et al* refer to the work posted on walls in project rooms as "coordination documents" [5].

There are drawbacks to the traditional paper-centric representation. Much of the information exists in the *relationships* between information chunks (Post-it notes). Because structure must be maintained manually, marks that the designers make about the data, such as links or annotations, often fall out of sync with the notes as they 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 is lost.

The designers in our studies also lamented that versioning is not feasible in a paper-only representation. The paper-only work practice also offers few opportunities for remote participants, whether at a desktop down the hall or in a meeting room across the world. Remote users have no way

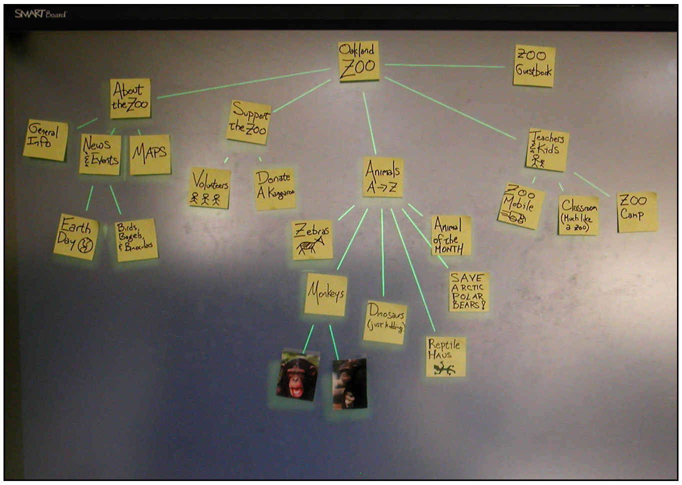

Figure 1 (top left). A web site information architecture using a combination of physical Post-it notes, physical pictures, and virtual links showing relationships between them.

Figure 2 (top right). DENIM, shown here in sitemap view, allows web site design by sketching. As seen here, physical information spaces created in Outpost can be electronically imported into DENIM, serving as baseline sitemaps.

Figure 3 (lower right). In DENIM's storyboard view, designers can continue working with an Outpost sitemap by sketching out the contents of a page.

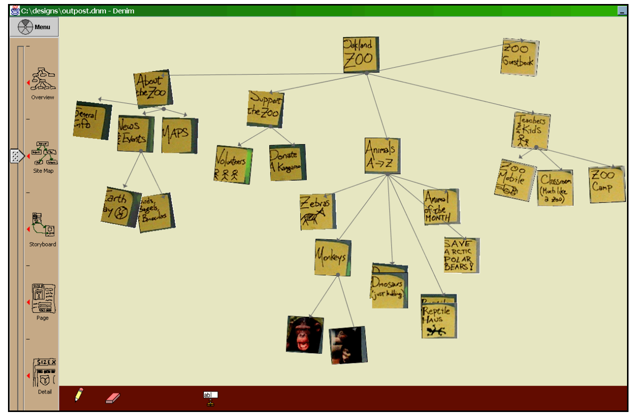

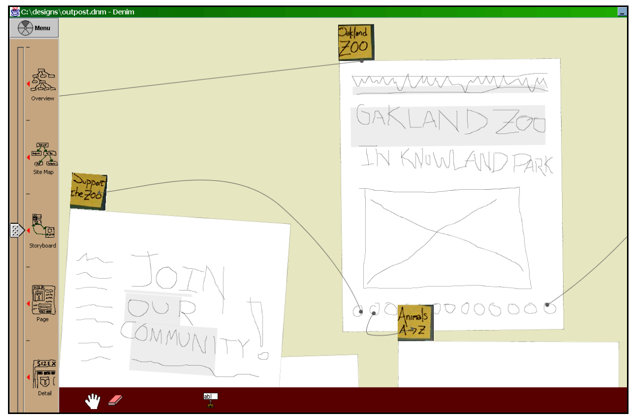

to update, or even access, the information. We also found, as others have, that the transition from the early paper-centric design stages to the later pixel-centric stages is highly problematic [12, 23]. As the site structure is changed during development, the early paper artifact drifts further and further out of date.

Supporting and Extending Practice with Outpost

The Designers' Outpost (see Figure 1) 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. Users have the same fundamental capabilities in the Outpost system as in a non-computational paper-based system; one can create new pages by writing on new Post-it notes and organize a site by physically moving Post-it notes around on the board. Thus, paper in the physical world becomes an input device for the electronic world. A rear-mounted projector outputs electronic information onto surfaces in the physical world. Through its electronic capture of designs, our system supports the transition from this early representation to later electronic tools, such as DENIM [13].

Outpost is part of our group's research on informal user interfaces [12]. Informal user interfaces support natural human input, such as speech and writing, while minimizing recognition and transformation of the input. These interfaces, which *document* rather than *transform*, better support a user's flow state. Unrecognized input embraces nuanced expression and suggests a malleability of form that is critical for activities such as early-stage design.

The rest of the paper is organized as follows. First, we discuss work related to Outpost in the areas of early stage web design and tangible user interfaces. Next, we describe the interaction techniques that give the core functionality of our system. Following this is a description of the design studies that we carried out to inform and validate our approach. A detailed discussion of the implementation of our computer vision algorithms follows, and finally, we describe our future work and conclusions.

RELATED WORK

Our research is inspired by previous work in two areas—early stage web site design and tangible user interfaces. We describe these two areas next.

Web Site Design: Tools and Practice

The goal of our earlier investigation into web design [17] was to inform the design of systems to better support actual practice. The study, consisting of interviews with eleven professional web site designers, provided us with three important insights. First, designers create many different intermediate representations of a web site. Examples of pervasive and significant intermediate artifacts include sitemaps, storyboards, page schematics, and mockups. 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 and visual design, each of which has its own tools, products, and concerns. Information architects are mainly concerned with the information and navigation design of a web site. Visual designers typically focus on interaction and graphic design. We found that information architecture is not well supported by current software tools. For example, the study found that sitemaps were regularly generated by placing Post-it notes on walls.

The results of these studies provided motivation for Outpost, and they also provided the impetus for the development of DENIM, a sketch-based tool supporting information and navigation design of web sites [13]. DENIM (see Figures 2 and 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 *page schematic* representations of a web site. While DENIM supports authoring sitemaps, it is best suited for storyboards and page schematics. Outpost targets sitemaps. Also, DENIM was designed as a single-user interface, whereas Outpost was designed for collaborative work.

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" [1]. This study helped motivate our interest in combining the physical and electronic worlds to gain these benefits.

Tangible User Interfaces

Tangible user interfaces (TUIs) allow "users to 'grasp and manipulate' bits in the center of users' attention by coupling bits in with everyday physical objects and architectural surfaces. Tangible Bits also enables users to be aware of background bits at the periphery of human perception using ambient display media" [9].

Desks are one area where research has yielded compelling TUIs; a seminal example is Wellner's DigitalDesk. The DigitalDesk used ceiling mounted cameras to track documents and hands on a physical desktop, with a ceiling mounted projector to electronically augment the real desk [23]. Outpost continues in the direction the DigitalDesk began, by augmenting physical paper with electronic information. MIT's Tangible Media Group later created the metaDESK [21], a digital desk employing tangible interfaces as the controls for and views of a map of the MIT campus. Outpost employs the metaDESK technique of tracking objects with a rear camera; it differs in that it targets an existing professional design practice.

Several influential systems for wall-scale interaction (both physical and electronic) have come out of Xerox PARC, including the LiveBoard [6, 15], and more recently, Collaborage. "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" [16]. Collaborage's computer vision capture of paper on walls inspired our work. One drawback of the Collaborage capture system is its long latency (8-10s on average). Outpost improves on this, with a location recognition latency of ~200ms, and an average capture latency of ~1.5s. This improvement is primarily due to our two-camera hardware approach, and software built on top of OpenCV, a highly optimized vision toolkit [4]. Outpost also incorporates other forms of input using styli and physical tools.

Several other researchers are investigating interaction techniques for large electronic display surfaces [20], the combination of these surfaces with physical objects [18], and multimodal interaction with paper [14]. This body of work motivates the concept that, for many tasks (especially in collaborative situations), manipulating physical objects on large surfaces is an intuitive and effective means of performing computer input.

OUTPOST INTERACTION TECHNIQUES

The physical, direct manipulation interaction techniques in Outpost provide for authoring content with standard pens on Post-it notes. The system tracks notes as users physically add, remove, and move them around the board. The system makes no attempt to recognize the content of the notes, tracking only their positions and relationships.

Outpost supports the following interaction techniques for working with paper on the board. We have combined these physical interactions with interactions that are better suited to an electronic medium, such as digital ink annotation and specifying relationships using virtual arrows.

Adding Notes: Users can write on a note with a standard pen and add it to the board. Our vision system recognizes it and updates its internal understanding of the board.

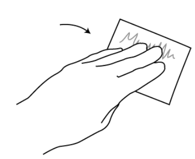

Creating Links: To link a pair of notes, the user draws a line from one note to another with the board stylus.

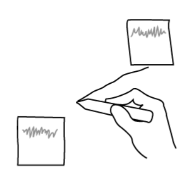

Removing Notes: To delete a note and its associated links, the user pulls the note off the board.

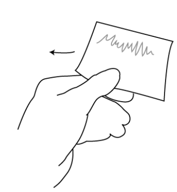

Moving a Note: To move a note and its links, the user picks it up and places it at its new location. This provides for a lightweight means of keeping the electronic data and the physical object coupled.

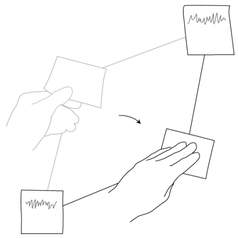

Context Menus: Tapping a note invokes an electronic context menu, enabling the manipulation of the electronic properties embodied by physical objects. *Sticky* replaces a physical note with an electronic image of the note. *Delete* removes a note (useful if the vision system misses a physical removal).

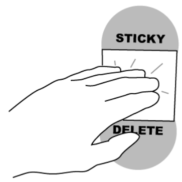

Outpost offers three physical tools for manipulating electronic content. We added ink annotation/erasing and export because of feedback we received from our study.

Freeform Ink: In addition to being a space for interacting with physical Post-it notes, Outpost is also an electronic whiteboard, supporting freeform drawing using board styli.

Move Tool: A physical move tool provides a means of interacting with the system after the physical content has become electronic, retaining haptic direct manipulation.

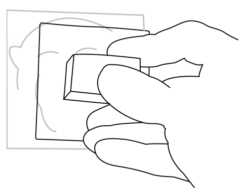

Physical Eraser: Working like a normal whiteboard eraser, the Outpost eraser removes ink on the board. It operates semantically, deleting each stroke it passes over.

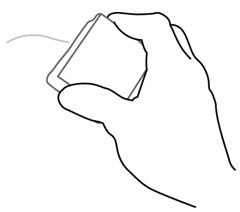

Two primary benefits to structured capture of informal artifacts are 1) later recall, and 2) export to other tools; saving enables both of these.

Saving the Board: Users can press save to save the board state to disk. Then they can open it later in DENIM or Outpost.

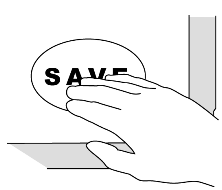

One important part of the Outpost visual design is that the board's background is black. Because the board does not emit light except in areas where the user has authored content, it minimizes the feeling of having a glowing presence in the room.

DESIGN STUDIES

To explore the viability of combining physical and electronic representations for web site information architecture, we undertook a series of three design studies. First, we evaluated the basic concept with a paper prototype. Next, we built mock-ups that envisioned the electronic augmentation of the physical design. Lastly, we created a wall-scale prototype for a set of participatory design sessions with professional designers.

Our paper prototype validated the general approach [10]. It indicated a need to minimize the extra user effort required to use the system, and encouraged us to allow the interaction to be as freeform as possible. Our next prototype [11], on a drafting size digital desk, fleshed out the interaction techniques using paper and computer images. It quickly became evident that a digital desk is too small a space for professional web site information architecture. 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.

Interface Prototype

Our interactive prototype for the study was implemented as a Java application running on a rear-projected 72" diagonal touch-sensitive SMART Board with a 1280x1024 LCD projector. With this prototype, we recognized the location of notes on the board using the board's touch sensor. Drawing a line from one note to another with the board stylus creates a link. The stylus is also used for creating freehand electronic ink on the board (see Figure 4).

Tapping on a note invokes a context menu (see Figure 5) that in this prototype 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.

Professional Design Study

We ran five design sessions with between two and five designers per session, for a total of fifteen professional web designers. 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

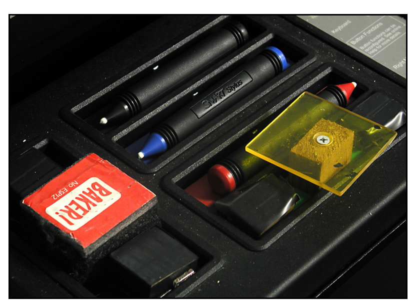

Figure 4. The board's tool tray: styli for drawing electronic ink, a clear plastic square for moving electronic content, and the eraser. *(Only the pens were available during the design study.)*

Figure 5. Tapping on a note invokes an electronic context menu for physical content.

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

Table 1. The five study groups: their size and primary role.

were composed of information architects, two groups were visual designers, and one group had individuals performing both roles (see Table 1).

Each session lasted roughly two hours. 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 to 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 6 and 7 are stills from those recordings.

Design Findings

Our findings from this participatory design study offered insight into the designers' collaborative work processes and suggested an appropriate interactivity model.

Existing Work Processes

Every participant currently works with groups on whiteboards early in the web site design process. The information architects all said that they currently create sitemaps by placing Post-it notes on the board, while the visual designers sketch page designs directly on the board. Capturing whiteboard designs was highly valued by all five teams. Three of the design teams currently use a digital camera for documenting their work, one uses a whiteboard capture device (the Virtual Ink Mimio), and one assigns a scribe to save information from design meetings. Also, every designer said that they currently use either the Visio or Inspiration structured drawing software for creating sitemaps. Sitemaps can get quite large; one firm said that two to three hundred nodes is typical.

Interactive Board Work Process

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

Phase I: Brainstorming

First, the designers brainstormed, quickly putting 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. Designers did not eliminate ideas nor link concepts together into any formal structure at this stage.

One designer commented that Outpost 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 from scratch with a new tool.

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.

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.

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 Ink

After the designers created a rough cut of a sitemap, 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.

A key design implication taken from this phase is the need for associating freeform ink with individual notes. 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." Two groups

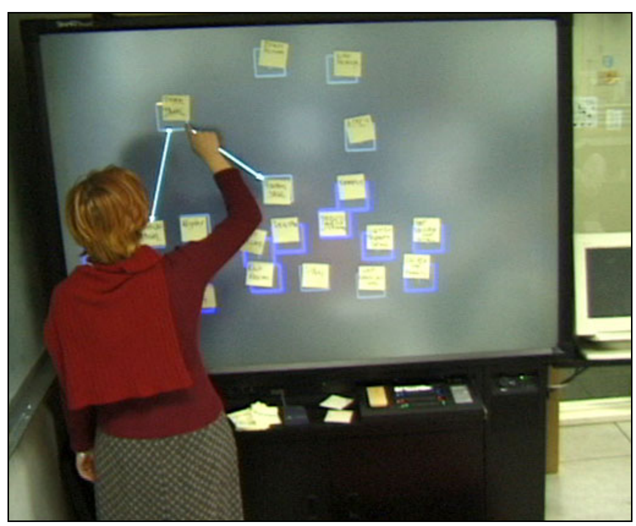

Figure 6. This is an example of the *facilitator* style; one person remains at the board guiding the group's process.

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 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?"

We observed two styles of interacting with the board. In the *facilitator* style, one person, usually the senior-most individual, stands at the board (see Figure 6). The entire group discusses the site, but as the discussion progresses, the facilitator creates notes that synthesize the discussion content. 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 (see Figure 7). This paradigm affords each person their own paper "input device," a working style we had not considered but that the designers regularly engaged in. 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. This result encouraged us to complete a computer vision system. Vision lends itself both to simultaneous input and to rich sensing capabilities (e.g., object size, color, orientation, and capture of its contents).

The post-test questionnaire asked: "How likely is it that you would integrate Outpost as a regular part of your web site design practice?" Participants rated their response on a five point Likert scale. Four participants rated the system the top value, *very likely*. Eight gave the second value, *somewhat likely*. Three gave the fourth value, *somewhat unlikely*; we believe the primary reason for these participants' negative feelings was the distracting visual feedback in this prototype. The current system is much calmer.

Figure 7. This is an example of the *open board* style; all participants directly express their ideas in the artifact.

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.

DESIGN IMPLICATIONS

This study underscored several important points about how calm [22] 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 was not useful, as they already knew the layout of the notes.

Furthermore, the group, note outline, and menu feedback was considered distracting during brainstorming. 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], a sketch-based GUI design tool.

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 into more explicitly creating a sitemap. Ideally, this functionality should be available but not automatic.

Sweet Spot on the Tangible/Virtual Spectrum

We have seen virtual wall-scale interfaces [15] as well as physical ones [16]. There are appealing aspects to both interaction paradigms; one of our goals with Outpost is to leverage the advantages of both.

Fluidity and Physicality

This series of design studies 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.

One of our concerns about a tangible/virtual interface was that designers would find it tedious to remove physical

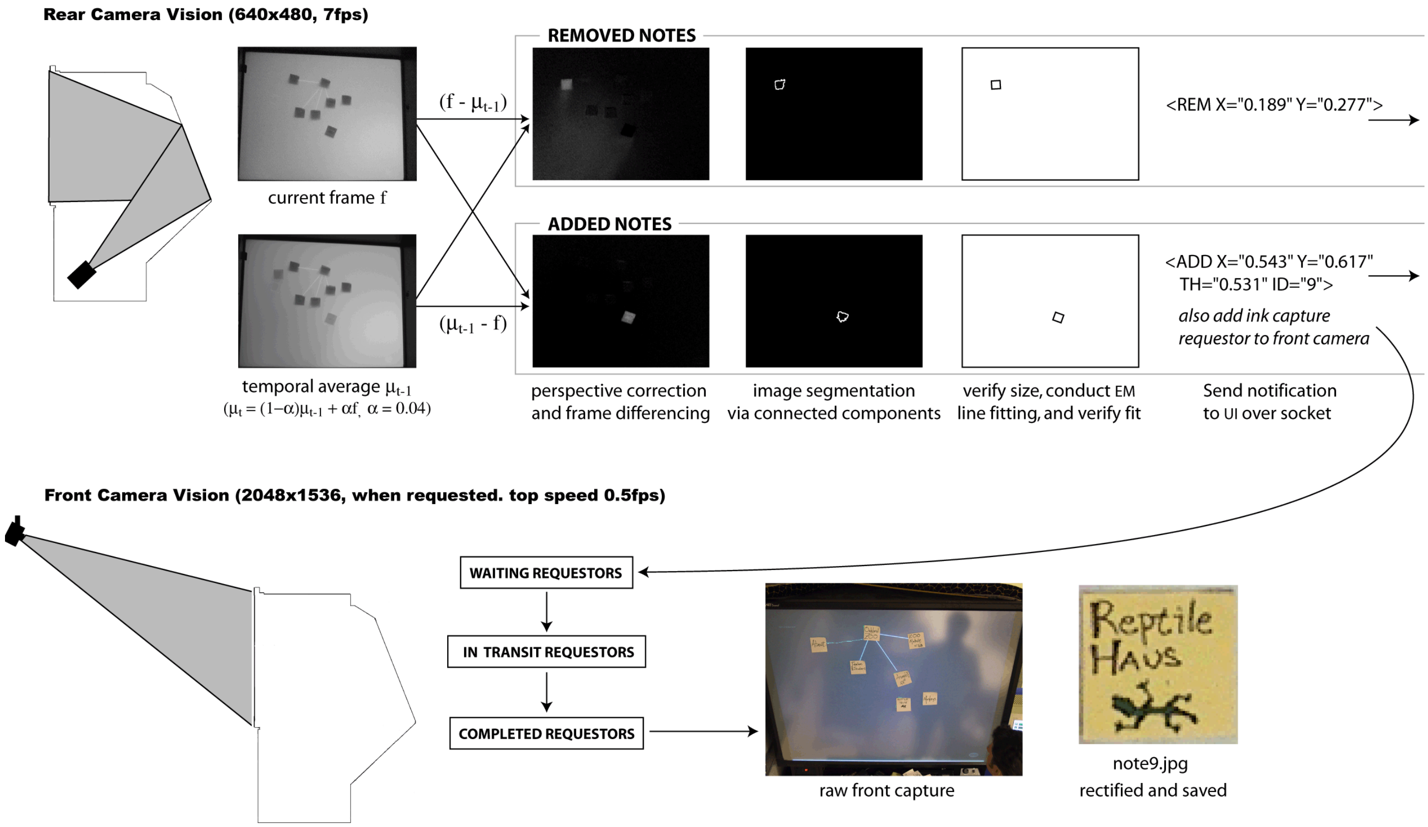

Figure 8. The Outpost vision pipeline at a frame where one note ("Reptile Haus") was added and another was removed.

objects from the board as they began to flesh out the design in detail. Somewhat surprisingly, the designers did not mind removing the physical objects. On the contrary, removing them was a natural signifier of a shift in the design process. To facilitate this, we now include functionality to make all notes electronic in one command.

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 system 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.

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 Design Process

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.

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 pen-based interface is intended for a single designer working at a PC. Outpost is most appropriate for creating sitemaps; whereas DENIM becomes more relevant when the design team starts to storyboard the specific pages and create schematics. The current Outpost system and DENIM read and write the same XML file format. This makes it possible for an individual to "save out a wall" from a collaborative design session, and then flesh out the design in a personal space. To support this, we augmented DENIM to handle images as page labels (see Figures 2 and 3). We plan to add networking socket code that will send XML design changes to registered Outpost/DENIM clients, that will enable real-time and asynchronous participation by remote members of a design team.

Long projects magnify the benefits of having a sitemap artifact remain in use throughout the entire design cycle. For example, 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. Through its electronic capture functionality, we hope Outpost will help design teams with such long term projects.

While an early interest in Outpost was to provide interactive support for information architecture design sessions, designers in our most recent studies found additional fruitful directions for our research. They encouraged us to refocus our efforts toward a more documentary interface, supporting free ink electronic annotations to sitemap pages, versioning of design artifacts, fluid transitions to tools such as DENIM, and supporting collocated and remote collaboration. We also found that the system is much more appropriate for information architects than for visual interface designers.

IMPLEMENTATION

The Outpost system consists of two main components. The interface component handles stylus, physical tool, and touch input on the board, and provides graphical feedback to the user. The computer vision component tracks and captures physical Post-it notes and pictures.

Physical Tools and Graphical Display

The physical tools input and graphical feedback are implemented in Java using SATIN, a toolkit for informal pen-based user interfaces [7], and the SMART Board SDK. In Outpost, we make use of SATIN's extensive support for ink handling, gesture recognition, and rendering. Free ink in Outpost is captured and saved as a *stroke* primitive. We use a tap *interpreter* for invoking context menus on existing notes. We also use a gesture *interpreter* for drawing links between pairs of notes.

The SMART Board's tool tray consists of four pen tool slots and one eraser tool slot. The hardware detects the presence of the tools via a photometer in each slot. The hardware defines the active tool to be the tool most recently removed from the slot. The tools themselves are passive. The SMART Board SDK uses callbacks to inform registered applications of the current tool. We use this mechanism to know when the move tool or the eraser is active.

Computer Vision Infrastructure

We now present the computer vision system that we have created for building wall-scale tangible interfaces. We use computer vision to precisely locate and capture the information on Post-it notes and images that users place on the board. This system can handle simultaneous input; essential for collaborative design. Our vision system is written in C++ on top of OpenCV [4], a highly optimized library of computer vision and image processing primitives. The vision system runs as a separate process, passing semantic events (e.g., ADD [x, y, θ, ID], REMOVE [x, y]) to the Outpost UI through a socket network connection.

In building this system, we realized that locating a note is a completely separate problem from capturing a note's ink. Dividing our task in this way enabled us to realize that the system architecture should have two cameras (see Figure 8). To obtain an occlusion free view of the board, we followed the metaDESK researchers [21] by mounting a video camera behind the board. Interactive frame rates are crucial for this camera. Because notes are fairly large (three inches square), standard video resolution (640x480) is acceptable for location and orientation 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. This two camera approach obviates the need for a mechanical pan/tilt/zoom camera and image stitching algorithms.

This system offers interactive rates (~7 frames per second) for detecting the location of notes with the rear camera, combined with background high-resolution capture (~1.5 second latency) for virtual display and transitioning to DENIM. This design achieves the multiple person low-latency input and capture that the designers in our study were interested in. One way to think about the board capture is as a direct manipulation scanner. One operation, placing a physical document on the board, specifies both the location of the document and that the document should be captured.

Interactive Vision Techniques

There are several processing steps that we perform with each new image from the rear camera. First, we employ spatial and temporal filtering techniques that help alleviate problems due to camera noise and lighting changes. This is a common and effective technique in many computer vision applications. Our temporal filtering computes a running image average $\boldsymbol{\mu}_t$, by averaging in each new frame $\mathbf{f}_t$ with weight $\boldsymbol{\alpha}$ (in our case, 0.04).

Each frame, we rectify the perspective camera view by bringing the board into a 2D plane using a projective transform matrix. There are more precise algorithms for camera calibration; we chose a simple perspective warp because it is fast, and works well for our purposes

Next, we construct two thresholded difference images: added notes are found in the $(\boldsymbol{\mu}_{t-1} - \mathbf{f})$ image and subtracted notes in the $(\mathbf{f} - \boldsymbol{\mu}_{t-1})$ image. We segment the two binary images using the connected components algorithm, finding note-sized components from changed pixels.

After segmentation, we compute the center of mass and the orientation of the note components. We use an expectation-maximization (EM) algorithm (for a good overview, see [3]) as a robust method for finding the best fitting square on the set of outline pixels of the note.

As a final step, we require that added objects be found in the same place for two consecutive frames. We added this step to reduce false positives to a negligible level. At the completion of this vision pipeline, we send the semantic information about board state changes over the socket to the Java user interface.

To minimize our computational overhead, the front camera only takes a picture when the rear camera has detected that a new note **n** has been placed on the board. For each **n** we add a *requestor* to the front camera's request queue with [x, y, θ, ID] as the location to capture. The camera takes a picture when it is next available. The system corrects for perspective skew upon receiving the picture. For each *requestor*, the system saves the rectified area of the board as a JPEG file. This method insures that note capture will complete soon after the note is placed on the board (capture completion time is bounded by twice the image transfer time). It also enables multiple notes to be captured from a single image.

Outpost can be run in a calibration mode, where it automatically detects the corners of the board and saves the calibration parameters to a file. We do this by capturing a frame of an entirely black board, capturing a frame of the board with a projected white outline, computing the thresholded difference image between the two frames, and finding the set of outline pixels.

Discussion

We designed the vision system to be highly robust at finding notes. The occasional errors we do have fall into three categories:

Missed actions: There are very few cases where the vision misses an add or remove action (~1%). As visual feedback, the UI displays a faint shadow around recognized objects. When objects are not recognized, the user must perform the action again.

False positives: Rarely (2-3%), the system reports an action that did not happen; this is nearly always because a user's closed hand is perceived to be a note. As a UI solution, we offer the delete option on the context menu.

Location and orientation misreporting: In this system, there are two kinds of accuracy: *resolution* and *calibration*. Our system performs adequately in both regards. As a point of comparison, most of the time our vision system is of higher accuracy than the board's capacitance sensor; a more sophisticated camera model could improve this further.

Computer Vision as a Sensing Technique

We have found computer vision to be an appropriate methodology for this task because it can provide automatic, untethered, and untagged tracking and capture of artifacts users place on the board. We can characterize many sensing systems as being either of the *AI* variety, "If it has some properties of a duck, it is a duck," or of the *tag* variety, "If it wears the tag we told ducks to wear, then it is a duck."

Tags can often be more robust, more accurate, and/or computationally cheaper; they are appropriate when the same object will be reused many times with the same system. The SMART Board's pen tray is an example of tagging. The main drawbacks to tagging are the monetary cost and a deployment time that is proportional to the number of objects. Because of this, there is a barrier to "suiting up" objects for use in the system.

While it can sometimes be less robust, more computationally intensive, and more laborious to develop, *AI* enables informally appropriating and including any members in the class of objects that are being sensed. This is ideal for the Post-it notes and pictures in our system, where free integration of paper artifacts is critical in supporting the flow state of a design session.

FUTURE WORK

Our earlier study into web design practice showed that designers desired a way to manage different versions of design ideas [17]. 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 manually photograph, photocopy, or scan an artifact to save the state, or abandon the current state and keep working. We have recently begun researching interfaces for the capture and access of design history and versions.

The comments made on a design are often more valuable later than the design itself. We are currently investigating how to capture and search the design rationale that is often given verbally while working in front of the board. We also plan to deploy Outpost at a design firm for an extended period of time to see how well it works in practice.

CONCLUSION

We have presented The Designers' Outpost, a tangible interface for collaborative web site information design. Its functions are informed by observations of real web site design practice, providing many of the affordances of current paper-based practice while offering the advantages of electronic media.

Outpost is implemented on top of a vision system that yields an interactive-rate system for robustly finding notes on a large surface. These results and those of other researchers show that computer vision is an effective methodology for informal, collaborative interaction with physical media on walls.

We validated our design with fifteen professional designers, showing that electronic whiteboards should be calm and that there is substantial merit in a system that is simultaneously tangible and virtual. The designers were enthusiastic about using Outpost and achieving the fluid transition from artifacts on walls to single-user tools such as DENIM.

ACKNOWLEDGMENTS

We thank Michael Thomsen for help with the physical tools; Ethan Phelps-Goodman for help with recent software improvements; Francis Li for video editing wizardry; John Canny and Marti Hearst for insightful discussions; Jimmy Lin and Jason Hong for Java programming advice; and Raecine Meza for help with the paper prototype. Thanks to Smart Technologies for donating a SMART Board, and to Intel for donating two workstations. This work is supported by National Science Foundation grant #IIS-0084367. Finally, a big shout out to all the designers who participated in our studies.

REFERENCES

1. Bellotti, V. and Y. Rogers. From Web Press to Web Pressure: Multimedia Representations and Multimedia Publishing. In Proceedings of *ACM Conference on 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. Bishop, C.M., *Neural Networks for Pattern Recognition*: Oxford University Press, 1995.
4. Bradski, G. Open Source Computer Vision Library, 2001. http://www.intel.com/research/mrl/research/opencv/
5. Covi, L.M., J.S. Olsen, E. Rocco, W.J. Miller, and P. Allie. A Room of Your Own: What do we learn about support of teamwork from assessing teams in dedicated project rooms. In Proceedings of *Cooperative Buildings: Integrating Information, Organization and Architecture: First International Workshop, Co'Build '98*. Darmstadt, Germany, 1998.
6. Elrod, S.*, et al.* Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations and Remote Collaboration. In Proceedings of *ACM CHI'92 Conference on Human Factors in Computing Systems*. pp. 599-607, 1992.
7. Hong, J.I. and J.A. Landay, SATIN: A Toolkit for Informal Ink-based Applications*.* UIST 2000, ACM Symposium on User Interface Software and Technology, *CHI Letters*, 2000. **2**(2): p. 63-72.
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 *ACM Conference on Human Factors in Computing Systems: CHI '97*. pp. 234-241, 1997.
10. Klemmer, S., M. Newman, R. Farrell, R. Meza, and J.A. Landay, *A Tangible Evolution: System Architecture and Participatory Design Studies of the Designers' Outpost*. Technical Report UCB//CSD-00-1116, UC Berkeley, Berkeley, CA, November 2000.
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 *Extended Abstracts of CHI: Conference on Human Factors in Computing Systems*: ACM Press. pp. 333-334, 2000.
12. Landay, J.A. and B.A. Myers, Sketching Interfaces: Toward More Human Interface Design*. IEEE Computer*, 2001. **34**(3): p. 56-64.
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 2000, Human Factors in Computing Systems, *CHI Letters*, 2000. **2**(1): p. 510-517.
14. McGee, D.R., P.R. Cohen, and L. Wu. Something from nothing: Augmenting a paper-based work practice via multimodal interaction. In Proceedings of *Designing Augmented Reality Environments*. Helsinor, Denmark: ACM Press. pp. 71-80, April 12-14, 2000.
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., E. Saund, W. van Melle, A. Gujar, K. Fishkin, and B. Harrison. Design and Technology for Collaborage: Collaborative Collages of Information on Physical Walls. In Proceedings of *Symposium on User Interface Software and Technology 1999*: ACM Press. pp. 197-206, November, 1999.
17. 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.
18. Rekimoto, J. and M. Saitoh. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In Proceedings of *ACM Conference on Human Factors in Computing Systems: CHI 1999*: ACM Press. pp. 378-385, 1999.
19. Rettig, M. Walls beyond whiteboards. *AIGA Advance for Design Summit*, Summer, 2000. http://www.enteract.com/~marc/rettig.walls.72dpi.pdf
20. Streitz, N.A., J. Geissler, T. Holmer, and S. Konomi. i-LAND: An interactive Landscape for Creativity and Innovation. In Proceedings of *ACM Conference on Human Factors in Computing Systems: CHI 1999*: ACM Press. pp. 120-127, 1999.
21. Ullmer, B. and H. Ishii. The metaDESK: Models and Prototypes for Tangible User Interfaces. In Proceedings of *ACM Symposium on User Interface Software and Technology: UIST 1997*: ACM Press. pp. 223-232, October, 1997.
22. Weiser, M. and J.S. Brown. The Coming Age of Calm Technology, 1996. http://www.ubiq.com/hypertext/weiser/acmfuture2endnote.htm
23. Wellner, P., Interacting with Paper on the DigitalDesk, *Communications of the ACM*, vol. 36(7): pp. 87-96, 1993.
24. 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.