

Interactive Gigapixel Prints: Large Paper Interfaces for Visual Context, Mobility, and Collaboration

Ron B. Yeh, Joel Brandt, Scott R. Klemmer, Jonas Boli, Eric Su, Andreas Paepcke
Stanford University HCI Group, Computer Science Department
Stanford, CA 94305-9035, USA
Contact: ronyeh@cs.stanford.edu

ABSTRACT

Creative professionals in many domains use large paper artifacts for their size and resolution, mobility, and flexible support for collocated collaboration. In this paper, we suggest that large-format printers can be thought of as displays with high *spatial* resolution but low *temporal* resolution. Conversely, digital displays have a lower spatial resolution but higher temporal resolution. This work introduces Interactive Gigapixel Prints (GIGAprints), a set of techniques and tools for imbuing large paper media with interactive input (through the location-specifying Anoto pattern) and using the complementary affordances of paper and digital together as an ensemble. We introduce a design space and interaction techniques for GIGAprints. We also describe the design tools for creating these prints, introducing a software architecture that provides an event-based programming model for ensembles that integrate paper and digital artifacts to provide a unified user experience. Finally, we present the results of a first-use study of three applications (12 participants), demonstrating that GIGAprints supports our goals of context, mobility, and collaboration.

ACM Classification Keywords

H.5.2: User Interfaces—*input devices and strategies; interaction styles; prototyping*. H.5.3: Group and Organization Interfaces — *collaborative computing*.

Keywords

Augmented paper, large displays, device ensembles.

INTRODUCTION

For centuries, people have used large paper information graphics as cognitive aids in tasks such as navigation, design, and architecture. This paper was inspired by the rich information environments present in large prints—from the high resolution photographs from the Gigapxl Project [10], through Kerouac’s *On the Road* [23] manuscript (typed on one roll of teletype paper 120’ long), to Imhof’s cartography [21], and Minard’s flow maps [40].

However, the value of paper tools—due to their reliability, mobility, flexibility, readability, and low cost—are diminished by the fact that paper lacks the computing capabilities we depend on (*e.g.*, multimedia and search). We introduce Interactive Gigapixel Prints (GIGAprints), computer controlled large paper displays that afford pen-based input. This research is both a *pragmatic* solution to the current affordance differential between paper and digital media, and

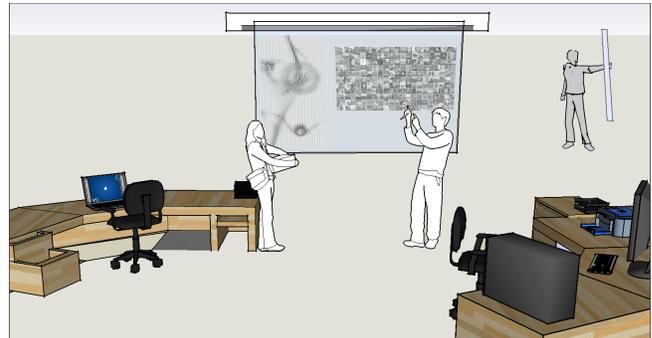


Figure 1. Interactive Gigapixel Prints provide a window into the user experience of future technologies, such as electronic paper. GIGAprints are large, flexible, mobile, and support collaboration.

an *epistemic* solution that provides a window into the user experience of future technologies such as electronic paper (see Figure 1). The GIGAprints platform enables designers to leverage centuries of study in printed information graphics [6, 53], and recent research into interactive visualization [7, 57].

Prior research (*e.g.*, [15, 37]) has shown that the physical affordances of paper—grasping, carrying, and folding—have encouraged its significant use, even in today’s technology-enabled world. Of particular relevance to this discussion is the observation that paper tools are readily accessible to a variety of domains; users neither have to “boot up” a paper tool, nor learn how to operate it. Additionally, paper prints are more mobile and reconfigurable; one only needs to roll up a print to transport it. Finally, content and annotations rendered onto sheets of paper are permanent. With respect to GIGAprints, the abundance of paper artifacts and their flexible arrangement in space enables users to author for themselves information rich environments that lower the expected cost of finding common information [24]. Additionally, the interactions in GIGAprints offer an appealing benefit comparison with traditional paper: at its core, using a GIGAprint *is* using paper, and the augmentation provides benefit on top of that. This means that with GIGAprints, users are *never worse off* than using paper alone. Finally, paper and pen provide a low threshold interface for handwritten annotations.

To use a GIGAprint, one or more users (each with a digital pen) interact with a large printed visualization. This print is augmented with the Anoto location-specifying pattern to enable pen-based input [1]. Output is provided through aux-

iliary displays such as LCDs and projectors, or audio displays such as a wireless earpiece. To build a GIGApriint, the designer/developer creates the visual design in PDF format (*e.g.*, from Adobe Illustrator), and adds active regions using the GIGApriint designer, an Acrobat plugin. Finally, he adds auxiliary devices and event handlers to the active regions through Java programming. The GIGA print design pipeline lowers the threshold for building interactive prints; and allows the process to be divided between multiple specialists: domain experts, designers, and developers.

The rest of the paper is organized as follows. First, we describe the results of interviews with biologists and architects. To illustrate the possibilities, we provide seven example GIGApriint applications. We then describe the interaction techniques and design space. Following that is a discussion of the technical contributions of the work, including the software architecture, programming model, and creation pipeline for GIGApriints. We then offer a broader discussion of the user experience and the results of a first-use study with 12 participants. The next section discusses related work, and we conclude with ideas for future work.

FIELD STUDY

To explore opportunities for design, we conducted interviews and observations with 12 interviewees in two disciplines: biology and architecture.

Our study of biologists draws on earlier work observing that field biologists considered paper notebooks an essential tool in their work [44]. The notebook facilitates a biologist's need to capture field data and transform it into an analyzable form. We continued this fieldwork, and found that in addition to notebooks, large-format prints were used in multiple tasks. In particular, large maps were used in research planning, data capture, and communication (see Figure 2). For planning, the Jasper Ridge camera trapping project used maps to determine where to place cameras to automatically take photos of passing mammals. For data capture, one biologist uses a large map in the field to capture her ant survey data. She writes observations on the map, effectively *geo-tagging* her research data as she travels between sites. Finally, we have observed biologists using large maps in the field to teach and communicate with incoming researchers.

For architecture, we conducted structured interviews with seven professionals in construction—six architects and one contractor. From this work, we found that GIGApriints can help at multiple stages, particularly in the early design and communicating stages (see Figure 2); this confirms earlier work in investigating tablet support for architects [7]. Currently, architects leverage paper-based visualizations for



Figure 2. Left to Right) The ant biologist's 4-foot by 3-foot map is folded up before it is taken out to the field. Maps are also used by biologists for teaching and communicating with students and new researchers. Architects leverage prints in multiple ways, such as explaining ideas to clients with this perspective view.

ideation and communication. Initial designs are sketched on paper; once the designs have been redrawn in CAD, they are printed onto large "whiteprints." The whiteprints are tacked onto walls, where teams can gather to draw revisions, either directly on the prints or on translucent information overlays. One interesting finding was that some veterans do not know how to use CAD; instead, they create on paper and give their designs to junior architects to input them into CAD. Another interesting finding was that four architects use Google Sketchup, an early-stage 3D modeling program; they report that it is more intuitive than AutoCAD, but lacks CAD's sophisticated capabilities.

INTERACTING WITH GIGAPIXEL PRINTS

The Gigapixel Prints ensemble leverages *wide-format inkjets* for high-resolution display, *digital pens* for user input on the printed surface, and *digital displays* for interactive content. Specifically, we use a 44" Epson Stylus Pro 9800 printer, the Nokia SU-1B Anoto digital pen, and myriad digital displays, including projectors and handhelds. Anoto pens determine their location by tracking a dot-pattern printed on a paper surface; and the SU-1B supports real-time streaming of this data.



GIGApriints interactions span multiple users and multiple devices. Any single interaction with a GIGApriint may require the input of multiple users or devices. This *device ensemble* approach was introduced by Schilit *et al.* [36]; we extend this term to include paper prints, digital pens, and input from multiple users.

Pen-based interaction with GIGApriints comprises four categories of actions: select, manipulate, associate, and transport (see Figure 3). *Select* is an input operation by which the user specify/designates content of interest. *Manipulate* actions are those that modify the specified content. *Associate* actions enable the creation of symbolic links: these links can be between physical elements, digital ones, or both. Finally, the sending of data to the receiving device is accomplished through *transport* operations. This taxonomy for

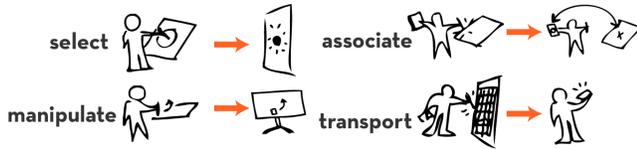


Figure 3. The user initiates a *select* to designate one or more items of interest on the receiver. A *manipulate* interaction enables the sender to modify content on the receiver. *Associate* creates a link between two pieces of information. *Transport* describes the movement of content between the sender and receiver devices.

GIGApriint interactions is adapted from prior work in defining graphical input primitives [9] and graphical input architectures [18, 33]. We subsume Foley & Wallace’s *pick, button, locate*, and *valueate* into our first two ensemble interactions (*select* and *manipulate*), and contribute *associate* and *transport* operators for reasoning about interactions that involve multiple devices (including paper). To demonstrate these interactions, and to attempt to address our goals of visual context, mobility, and collaboration, we developed seven GIGApriint applications.

Visual Search—Photo Wall and Twistr

We have developed a Photo Wall, inspired by the value of high-visibility artifacts to publishing houses [3] and design studios [27]. The application is a wall-sized print that hangs as an always-on content display. Users *select* a photo (by tapping the pen to the region under each photo) to *transport* it to their mobile device (prototyped with an OQO handheld). For organizations such as news agencies that operate on short time cycles where finding the right asset for the front page can be difficult, a large, high-resolution, intermittently-updated display can prove invaluable.

Related to the Photo Wall is Twistr, a game which uses a large sheet of photos on a wall or a table. In this game, two players each carry two pens, one per hand. The display shows four photos per turn, one per player per hand. The object is to tap your own photos before the other player taps his. This game demonstrates three qualities of GIGApriints.



First, the system supports more than one user (and multiple pens). Second, this print allows us to test the ability for large prints to provide visual context in a search task. Finally, it is an interesting exercise in the study of proxemics [13]. We find that players *must* engage in what Hall deems the *intimate distance*, reaching around each other to acquire their own photos.

Based on technology developed for Photo Wall and Twistr, we designed a GIGApriint of bobcat photos for our camera trapping biologists. This print supports collaborative analysis of photographs. Tapping one photo retrieves it to a dis-

play, allowing for closer inspection; tapping two photos compares them. We demonstrated the sketches and the two related prints (Photo Wall and Twistr) to three biologists. These biologists commented that they could find the photo print useful for sorting their research content.

Remote Collaboration—BuddySketch

Inspired by our need-finding, we developed a print for remote users to discuss printed content (see Figure 4). In this application, participants at each remote location print out the same sheet containing content they want to analyze. For example, Karen, a biologist, prints out bobcat photographs from her camera project. Roger and John print out the same content at the other site, and then call Karen. When Karen wants to refer to a photo, she selects it with a tap of her pen. The photo is displayed on screens at both locations. If Roger now selects a photo, it will be displayed next to Karen’s. They can also trade ink annotations written on their sheets. The ink and photographs are *transported* back and forth by the software.

In addition to supporting remote collaboration, BuddySketch provides mobility, as the participants do not need to remain seated with a camera pointed at their print. The wireless pen frees the biologist from the “conferencing tether” so that one person can be seated at a table, while the other can be standing next to a wall print, or be working out in the field. This application, in augmenting *paper* prints with *electronic* information, offers a converse approach to the prior work of augmenting *electronic* displays with *paper* artifacts such as sticky notes [27].

Collocated Collaboration—Wall Posters

Robert Horn is a political scientist who produces what he calls *issue* and *mess maps*—wall-scale graphics that provide a visual summary of debates or current events. Upon hearing about GIGApriints, Mr. Horn provided us with two of his information graphics on the topic of Avian Flu. His timeline is a 13’x3’ print that discusses a hypothetical out-



Figure 4. BuddySketch supports remote collaboration. Users draw on a shared virtual space by writing on actual paper. They pass photos by tapping on printed photos with the digital pen. This screenshot was captured from one of our user study sessions.

break. The mess map is a 4'x2' print enabling readers to discuss the potential effects of an outbreak on the economy, transportation, and food. We augmented these prints with several types of active regions. One type enables a user to retrieve a source article by tapping on a box underneath an excerpt; the webpage loads on a nearby display. A second enables a user to add ink comments; tapping the upload box sends this ink to a wiki. With the same goal of supporting collocated communication, we also recently augmented four of our conference posters with these techniques.

Visualizing Large Datasets—Network Monitoring

Large public visualizations can support the input of multiple people. To examine the collaborative possibilities of GIGAprints, we created a network monitoring print (see Figure 5). This 3'x6' GIGAprint hangs in the hallway of our lab. The print comprises three types of content. First, it displays a grid of 225 graphs arranged in a small multiples pattern; each shows the network activity of a single machine. Second, the print displays 135 charts representing ISPs that our machines communicate with. In total, this visually represents 561,000 data points. Third, the print displays regions (buttons and map) that enable users to invoke queries. Feedback is presented through a handheld and overlaid projection. Users can alert the system administrator about suspicious activity by tapping *Email* and then specifying a PC. Users may also annotate the print with ink. These two interactions facilitate asynchronous collaboration.

The next set of interactions includes the mobile device. While computer names are anonymized on the printed public visualization, a network administrator can approach the print and select any computer; real-time private data is then *transported* to his handheld. Finally, richer interactions are available through georeferenced integration of the printed display with a projector. The projector acts as a targeted spotlight to provide real-time information. In ambient mode, the projector highlights machines as they engage in activity. When the user interacts, the highlights present the results of queries (e.g., the user can select a region on the map to see which locations the machines are talking to).

Mobile Queries and Data Capture—Augmented Maps

Maps have long been one of the prime benefactors of the resolution, size, and flexibility of print media. To explore



Figure 5. A user examines network data. The projector highlights graphs of machine activity that match the user's queries, while the handheld reveals private details. The user can also query by selecting a region on the map.



Figure 6. Two researchers use a map-based query to retrieve a geo-tagged photo on a handheld display.

the benefits of augmenting maps, we built several map-centric interactions. One enables users to select geo-tagged photographs from a database by tapping on a location on the printed map; matching photographs are presented on the mobile display (see Figure 6). We also built an augmented tourist map called AudioGuide that is shared between two users; this application includes a pen, Bluetooth audio earpiece, and handheld display. When a user selects a region on the map, an audio description of that region plays into the earpiece. This application studies the possibility for using GIGAprints out in the field, with *no digital visual display*. For example, a field biologist might use this to retrieve information about her experimental site, without having to read it from a digital screen. Finally, we have designed an augmented map for ant surveys, and demonstrated its interactions to five biologists. This map includes the Bluetooth earpiece and enables a user to capture content written directly on the map. When the user finishes writing, she hears a voice confirmation generated through text-to-speech. This print supports the mobility of field biologists, requiring only two hands (map in one; pen in the other). In future iterations, we will add a handheld display and integrate it with GPS, which may be useful if the biologist has a collaborator (to hold the extra devices). The domain experts felt that the audio interface would work, but suggested that the map might be cut into pieces for easier manipulation. We see that the flexibility of GIGAprints (supporting cut-ting) can satisfy a biologist's mobile field requirements.

Intermittent Updates—Blog Reader

This application displays ten weblog entries from a Really Simple Syndication (RSS) feed in a public place, and provides space for written comments. Individuals who routinely pass the display are able to casually scan it for new comments on stories that interest them. Additionally, the comments are captured digitally, and published as images to a second RSS feed. Thus, this application leverages both asynchronous and remote collaboration.

Navigating Virtual Views—Interactive Floor Plans

Finally, we demonstrated a Wizard-of-Oz version of an interactive floor plans application to one professional architect/developer. This architect used the application to navigate a Google Sketchup rendering by moving the pen

around the whiteprint. She jumped to different areas by tapping on the print, rotated the virtual view by pivoting the pen, and annotated the print with ink. The user remarked that the *walk-through* and *look-around* interactions were intuitive, and played like “an arcade game.”

A DESIGN SPACE FOR GIGAPRINTS

From our experience in building Gigapixel Prints, we suggest four axes that GIGApriint designs span. Here, we describe our design space, which frames our applications, and directs us toward future research.

Scale: While the name “GIGApriints” stresses the large-scale, the actual prints can span a range of sizes. The largest print we have constructed is 13' wide; the smallest is 3.5' wide. An even smaller print can be folded into a breast pocket (providing mobility). A large print, however, is most usable when displayed on a wall or table (supporting collaboration). Relating this to prior work, one can see that GIGApriints fall into the *yard* category of Weiser’s “computing by the inch, foot, and yard” [42]. Examples of related systems that fall into these categories include WebStickers [17], Books with Voices [26], and DigitalDesk [43] for the inch, foot, and yard, respectively.

GIGApriints can also vary in the *spatial coordination of paper and digital content*. We use Klemmer’s taxonomy [25]: georeferenced, collocated, non-collocated, and non-visual. When digital and paper content are georeferenced, the graphical feedback is overlaid on the print (*e.g.*, the network monitoring app). With collocated systems, the paper and digital content occupy adjacent physical space (*e.g.*, the Photo Wall). In non-collocated systems, the physical and digital do not occupy a shared space. The Blog Reader exemplifies this, as the published RSS feed lives apart from the print. Non-visual output describes systems with other digital modalities of feedback (*e.g.*, AudioGuide).

Temporal coordination represents the measure of time it takes for a paper-based action to invoke a change in the digital content. This coordination can be synchronous (*e.g.*, real-time audio feedback in AudioGuide) or asynchronous (*e.g.*, batched handwriting captured from the augmented map).

The final design axis for GIGApriints is their *update frequency*, as GIGApriints leverages printing as a computationally controlled display. This varies from *static*, to *intermittent*, to *real-time*. A wall-sized poster may never be updated; the blog reader can be updated intermittently (daily); a streaming display of incoming photos can be printed continuously, in real-time.

This design space draws upon synthetic work on taxonomies for tangible user interfaces (TUIs). In particular, our design space is similar to Klemmer’s work [25], with the distinction that our formulation is concerned with the *interaction techniques* rather than the *software architecture* for implementing these systems. The MCRit model’s division of a TUI into intangible and tangible components is subsumed

by our *spatial coordination* axis [41]. Fishkin’s framework [8] suggests two primary concerns: *metaphor* and *embodiment*. Embodiment is also subsumed by the *spatial coordination* axis; and metaphor, while applicable to *user experience*, is less relevant for describing the *interface* embodying that experience.

THE GIGAPRINTS SYSTEM

We now describe the building blocks that compose a GIGA print application, the core components of both the user’s conceptual model of the application, and the designer’s system model. The central visual component of a GIGA print is the large print (there may be in fact more than one, but we assume one in this example). A user interacts with this print with one or more pens. She receives output—real-time or batched—through auxiliary devices: LCDs, projectors, handhelds, audio headsets, and more. The user may ask the application to generate more prints; this is useful if new data arrives, and the print needs to be updated.

These components—the prints, pens, and auxiliary devices—form the core of every GIGA print application. The designer and developer must 1) generate appropriate paper-based visualizations, 2) decide upon pen interactions that suit the task, and 3) connect the devices in a way that appropriate output is triggered for every input.

Building GIGApriints

The first step in creating a GIGA print is to decide on the content. The designer first creates a template for the print based on this content. The template includes background graphics, and content that does *not* need to be updated periodically (and can be arranged by hand). At this stage, the designer uses any tool he or she is familiar with, as long as it exports to a PDF document. In our work, we use Adobe Illustrator to design many of our GIGA prints.

The next step is to augment the PDF with active regions. If the print will contain few active regions, the designer may use the GIGA print design tool, a direct manipulation interface for specifying regions. Alternatively, a developer may add active regions through a Java API (the preferred method for complicated prints). These two approaches provide the potential for designers or domain experts to create an initial, functional prototype; this early prototype may serve a pragmatic role for the users and an epistemic role in creating a useful communication tool for working with developers, who could then create a richer version.

Finally, the developer must architect the software portions of the GIGA print. In Java, the developer can add event handlers to active regions, register input devices (*e.g.*, pens), and attach output devices (*e.g.*, projector). At runtime, the GIGA print program listens for real-time or batched input from pens. When input arrives, the GIGA prints system directs the pen events to the active regions. The event handlers (written by the developer) are invoked, and the result can be as simple as playing a sound, or as intricate as updating a wiki to include the most recent pen annotations.

Implementation

The rendering of the Anoto pattern is handled automatically by the GIGAprints system; pattern is layered onto any region that contains an event handler. Thus, the developer never needs to manage low-level transformations from pattern coordinates to locations within an active region. This information is encapsulated by the PenEvent object that is passed to the event handler. The GIGAprints system is built primarily on Java SE 6.0, but also contains a Microsoft .NET 2.0 component that interacts with the Anoto SDK. The system integrates with the Adobe tools through HTTP and XML. The GIGAprints Flash 9 plug-in is written in ActionScript 3.0, and the GIGAprints Acrobat plug-in is written in JavaScript. GIGAprints are rendered using the iText library.

DISCUSSION

To articulate the merits of the GIGAprints ensemble approach, we analyzed GIGAprints through three lenses. First, we examine how GIGAprints supports ubiquitous computing. Second, we analyze the tools by applying the cognitive dimensions framework [4, 11]. Finally, we examine how GIGAprints support collaboration by acting as artifacts for distributed cognition [20, 34].

Ubiquitous Computing Themes

Klemmer *et al.* describe five themes for describing the importance of incorporating the richness of human physical activity in interactive systems [14]. We analyze how GIGA prints fits into these themes. Since a primary part of the interface is paper and pen, the interface supports tasks that engage the user in *thinking through doing*. The pen encourages sketching and the mobility of the ensemble enables a user to work out in the world. As for *performance*, we note that pen input supports high performance sketching. GIGA prints can also support bimanual (dual pen) input, for increased performance. One limitation is that handwriting on a print is slower than entering text on a keyboard.

Paper prints excel in tasks that require *visibility* of content. GIGAprints created for design studios can support situated learning, so that an apprentice may learn from a master of the trade. The display's permanence (it does not need batteries) promotes visibility, and in turn supports collocated collaboration. Paper and digital representations each have different *risk* characteristics: most notably, changes made with the inking pen cannot be undone as within a GUI. This risk can prove to be both a benefit (providing confidence as to the permanence of changes) and a hindrance (reluctance to mark up a print). Finally, GIGAprints support *thick practice*, as they fit well into the many disciplines that leverage large paper maps, posters, and diagrams. An added benefit is that in many cases, if the technology fails, the interface degrades gracefully and becomes an inking pen on a piece of paper. For example, if a biologist's GIGAprint map fails, it is no worse than the paper map she uses today.

Cognitive Dimensions

Cognitive dimensions are design principles developed for programming language design [11]. As it lowers the threshold for discussing our system's usability, we adapt them to evaluate our system. We cluster the 14 dimensions into five groups, and apply each to the prints and/or the tools.

Visibility and Juxtaposability: For activities such as exploration and design, GIGAprints can support our need for visibility and juxtaposability. The paper can permanently reside on a wall. When a user needs to look at information, it will be there, without navigating any folders.

Secondary Notation, Mapping, and Consistency: One benefit of the pen and paper interface is that adding *secondary notation* (handwritten annotations) is easy. When a user adds sketches or writing to a print, he is adding supplementary information, which may mean nothing to the computer, but everything to a collaborator. Additionally, these interfaces *map* closely to domains that value paper tools (*e.g.*, architecture). Because of this close mapping, and because of the parallels between our pen-based interactions and mouse or stylus-based GUI interactions, the interfaces can be consistent with a user's expectations. For example, pen taps (like mouse clicks) signify selections.

Viscosity, Progressive Evaluation, and Commitment: These guidelines apply to the design tools. While the Java-based tools provide a flexible (low *viscosity*) environment for creating augmented prints, the process of creating a GIGAprint currently requires *some premature commitment*. The designer of the print must plan the look of the paper interface ahead of time, so that the event handlers and pattern can be added to visible regions on the print. Additionally, iterative debugging is currently a viscous process. Modifications to the interface must be sent to a printer; with current technology, rendering and printing the patterned interfaces takes from 10 minutes to half-an-hour, depending on the print's complexity and the printer's speed.

Terseness, Expressiveness, and Abstraction Gradient: Terseness applies to both the interactions and the tools. Since pen and paper interactions map to real-world tasks, input can be expressed as tersely as a check mark in a box. As for the tools, the architecture allows for compact expression of rich interactions (terse yet expressive), with complexity mitigated by Java's abstraction-friendliness. For example, one of our conference posters with five active regions contains 56 lines of source code.

Errors, Mental Operations, and Dependencies: The API for event handling is similar to Java Swing; we hope that the familiar constructs will help developers to avoid errors, build richer interactions (hard mental operations) and minimize hidden dependencies between pieces of code. However, we cannot evaluate this until more developers use the tools; we leave a tools study for future work.

Distributed and External Cognition

The interfaces that designers can create using the GIGApriint tools can aid collaborative reasoning by supporting distributed and external cognition [20, 35]. GIGApriints succeed with respect to these views. For example, the static print excels at *explicitness*, *visibility*, and *cognitive tracing* [35]. Cognitive traces such as annotations can aid a user in reasoning about visualizations. These traces (like *secondary notations*) may have no semantic meaning for computer systems, but have tremendous value for the end users. If we consider how the size of the prints may encourage collaboration, we can quickly see that GIGApriints allow users to offload computation not only to external paper representations, but to the mental space of their collaborators.

The frameworks we have applied suggest that Interactive Gigapixel Prints can provide flexible support for in-the-world tasks, and can aid collaborative processes by serving as external artifacts for cognition. However, this analysis has its limits; it cannot fully predict how real users will respond to the prints and accompanying interactions.

EVALUATION

To evaluate the user experience, we conducted a first-use study of three GIGApriints — Twistr, BuddySketch, and AudioGuide. The study focused on four hypotheses:

- H1 Large prints provide rich *visual context* for graphics-heavy applications, aiding in search and comparison.
- H2 The large scale, multiple pens, and devices support *collaboration*.
- H3 The robustness of pen and paper, the modes of input handling (batched and real-time), and the ability to use real-time audio displays aids in *mobility*.
- H4 The pen and paper aligns with current practices, and even if the technology fails, GIGApriints will *degrade gracefully*; they are never worse than paper-only tools.

Each session included two simultaneous participants, and lasted 1.5 hours. The 12 participants (six male; six female) were introduced to the technology, and then completed three tasks—one per print. They then filled out a questionnaire evaluating the applications and the GIGApriints concept in general. Each participant received a \$25 gift card.

For the first task, the two users engaged in three rounds of

Twistr, each lasting 15 turns (Figure 7). The users listened to a brief tutorial of the game, and then began. Whoever scored more points was declared the winner. For the second task, the users engaged in a simulation of remote video conferencing. We seated the users across the room from each other, each with a laptop and paper sketching interface (BuddySketch); they were far enough apart so that each could not see the content of the other’s laptop or print. Thus, the task simulated a conferencing setup with perfect fidelity sound and visuals; we sought to test how BuddySketch’s ability to transport photos and handwriting could improve this otherwise perfect configuration, supporting synchronous remote collaboration. For the third task, the goal was take a walking tour with AudioGuide. When the “tourists” tapped on a building on the map, the user with the earpiece could hear a description of that building.

RESULTS

Our evaluation revealed that GIGApriints has the potential to support our goals of visual context, mobility, and collaboration. This section highlights quantitative results (Figure 8), includes observations from the sessions, and discusses how these observations will impact our future directions.

From the data, we find that the most obvious benefit of GIGApriints to users was the fact that they could visualize lots of content at once. This factored into Twistr, when users remarked that they could start remembering, after a few rounds, where certain photos were. Some commented that nature photos were difficult to distinguish, whereas photos with faces and sharp colors jumped out. During one round, one user remarked “I just saw this one!” and pruned her visual search space by looking in only one area.

During BuddySketch, we observed users moving laptops and prints around to work on different areas of their drawings. With the AudioGuide, users walked around outside while interacting with the print and audio feedback. The print worked well in the sun, where the handheld’s display (which we did not leverage) was rendered unreadable. Users reported that AudioGuide enhanced their mobility (6.4 out of a 7-point scale). With AudioGuide, one biologist exclaimed to her collaborator, “This feels very much like we are doing our ant survey.” One other user received a phone call during his session; he was able to quickly take off his



Figure 7. Left) A participant reaches across his opponent to select a photo in *Twistr*. Middle) A participant gestures to describe his drawing of a camera mount in *BuddySketch*. Right) One collaborator selects a region on the *AudioGuide* while the other listens to the description.

earpiece, let the caller know he would return the call, and replaced his audio interface. These observations support the mobility aspects of GIGAprints.

As the Twistr game fostered *competition*, we observed that users would stand close to each other and reach in front of the other to acquire their photos. This illustrates how large prints can support multiple simultaneous users, and the same qualities can instead support collocated collaboration. For AudioGuide, we observed users working with each other to determine which areas of the map to listen to. One common situation was that User A would carry the pen and map, while User B carried the handheld and earpiece. In this fashion, the “driver” would click on parts of the map, and the “listener” would explain what the AudioGuide was saying. BuddySketch supported collaboration by enabling “remote” users to see each other’s drawings. One group ended up talking about and analyzing football plays.

One pair of users tested the strength of the connection between the handheld and digital pen. The user with the handheld and earpiece walked away while the owner of the map repeatedly tapped on audio regions. Eventually, the connection broke, and the tool gracefully degraded to *no worse* than a regular map. Ink could still be written on it, and the tour could still be taken (albeit without audio descriptions).

There were several limitations with the applications we studied. As we strove for an interface that could be reconfigured easily, BuddySketch did not provide a projection overlay (as with network monitoring) so that a user could see digital versions of the other’s handwriting georeferenced with his own. Users reported this as the main problem with BuddySketch, and suggested a number of improvements, including a simple grid with letters and numbers on either axis (*e.g.*, square A-5, as in some maps).

One person mentioned that if he were actually a tourist, he could benefit from GPS capability in the AudioGuide, to help him find his bearings. We did anticipate this desire; however, at the time of the study, AudioGuide’s “You Are Here” region only returned pre-rendered speech, regardless of where one was standing. One user suggested a manual approach that did not require GPS. A pen gesture could tell the system which way one was facing. The audio could then be tailored to a person’s orientation.

RELATED WORK

The GIGAprints research draws on prior work in three areas: augmented paper interactions, mobile computing, and software architectures. In this section, we explain how this work contributed to our system’s design.

Augmented Paper Interactions

Since the introduction of DigitalDesk [43], there has been significant interest in integrating paper and digital interactions, *e.g.*, [6, 16, 22, 29]. Video Mosaic demonstrated how paper storyboards can be used to *manipulate* temporal information such as video [30]. Subsequent work, Ariel, used overlaid projection and vision-based tracking of light pens

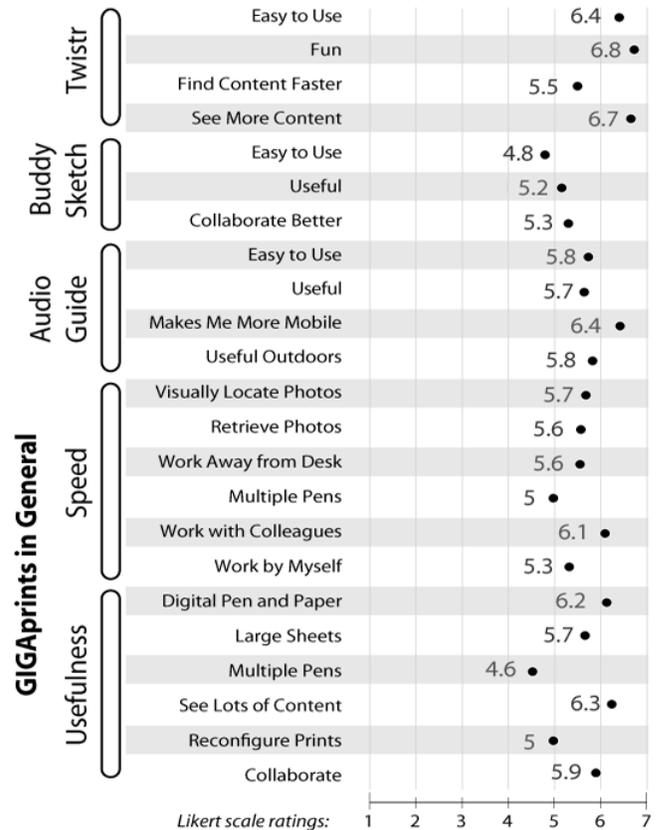


Figure 8. The study revealed several highlights: Twistr was fun and provided good visual context. AudioGuide’s mobility was appreciated by participants. The most useful feature of GIGAprints was its ability to show lots of content at the same time.

and LEDs to augment engineering drawings, achieving fluid review and annotation [31]. Gigapixel Prints extend the ideas presented in Video Mosaic and Ariel in two ways. First, the print can be updated computationally and at regular intervals, with the paper conceptualized as a display for the computer. Second, GIGAprints are mobile; a field worker can fold up a print and use it outside, in multiple locations, or various configurations. An auxiliary device provides the real-time feedback in the mobile scenario. The prior work, including DigitalDesk, was designed for a fixed physical location.

Two mobile examples include A-book, which integrates a laboratory notebook with a PDA [32], and ButterflyNet, which provides data capture for field biologists [44]. A-book’s “interaction lens” enables biologists to author links between pages and links to external sources. ButterflyNet’s “hotspot” enables authoring of links between field photos and note pages. GIGAprints adopts these *associate* interactions for working with mobile devices, and was founded on a similar view of respecting thick practice. GIGAprints also leverages ButterflyNet’s mixed-mode input processing (some in-the-field interactions provided real-time output, while others, like the notes capture, only provided feedback back at the field station). GIGAprints generalizes support for

mixed-mode processing by including (in the core architecture) event handlers that process both batched input and real-time input. Another augmented paper system, Audio Notebook, introduced a paper notebook where tapping parts of a written page retrieved the audio recorded when those notes were written [38]. Similarly, Books with Voices, a system for oral historians, introduced paper transcripts as an input medium for browsing video [26]. GIGAprints was inspired by the qualities of these systems, and contributes larger paper displays, and explicit support for multiple simultaneous input/output devices through the design tools.

PADD demonstrated techniques for manipulating documents in either digital or physical form [12]. PapierCraft included a set of gestures for interacting with these documents, allowing users to copy edit digital documents by working with the physical copy [28]. GIGAprints expands the PADD concept of digital/physical duals by including the view that the physical and digital counterparts may be disparate in content and span multiple artifacts, yet complement each other to support a unified end goal.

Rasa and NISMap demonstrated a class of multimodal interactions with paper [5]. The systems contributed fusion techniques to integrate speech with pen strokes; additionally, they demonstrate paper tools' robustness and support for graceful degradation. GIGAprints currently does not support speech input, but contributes multiple pen support and bimanual interactions (pen and device or dual pens).

Mobile and Ubiquitous Computing

Our motivation for developing Gigapixel Prints was driven by the goal of ubiquitous computing [42]. To approach Weiser's vision, many projects have tried to integrate paper and digital components. One example is WebStickers, which leveraged sticky notes to allow users to share Web bookmarks [17]. Listen Reader is an augmented paper book that allows a reader to manipulate audio streams by moving her hands onto different parts of a page [2]. Designers' Outpost enabled synchronous, collocated design work where sticky notes were integrated with an electronic wall-sized display [27]. GIGAprints takes another step toward this end goal, by enabling us to improve *existing tools and practices* by augmenting artifacts such as posters, architectural plans, maps, and visualizations with interactivity.

Software Architectures

The design of the programming model and architecture was informed by two existing toolkits for *graphical* user interfaces, subArctic [19] and Java Swing [39]. In particular, we adapted subArctic's idea of maintaining state in dispatch agents, and composing agents to achieve extensibility. For example, our marking gestures handler receives its input from the default pen handler, which in turn receives its input from the low level pen servers. However, to maintain a low threshold for Swing programmers, our handlers are employed in a fashion similar to Swing *listeners*. Beyond these models of GUI input handling, GIGAprints contributes

architecture support for multiple simultaneous input devices (the handler can query the pen's ID) and multiple output devices (the programmer adds *Devices* to an *Application*).

CONCLUSIONS AND FUTURE WORK

In summary, we have contributed to the ubiquitous and mobile design space in four ways. First, we introduced Interactive Gigapixel Prints, a class of interfaces where large prints are augmented with digital devices. Second, we introduced ensemble interaction techniques that enable us to leverage the print and digital media, and illustrated these techniques through example applications. Third, we presented a design space and programming model for building GIGAprints. Finally, we reported feedback from our first-use study and domain expert reviews of GIGAprints, which suggest that these techniques can benefit tasks that require visual context, mobility, and collaboration.

We have released the GIGAprint applications and tools for designing and building GIGAprints as open source software, and are currently working toward longitudinal evaluations with the domains we have studied. Looking forward, we are continuing research into flexible software architectures for paper-plus-digital applications and interactions.

ACKNOWLEDGEMENTS

We thank Terry Winograd for his insightful comments; Rodolfo Dirzo, Trevor Hebert, and Mor Naaman for access to their photo databases; Jory Bell for the OQO; and members of the Stanford Graphics Lab for feedback. NSF Grant IIS-0534662 supported this work.

REFERENCES

- 1 Anoto AB, Anoto Technology. <http://www.anoto.com>
- 2 Back, M., J. Cohen, R. Gold, S. Harrison, and S. Minneman. Listen Reader: an Electronically Augmented Paper-based Book. CHI: ACM Conference on Human Factors in Computing Systems. pp. 23–29, 2001.
- 3 Bellotti, V. and Y. Rogers. From Web Press to Web Pressure: Multimedia Representations and Multimedia Publishing. CHI: ACM Conference on Human Factors in Computing Systems. pp. 279–86, 1997.
- 4 Blackwell, A. and T. Green, Notational Systems—The Cognitive Dimensions of Notations Framework, in HCI Models, Theories, and Frameworks: Toward a Multidisciplinary Science, J.M. Carroll, Editor. Morgan Kaufmann. p. 103–33, 2003.
- 5 Cohen, P. R. and D. R. McGee. Tangible Multimodal Interfaces for Safety Critical Applications, Communications of the ACM, vol. 47(1): pp. 41–46, 2004.
- 6 Dymetman, M. and M. Copperman. Intelligent Paper. EP'98: International Conference on Electronic Publishing: Springer-Verlag GmbH. pp. 392–406, 1998.
- 7 Elliott, A. and M. A. Hearst. A Comparison of the Affordances of a Digital Desk and Tablet for Architectural Image Tasks. International Journal of Human Computer Studies 56(2). pp. 173–97, 2002.

- 8 Fishkin, K. P. A Taxonomy for and Analysis of Tangible Interfaces. *Personal and Ubiquitous Computing* 8(5). pp. 347–58, 2004.
- 9 Foley, J. D. and V. L. Wallace. The art of natural graphic man-machine conversation. *IEEE* 62(4). pp. 462–71, 1974.
- 10 Gigapxl Project, Gigapxl Project. <http://www.gigapxl.org>
- 11 Green, T. R. and M. Petre. Usability analysis of visual programming environments: A 'cognitive dimensions' framework., *Journal of Visual Languages and Computing* 7. pp. 131–74, 1996.
- 12 Guimbretière, F. Paper Augmented Digital Documents. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 51–60, 2003.
- 13 Hall, E. T., *The Hidden Dimension*. New York: Anchor. 201 pp. 1966.
- 14 Hartmann, B., S. R. Klemmer, and L. Takayama. How bodies matter: five themes for interaction design. *DIS 2006: ACM Symposium on Designing Interactive Systems*. pp. 140–49, 2006.
- 15 Heath, C. and P. Luff, *Technology in Action (Learning in Doing: Social, Cognitive & Computational Perspectives)*: Cambridge University Press. 286 pp. 2000.
- 16 Heiner, J. M., S. E. Hudson, and K. Tanaka. Linking and Messaging from Real Paper in the Paper PDA. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 179–86, 1999.
- 17 Holmquist, L. E., J. Redström, and P. Ljungstrand. *Token-Based Access to Digital Information. Handheld and Ubiquitous Computing*: Springer-Verlag, 1999.
- 18 Hudson, S., J. Mankoff, and I. Smith. Extensible input handling in the subArctic toolkit. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 381–90, 2005.
- 19 Hudson, S. E., J. Mankoff, and I. Smith. Extensible Input Handling in the subArctic Toolkit. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 381–90, 2005.
- 20 Hutchins, E., *Cognition in the Wild*: MIT Press. 408 pp. 1995.
- 21 Imhof, E., *Cartographic Relief Presentation*: Walter De Gruyter, Inc. 388 pp. 1982.
- 22 Johnson, W., H. Jelinek, L. K. Jr., R. Rao, and S. Card. Bridging the Paper and Electronic Worlds: The Paper User Interface. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 507–12, 1993.
- 23 Kerouac, J., *On the Road*. Reprint ed: Penguin Books. 307 pp. 1991.
- 24 Kirsh, D. The Intelligent use of Space. *Artificial Intelligence* 73(1-2): Elsevier. pp. 31--68, 1995.
- 25 Klemmer, S. R., *Tangible User Interface Input: Tools and Techniques*, Unpublished PhD, University of California, Computer Science, Berkeley, CA, 2004. <http://hci.stanford.edu/srk/KlemmerDissertation.pdf>
- 26 Klemmer, S. R., J. Graham, G. J. Wolff, and J. A. Landay. Books with Voices: Paper Transcripts as a Tangible Interface to Oral Histories. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 89–96, 2003.
- 27 Klemmer, S. R., M. W. Newman, R. Farrell, M. Bilezikjian, and J. A. Landay. The Designers' Outpost: A Tangible Interface for Collaborative Web Site Design. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 1–10, 2001.
- 28 Liao, C., F. Guimbretière, and K. Hinckley. *PapierCraft: A Command System for Interactive Paper*. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 241–44, 2005.
- 29 Luff, P., C. Heath, M. Norrie, B. Signer, and P. Herdman. Only Touching the Surface: Creating Affinities Between Digital Content and Paper. *CSCW: ACM Conference on Computer-Supported Cooperative Work*, 2004.
- 30 Mackay, W. E. and D. S. Pagani. Video Mosaic: Laying Out Time in a Physical Space. *MULTIMEDIA: ACM International Conference on Multimedia*. pp. 165–72, 1994.
- 31 Mackay, W. E., et al. Ariel: Augmenting Paper Engineering Drawings. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 421–22, 1995.
- 32 Mackay, W. E., G. Pothier, C. Letondal, K. Bøegh, and H. E. Sørensen. The Missing Link: Augmenting Biology Laboratory Notebooks. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 41-50, 2002.
- 33 Myers, B. A New Model for Handling Input. *ACM Transactions on Information Systems* 8(3): ACM Press. pp. 289–320, 1990.
- 34 Perry, M., *Distributed Cognition, in HCI Models, Theories, and Frameworks: Toward a Multidisciplinary Science*, J.M. Carroll, Editor. Morgan Kaufmann. p. 193–223, 2003.
- 35 Scaife, M. and Y. Rogers. External cognition: how do graphical representations work? *International Journal of Human-Computer Studies* 45. pp. 185–213, 1996.
- 36 Schilit, B. N. and U. Sengupta. Device Ensembles. *Computer* 37(12). pp. 56–64, 2004.
- 37 Sellen, A. J. and R. H. R. Harper, *The Myth of the Paperless Office*. 1st ed: MIT Press. 242 pp. 2001.
- 38 Stifelman, L., B. Arons, and C. Schmandt. The Audio Notebook: Paper and Pen Interaction with Structured Speech. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 182–89, 2001.
- 39 Sun Microsystems, *Swing: Java Foundation Classes*. <http://java.sun.com/>
- 40 Tufte, E. R., *The Visual Display of Quantitative Information*. 2 ed: Graphics Press. 197 pp. 2001.
- 41 Ullmer, B., H. Ishii, and R. J. K. Jacob. Token+Constraint Systems for Tangible Interaction with Digital Information. *ACM Transactions on Computer-Human Interaction* 12(1). pp. 81–118, 2005.
- 42 Weiser, M. The Computer for the 21st Century. *Scientific American*. pp. 94–104, 1991.
- 43 Wellner, P. Interacting With Paper on the DigitalDesk, *Communications of the ACM*, vol. 36(7): pp. 87–96, 1993.
- 44 Yeh, R. B., et al. ButterflyNet: A Mobile Capture and Access System for Field Biology Research. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 571–80, 2006.