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

Ron B. Yeh, Joel Brandt, Jonas Boli, Scott R. Klemmer
Stanford University HCI Group
Computer Science Department
Stanford, CA 94305-9035, USA
{ronyeh, jbrandt, jboli, srk}@cs.stanford.edu

ABSTRACT
For centuries, large paper information graphics such as maps have been important cognitive artifacts in navigation, architecture, design, engineering, and scientific work. Paper-based work practices leverage the high resolution, low cost, reliability, mobility, and flexibility of paper—yet lack the interactivity afforded by digital technologies. This paper introduces Interactive Gigapixel Prints (GIGAprints), computer controlled large-scale paper displays that afford direct pen-based input. These paper prints are augmented with digital displays, integrating the high *spatial* resolution but low *temporal* resolution of wide-format printing with the lower spatial resolution but higher temporal resolution of digital displays. Using large paper displays and digital devices together as an ensemble leverages the relative benefits of each medium; GIGAprints afford both ambient awareness and simultaneous viewing and input from multiple users. The pen-based interaction includes selection, progressive information disclosure, filtering, and annotation. This paper contributes a design space for integrated paper and digital interactions, an infrastructure for creating interactive ensembles of large paper displays and digital devices, and four applications built using our infrastructure, each illustrating points in the design space.

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
This paper was inspired by the rich information environments present in large paper prints, maps, and scrolls—from the incredibly high resolution photographs from the Gigapixl Project [13], through Kerouac's *On the Road* [26] manuscript (typed on a single roll of teletype paper 120' long), to Imhof's cartography [21], and Minard's flow maps [50]. By integrating printed displays with overlaid digital projection or adjacent mobile devices, our research leverages centuries of study in printed visualization [6, 50], and the recent decades of research into interactive visualization [7, 54].

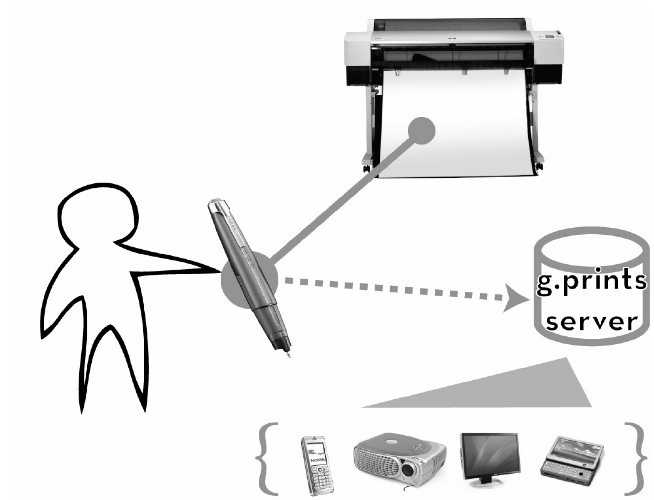

Figure 1. The Interactive Gigapixel Prints ensemble: wide-format inkjet printers, digital pens, and digital displays. GIGAprints interactions enable users to partition activities across large-scale paper-digital displays, digital pens, and mobile devices. These ensembles support ambient awareness, visual context, and collaboration.

Prior work (*e.g.*, [15, 47]) has shown that the physical affordances of paper — flexible input, grasping, carrying, and folding—have encouraged its significant use, even in today's technology-enabled world. Paper-based visualizations offer users an immediate sense of *familiarity,* stemming from our everyday use of printed paper displays. They are *ubiquitous*, yet their calm presence affords their ambient qualities in a room when not being used. Paper also affords *mobility* — for example, an archaeologist can fold a high resolution map into a small form factor (with infinite "battery life"!), and sketch directly on the map, as it offers *flexibility* of input. Print also provides *robustness*, as paper can survive heavy abuse. The *high resolution* of printing reveals nuanced imagery, and enables us to achieve detail on demand, not by manipulating a graphical slider, but by drawing nearer to the display. Likewise, the large *size* of some printed visualizations renders zooming unnecessary. Large physical sizes also enable synchronous and collocated *collaboration* around a tabletop or a wall. And finally, the *physicality* and *tangibility* of paper enables users to leverage their spatial reasoning and navigation abilities.

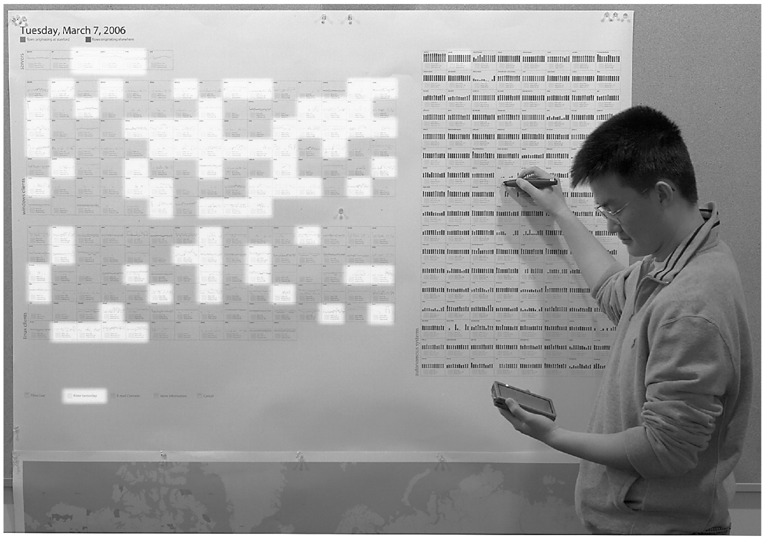

Figure 2. A user examines live network data on an Interactive Gigapixel Print. The projector highlights graphs of machine activity that match the user's query; the mobile device reveals detailed information about the selected machine.

Bellotti and Rogers' [5] study of (print) publishers publishing on the Web found that paper artifacts were critical to ensure quality in the process, because they supported multiple representations of a newspaper page. Interactive Gigapixel Prints combine the affordances of paper with the affordances of digital media to enable people to benefit from the advantages of both.

Contribution Overview and Road Map

Interactive Gigapixel Prints (GIGAprints) are computer controlled large-scale paper displays that afford direct pen-based input. These prints are augmented with digital displays, to integrate the high *spatial* resolution but low *temporal* resolution of wide-format printing with the lower spatial resolution but higher temporal resolution of digital displays. GIGAprints introduce techniques for augmenting large paper displays with pen input (using the Anoto technology [1]) and mobile devices. These pen-based interaction techniques include selection, filtering, progressive information disclosure, and annotation. GIGAprints also provides an infrastructure for creating and interacting with these printed displays. A video demonstrating GIGAprints is available online, linked from http://hci.stanford.edu/g-prints

In the following section, we first introduce a design space for augmented paper interactions and subsequently describe in more detail the subspace of these interactions spanned by GIGAprints. We then explain the design of ensemble interaction techniques [46] that can take advantage of hybrid paper-digital media. Finally, we discuss four applications we have built that demonstrate different points in the design space of GIGAprints.

INTERACTING WITH GIGAPIXEL PRINTS

The Interactive Gigapixel Prints ensemble leverages *wide-format* inkjet printers for high-resolution display, *digital pens* for user input directly on the printed surface, and *digital displays* for interactive content (see Figure 1). The GIGAprints prototype employs a wide-format inkjet printer (we have used both a 44" Epson Stylus Pro 9800 and a 36" HP DesignJet 1055cm plus), and Anoto digital pen technology [1]. With Anoto, pens determine their location by tracking a faint dot-pattern printed on a paper surface; the pens we use (the Nokia SU-1B [39]) stream data in real-time to the GIGAprints server. The GIGAprints technology can incorporate digital displays in one of several ways. For augmenting large prints with personal information and interaction, GIGAprints integrates with mobile devices (we use the OQO model 01+ [40]), LCDs, and projected displays for real-time output.

A Taxonomy of Ensemble Interactions

In 1974, Foley and Wallace suggested that all interactions with a graphical interface fall into one of four categories: *pick*, *button*, *locate*, and *valuate*.

Over the years, variations on this taxonomy have been proposed of between four and six categories, and have been used to guide input architectures such as Interactors [37]. Broadly speaking, these taxonomies specify a *select* operator and several *manipulate* operators. Interacting with device ensembles can be described similarly, with the addition of two operators that facilitate inter-device communication (see Figure 3). We refer to the device initiating the action as the *sender*, and the device executing the action as the *receiver*.

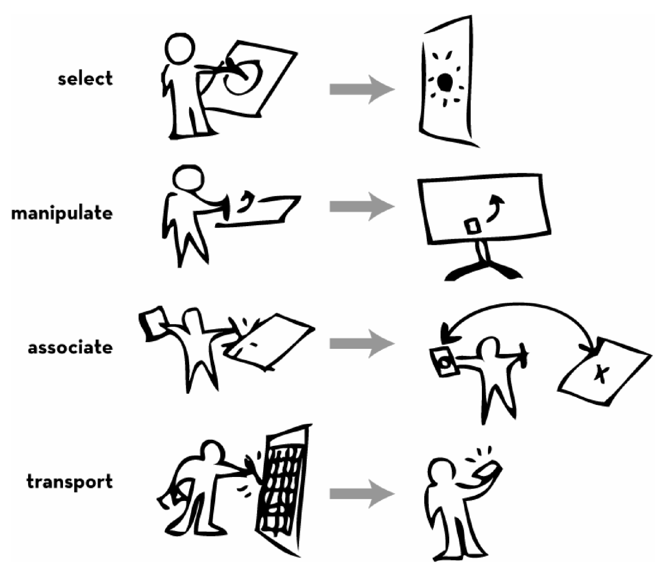

Figure 3. The user initiates a *select* operation on one device (the sender) to designate one or more items of interest on the other device (the receiver). A *manipulate* interaction enables the sender to modify content on the receiver. An *association* operation creates a symbolic link between two pieces of information (one on the sender and one on the receiver). The transport operation describes the movement of content from sender to receiver.

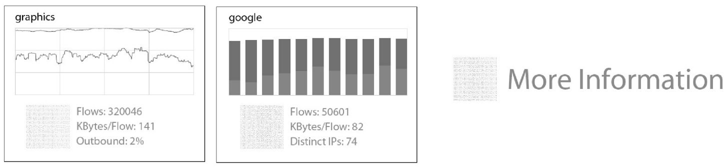

Figure 4. The network awareness tool displays three types of content: a composite metric of network activity on an individual machine (*left*), the in- and outbound flows of an ISP (*middle*), and paper buttons for users to invoke commands (*right*).

- With the *select* operation, a user employs the sender device to designate one or more objects of interest on the receiver.
- *Manipulate* is an umbrella term for the set of operations that enable the sender to modify content on the receiver.
- An *association* operation creates a symbolic link between a piece of information the sender device and a second piece on the receiver.
- The *transport* operation describes the actual movement of a piece of content from the sender to the receiver.

In the GIGAprints prototype, device ensembles comprise Anoto-enabled *paper print(s)* and nearby *computer display(s)* such as a mobile phone, LCD, or projector. As prior work has explored the techniques for forming connections between devices [3, 10, 17, 41], this paper concentrates on interactions once the ensemble is already formed.

Applications

In this section, we demonstrate the GIGAprints ensemble interaction techniques through four examples: a network awareness tool, a photo monitor, a collaborative sketching application, and a map-based query system.

In all examples, the *pen* component is the sender device, and is used to perform selection, progressive information disclosure through manipulation and transport, filtering, and annotation. The *printer* is used to print slow-moving content that benefits from high resolution, such as maps, photographs, and visualizations. The receiver is either the projector or the mobile device. For example, the sketching application explained in this section demonstrates the use of large paper as an input surface whose content is mirrored on a remote projector.

Together, this ensemble provides three benefits: 1) ambient awareness through its large physical size and persistent presence, 2) visual context, enabling users to rapidly see large quantities of data at high resolution, and 3) support for collaboration through its wall-scale form factor. Also, these applications demonstrate how a physical/digital ensemble can be used to partition information across public and private displays. Through these applications, we demonstrate that GIGAprints offers value in its basic case of pen input on large paper displays, and that this value can be enhanced through an array of digital display devices from personal mobile devices through geo-referenced wall-scale projection.

Network Awareness

Large, public, printed visualizations are well suited for tasks that may require the input of many people, or tasks that are not the responsibility of any single person, but rather are dependent on the community as a whole. Awareness of the computer network is such a task in our CS department. Each area within the department has a part-time network administrator that spends the majority of their time dealing with network intrusions. Our colleagues in information visualization are currently researching network visualization techniques (*e.g.*, [42]), and the design of this application is the result of extensive discussion with them. Most notably, the second author is a researcher on both this work and the network security work, and the design of the network activity graph is an application of the visualization group's research.

To increase the visibility of network activity, potentially reducing the workload of this task, we have designed an Interactive Gigapixel Print that hangs in the hallway of our laboratory (see Figure 2). The print is three feet wide and six feet tall, and comprises three types of content (see Figure 4). First, the print displays a grid of 225 graphs arranged in a small multiples pattern [50]; each graph shows the recent network activity of a single machine. Second, the print displays a grid of 135 ISPs that these machines have had significant communication with. In total, this visually represents 561,000 data points describing the state of network activity. Third, this print presents paper "buttons" that enable users to invoke commands. With this application, we explored pen-only, pen + mobile, and pen +

Figure 5. The network awareness GIGAprint provides multiple levels of interaction. From top-left to bottom-right: 1) It projects an *ambient* display showing the activity of different machines in the network. 2) In *close reading*, a user may inspect the printed information. Within arm's reach, the user enters the *interaction* phase. The user may also inspect *private* information with his mobile

projector interactions; we discuss each in turn.

The pen-only interaction provides the user the ability to report an unusual machine by first tapping a paper widget (*e.g.*, "Email concern") and then selecting the machine. This interaction demonstrates an asynchronous collaborative interaction between members of the lab. In addition to the interactions which have meaning in a computational sense, users may perform any interactions they can with normal paper, such as highlighting and annotation, which may benefit asynchronous collocated collaboration, even if it is not represented digitally. This flexibility leverages an affordance of paper that earlier work [15, 47] has found to be important and a missing element in graphical user interfaces.

Second, we built ensemble interactions involving a mobile device, the digital pen, and the printed display (see Figure 5). While the computer names are anonymized on the printed visualization (a public display), a network administrator who has sufficient privileges must be able to access the details. To do this, he walks up to the print and taps on a computer with his pen. Real-time information regarding that computer is displayed on his handheld device

Third, richer interactions are available through geo-referenced integration of the paper display with a projector. In the network application, the projector acts as a targeted spotlight to provide real-time information both *ambiently* and in *response* to user queries.

This large visualization provides three levels of engagement for interacting at different distances and time scales (see Figure 5). This multi-level approach synthesizes work in the information visualization community (*e.g.*, Tree-Maps [23]) with large-scale display research in the HCI community [43, 53].

At the *ambient* level, this visualization highlights machines as they engage in network activity. The highlights fade away over time. When the user engages in *close reading*, he can read the details of the printed graphs. Upon *interacting*, the user can tap on paper buttons to explore the network traffic, and perform range queries on the map to see which machines are talking to which remote locations (*e.g.*, Europe).

Photo Waterfall

Our second application illustrates a different style of GIGAprints in that it employs the wide-format printer as a *streaming display* of online content. The photo waterfall (see Figure 6) continually prints photographs from the Flickr photo sharing website [57] as they are posted. This prototype was inspired by the visibility of work in traditional publishing houses [5] and design studios [29]. For organizations such as newspapers that operate on very short time cycles and value content such as images where finding the "right" asset for the front page or a news story can be difficult. A large-scale, high-resolution, streaming display may prove invaluable in these settings.

Users can engage the photo waterfall in three ways. As an *ambient* display, photographs are immediately viewable as they come in. Multiple users can discuss photos in front of the display in a synchronous, collocated manner. Second, the user can *select* a photograph with a digital pen to mark it as an item of interest, which *transports* it to her mobile device. Finally, the user can *write* a descriptive word beneath the photo. The photo will be tagged with that word, with a confirmation displayed on the mobile device. One advantage of this display is that users do not need to navigate web pages to explore Flickr; browsing and paging turns into a visual search task on this Interactive Gigapixel Print.

This prototype takes advantage of several ensemble interactions. First, the tap-and-transfer of photos is an example of a *transport* interaction, where a photo is transferred from the physical print to the digital display. Second, the write-and-tag action *manipulates* the photograph's metadata, providing asynchronous, remote collaboration. As these tags are created, the mobile device provides confirmation for the user.

The current prototype downloads photos from the Flickr Interestingness blog using the Flickr API. Our wide-format printer is instructed to intermittently print out a single row of new photos with Anoto pattern. The mapping between pattern and Flickr URLs is stored at a static location on the Web. When a user selects the pattern associated with a photo, GIGAprints *transports* the selected URL, loading the photograph on her mobile device (we use the OQO 01+ handheld Windows XP computer to prototype a future mobile phone).

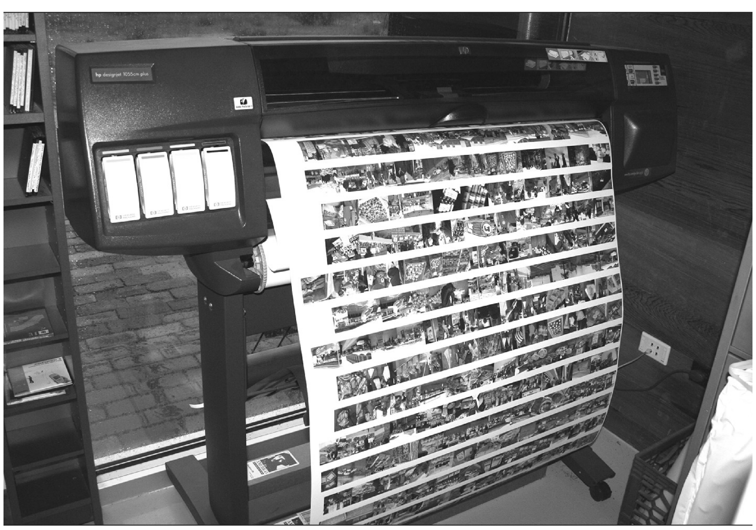

Figure 6. The photo waterfall streams photographs and enables users to take photos away with them on their mobile device.

Collaborative Sketching

The collaborative sketching application comprises Anoto paper augmented with digital projection, deployed at multiple locations. When one party sketches with the Anoto pen on their large-scale print, the ink is captured and projected at the remote party's location. This prototype—in augmenting paper displays with electronic information—offers the converse approach to prior work that augmented wall-scale *electronic* displays with *paper* artifacts such as Post-it notes [29].

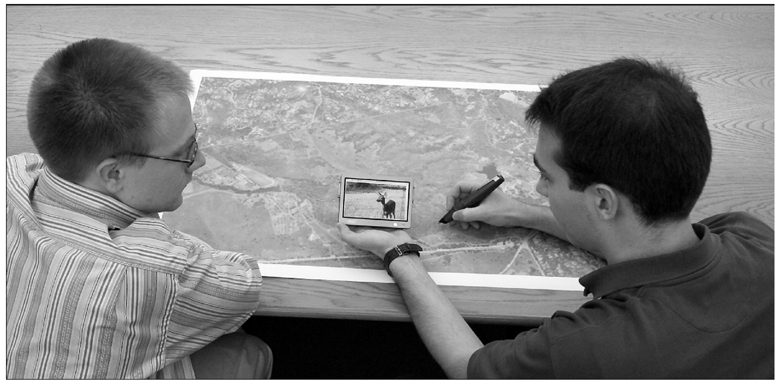

Figure 7. Two researchers use map-based queries to retrieve photos on a handheld display.

Map-based Queries

Maps have long been one of the prime benefactors of the high resolution, large size, and flexibility of print media. To explore the benefits of augmenting maps, and specifically using location data as a query, we prototyped two map-centric interactions (see Figure 7).

The first, a location-based query, enables users to *select* geo-tagged photographs from a database by circling a location on the map; the matching photographs are presented on the mobile display.

The second is a tool for building architects who want to show their clients a 3D walkthrough; the digital device in this technique is a PC. To specify the camera path of a 3D walkthrough, the user draws a line on a floor plan to direct a 3D walkthrough occurs real-time on a nearby digital display. Camera control and 3D rendering in this application are implemented using Java Robot API to control a character in a video game. The map *manipulations* are translated to keyboard and mouse movements on the mobile device.

RELATED WORK

Since the introduction of Wellner's DigitalDesk [56], there has been significant research interest in integrating physical and digital interactions. In the subsequent 15 years since this work, there has been considerable progress in terms of applications that integrate paper and digital interactions. In this section, we discuss the relationship of GIGAprints to prior work through the introduction of a design space for integrated paper-digital interactions. The design space indicates that an unexplored area prior to this paper is large-scale, computationally controlled paper displays. The design axes comprise *physical scale*, *interactions with paper*, spatial and temporal *paper-pixel coordination*, *digital modality*, *interaction proximity*, and *update frequency*. While space precludes a complete listing and discussion of all integrated paper-digital interfaces, we mention representative work in each area.

Scale

Systems that integrate paper and digital interactions occupy a range of scales, from the small to the very large; we use Weiser's "computing by the inch, foot, and yard" [55] to describe this axis.

Inch

At the inch scale, two prior efforts demonstrate how small pieces of paper can facilitate collaboration. WebStickers leveraged Post-it notes to allow users to create, store, and share Web bookmarks [18]. The user would *associate* bookmarks to the paper notes. The advantage of the bar-code-enhanced notes was that the small scale enabled users to share by placing notes at the most appropriate locations; a link to a document could be posted on the door of the person you wanted to give the document to. Second, Mackay's work studying paper flight strips for air traffic control [34] shows that thin paper strips enable efficient sharing of handwritten annotations.

Foot

Foot-scale displays are comfortable for personal reading and writing. At the foot scale, prior works include Books with Voices [28], Audio Notebook [49], and Listen Reader [2], which all enable users to retrieve media by pointing to a paper page. ButterflyNet [59] and PapierCraft [32] demonstrate that paper can be leveraged as a both a capture and association medium (using handwriting and gestures, respectively). Paper PDA [16], The XAX Paper Interface [24], and PADD [14] show how different ways to manage the transition of documents between the digital and physical worlds.

Yard

The class of yard-scale displays enables collocated collaboration and provides visual context for the displayed information. DigitalDesk [56] and URP [52] demonstrate that large projected displays can be used to provide context to the paper and other physical objects. Collaborage [36], Outpost [29], and Insight Lab [31] provide collaborative environments which enable capture of notes written on paper or whiteboards. Finally, Rasa [35] is a large-scale display that provides collaborative command and control using handwriting capture and speech interactions. GIGAprints is a yard-scale display that—in augmenting *paper* displays with *electronic* information—offers the converse approach to this prior work.

Paper-Pixel Coordination

This axis describes both the spatial and temporal relationships between paper and pixels.

Spatial

The spatial paper-pixel coordination borrows its four types from Klemmer's taxonomy [27]: geo-referenced, collocated, non-collocated and non-visual output.

In *geo-referenced* systems, graphical feedback is overlaid on top of the paper; examples of this type are the Designers' Outpost [29], Urp [52], and Rasa [35]. Fitzmaurice's work on spatially aware mobile devices [12], and subsequent Peephole [58] and Total Recall [19] enable mobile devices to augment the physical world in a geo-referenced manner. In these systems, the relative locations of the digital and physical displays were tightly-coupled, so that a user can treat a palmtop device as a window onto a larger virtual display.

In its approach of geo-referencing two unique display devices to combine their benefits, GIGAprints draws inspiration from Baudisch *et al.*'s Focus Plus Context Screens [4]. In that work, a traditional desktop display occupies the user's focus, and this display is placed in the center of a much larger but lower-resolution projector that provides contextual information.

With *collocated* systems, the physical paper and digital display occupy adjacent physical space. Examples of this type include Books with Voices [28], Paper Flight Strips [34], Palette [38], and WebStickers [33].

With *non-collocated* systems, the paper and digital displays do not occupy a shared space. Examples include Paper PDA [16] and XAX [24].

Systems with *non-visual output* include Listen Reader [2] and Audio Notebook [49]. These systems leverage alternate digital modalities, primarily speech and non-speech audio. AudioNotebook provided audio capture and retrieval based on paper-based pen input. Anoto digital pens [1] provide *vibration* output when interacting with paper-based widgets. These alternate modalities are useful when visual output is not convenient (in the case of a mobile device).

The GIGAprints applications explore all four categories on this design axis. The network awareness application demonstrates *geo-referenced* projector feedback, as well as *non-visual* audio output. The photo waterfall illustrates the benefits of a *collocated* mobile device. Lastly, the collaborative sketching application demonstrates GIGAprints for *non-collocated* collaboration.

Temporal

The temporal coordination axis is a measure of the time it takes for a paper-based action to invoke a change in the digital display. This coordination can be synchronous or asynchronous.

GIGAprints spans this axis, as it can provide synchronous paper-pixel coordination (selecting a machine brings up the network information on the mobile device), and asynchronous coordination (paper-based sketches can be saved and retrieved).

Interactions with Paper

This axis expresses how paper content is created and how the user interacts with it. Content can be *statically preprinted*, created by users *marking directly* on paper, or printed under *computational control*.

Listen Reader [2] and Books with Voices [28] are examples of interfaces where the paper content is statically preprinted. With Back *et al.*'s Listen Reader, a user controls an audio soundtrack by moving his hands over a preprinted paper book. With Klemmer *et al.*'s Books with Voices, a user leverages a handheld device to retrieve multimedia from paper transcripts of oral interviews.

Audio Notebook [49], ButterflyNet [59], PADD [14], and PapierCraft [32] leverage paper's affordances for authoring (writing and sketching directly on paper). With Audio Notebook, a student takes notes directly on paper, while audio of a lecture is recorded and associated with the handwritten notes. The student later retrieves the audio by interacting with the paper notebook. With ButterflyNet, a biologist can take field notes on paper notebooks. These notes are automatically correlated with digital photos to enhance later retrieval of the research content. With PADD and PapierCraft, users can annotate and edit documents by marking directly on paper.

Paper PDA [16] and XAX [24] enable users to create content that is then offloaded to the system for processing. In each system, the paper is considered a form that the user fills out and presents to the form processor (computer). The Anoto Forms toolkit [1] also enable similar form-filling interactions. These systems (and PADD, described earlier) enable the user to print content under computational control.

GIGAprints spans this axis. All GIGAprints are printed under *computational control*. Additionally, GIGAprints like the map are *statically preprinted*; and applications such as collaborative sketching demonstrate direct paper *marking*.

Interaction Proximity

This axis describes the how users can leverage the paper-digital interface at multiple distances, from *near* to *far*.

From a distance, users can take advantage of the ambient qualities of the printed visualization; in this case, the details of the display are not the focus of attention.

As the user approaches the paper display, he steps into the *close reading* range. At this stage, he can read text that is printed on the page, but cannot yet touch the display.

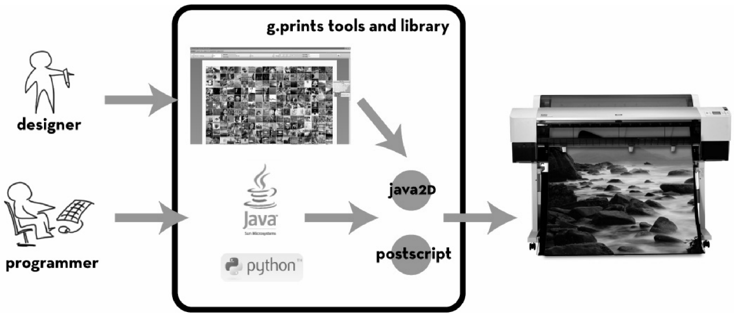

Figure 8. Designers and programmers can use the GIGAprints tools and architecture to build Interactive Gigapixel Prints.

At the nearest distances, the user enters a *direct interaction* phase, where he can leverage pen-based or other forms of direct input.

GIGAprints can span this range. Particular applications can include ambient interactions (as in the ambient highlighting in the network awareness tool), or close reading (when a user walks closer to inspect the graphs in the network awareness tool), or direct interaction (when a user employs the pen to query for computers).

Vogel and Balakrishnan's work [53] couples a user's proximity with the transition of the graphical display from public to private access. Instead of tying distance to public/private, GIGAprints instead uses the mobile device to provide the partition. GIGAprints also leverages multi-level visualizations (*e.g.*, Tree-Maps [23]) to provide the *ambient*, *close reading*, and *direct interaction* phases.

Update Frequency

As we have introduced printing as a display option, one can now think of displays in terms of their update frequency, which can vary from *none* (*e.g.*, a static display), to *intermittent*, to *real-time*. A wall-sized printed visualization might never be updated. In contrast, projector displays (*e.g.*, [9, 44]) refresh at interactive rates. Many of our examples lie in between, as GIGAprints can be thought of as a high-spatial resolution, low temporal resolution display. These prints can be updated very slowly (*e.g.*, once per day). They can also refresh slowly throughout the course of a day. To refer to the various update frequencies of paper-based displays, we use terms such as *streaming displays* (for seismograph-like update rates) or *static posters* (for displays that rarely, if ever, change).

Collaboration

One of the most important benefits of a large interactive surface is its affordances for collocated collaboration. Collaboration can be remote/collocated, and synchronous/asynchronous. ClearBoard [22] enabled synchronous remote shared drawing. Designers' Outpost enabled synchronous, collocated design work around a shared wall-sized display [29] (a later version supported remote collaboration). GIGAprints have the potential to support these types of collaborations, including asynchronous ones. One example of an asynchronous collocated interaction is a bulletin-board model of interacting with a printed display. One user can leave a voice note on the bulletin board by recording into his mobile phone, and writing a short description on the large paper surface. Later on, friends can retrieve the voice note by tapping on that handwritten note, and listening to their mobile phone replay the message. People who are not friends are barred from retrieving the audio.

Related Taxonomies

The design space presented here draws on related synthetic work on taxonomies for tangible user interfaces [11, 12, 20, 27, 48, 51]. The framing is most similar to Klemmer's work [27], with the distinction that our design space is concerned with the *interaction techniques and technologies* rather than the *software architecture* used to implement these systems. Ullmer and Ishii's MCRit model —which describes the division of an interface into digital and physical components—is subsumed by our two paper-digital spatial coordination axis. Fishkin's suggests two primary concerns for tangible interaction: *metaphor* and *embodiment*. The embodiment axis is also subsumed by paper-digital coordination, and the metaphor axis—describing the relationship between the interface and an analogous real-world behavior — while useful for describing user *experience*, is less relevant for describing the *interface* embodying that experience. Similarly, Shaer *et al.*'s TAC syntax provides a useful vocabulary for describing the *application logic* of a TUI, which again is a largely orthogonal concern to the *interaction techniques*.

IMPLEMENTATION AND TOOLS

The GIGAprints infrastructure enables programmers and designers to design and build Interactive Gigapixel Prints (see Figure 8).

GIGAprints provides programmers with two rendering engines to generate the paper display—Java to print Java2D, and Python to print PostScript. GIGAprints renders the Anoto pattern and manages the associations between content, pattern coordinates, and actions. Our architecture provides two ways to manage the association between locations on the printed surface, the content (*e.g.*, photos, graphs), and the types of actions that bind the two. First, the application software can directly manage a low level list of printed widgets and their associated actions (*e.g.*, button-click, range selection, slider). Pen information is streamed in real-time to the application, which can invoke the action logic. Alternatively, the application software can leverage an API using callbacks to monitor higher level events. An application registers listeners for these paper-based actions (built on top of Java Swing's Action architecture). Our low-level pen listeners and our high-level action listeners are managed using a pen server and an action server. Multiple clients (both local and remote) can connect to these servers

Figure 9. The GIGAprints designer provides techniques for managing paper-based Wizard-of-Oz user studies. Here, the Wizard, *right,* is controlling an interface for a participant, *left*. The Wizard's interface includes a remote action command, which triggers behaviors on the participant's mobile device.

to register themselves as listeners. This enables events generated on one device to trigger actions on another, an important capability to facilitate ensemble interactions. Because the event generators and event listeners are decoupled, we can support real-time input from multiple people (*i.e.*, multiple pens) interacting on one printed display, or multiple printed displays (remote collaboration).

The libraries we have developed enable programmers to build interactive paper-based visualizations. In addition, this architecture serves as the beginnings of a toolkit for interactive printed interfaces. To achieve this, we will need to refine our abstractions into efficient and usable techniques for building interactive paper-based visualizations.

The GIGAprints designer

Creating a GIGAprints application entails two tasks. The first is setting up the device ensemble and creating a physical print that is augmented with the Anoto pattern. The second is specifying the *digital behavior* invoked by the print, and – if needed — developing the *computational control* of the printing. To facilitate these tasks, we have created the GIGAprints designer. As the infrastructure setup and Anoto pattern are common across all our GIGAprints, this step is handled automatically by the design tool. As specifying the application logic is unique to each GIGAprint, the design tool enables designers to control the *correspondence* between pen input and application behavior through a Wizard of Oz interface [8, 25, 30].

To use this WOz technique, the designer starts the WYSIWYG interface designer (see Figure 9), and imports graphics that he or she has designed in an external graphical editor. During a use test, the designer can invoke actions by clicking on a tray of possible actions (more can be defined through the GUI). This tool contributes a technique for capturing and analyzing paper-based user tests. As the user interacts with the paper interface, the clicks and strokes are displayed on the designer's wizard view, and are captured for future analysis. It works for testing paper-digital interfaces, and also paper prototypes of traditional GUIs [45].

CONCLUSIONS AND FUTURE WORK

We have presented three main contributions. First, we introduced Interactive Gigapixel Prints, a class of paper-based interfaces where large printed surfaces are augmented with digital displays. Second, we introduced ensemble interaction techniques that enable users to leverage the benefits of both the print and digital media, and illustrated these techniques through four example applications. Third, we articulated a design space to describe how GIGAprints relate to prior integrated paper-digital interfaces.

The design tool and Java framework we developed enable designers and programmers to prototype and build Interactive Gigapixel Prints. As future work, we are exploring improved tool support for creating Interactive Gigapixel Prints. Our software is open source, and available at http://hci.stanford.edu/gigaprints.

ACKNOWLEDGEMENTS

We thank Terry Winograd for his insightful comments; John Gerth for his advice on the network monitoring application; Mor Naaman and Trevor Hebert for helping us acquire data sets; Jory Bell for the OQO; and the members of the Stanford Graphics Lab for providing great feedback on our ideas. NSF Grant IIS-0534662 supported this work.

REFERENCES

1 Anoto AB, *Anoto Technology*. http://www.anoto.com

2 Back, M., J. Cohen, R. Gold, *et al.* Listen Reader: an Electronically Augmented Paper-based Book. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 23–29, 2001.

3 Ballagas, R., A. Szybalski, and A. Fox. Patch Panel: Enabling Control-flow Interoperability in Ubicomp Environments. In Proceedings of *PerCom: Pervasive Computing and Communications*. p. 241–52, 2004.

4 Baudisch, P., N. Good, and P. Stewart. Focus Plus Context Screens: Combining Display Technology with Visualization Techniques. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 31–40, 2001.

5 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.

6 Bertin, J., *Semiology of Graphics*: University of Wisconsin Press. 415 pp. 1983.

7 Card, S. K., J. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*: Morgan Kaufmann. 712 pp. 1999.

8 Dahlbäck, N., A. Jönsson, and L. Ahrenberg. Wizard of Oz Studies: Why and How. *IUI: International Conference on Intelligent User Interfaces*. pp. 193–200, 1993.

9 Dietz, P., R. Raskar, S. Booth, *et al.* Multi-projectors and Implicit Interaction in Persuasive Public Displays. *AVI: Working Conference on Advanced Visual Interfaces*. pp. 290–17, 2004.

10 Edwards, W. K., M. W. Newman, J. Z. Sedivy, *et al.* Using Speakeasy for Ad Hoc Peer-to-Peer Collaboration. *CSCW: ACM Conference on Computer Supported Cooperative Work*. pp. 256–65, 2002.

11 Fishkin, K. P. A Taxonomy for and Analysis of Tangible Interfaces. *Personal and Ubiquitous Computing* **8**(5). pp. 347–58, 2004.

12 Fitzmaurice, G. W. Situated Information Spaces And Spatially Aware Palmtop Computers. *Communications of the ACM* **36**(7). pp. 39–49, 1993.

13 Gigapxl Project, *Gigapxl Project*. http://www.gigapxl.org

14 Guimbretière, F. Paper Augmented Digital Documents. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 51–60, 2003.

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 Hinckley, K., G. Ramos, F. Guimbretiere, *et al.* Stitching: Pen Gestures that Span Multiple Displays. *AVI: the Working Conference on Advanced Visual Interfaces*: ACM Press. pp. 23–31, 2004.

18 Holmquist, L. E., J. Redström, and P. Ljungstrand. Token-Based Access to Digital Information. *Handheld and Ubiquitous Computing*: Springer-Verlag, 1999.

19 Holmquist, L. E., J. Sanneblad, and L. Gaye. Total Recall: In-place Viewing of Captured Whiteboard Annotations. *Extended Abstracts of CHI*. pp. 980–81, 2003.

20 Hornecker, E. and J. Buur. Getting a Grip on Tangible Interaction: A Framework on Physical Space and Social Interaction. *CHI: ACM Conference on Human factors in computing systems*: ACM Press, 2006.

21 Imhof, E., *Cartographic Relief Presentation*: Walter De Gruyter, Inc. 388 pp. 1982.

22 Ishii, H. and M. Kobayashi. ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact. *CHI: ACM Conference on Human Factors in Computing Systems*, 1992.

23 Johnson, B. and B. Shneiderman, Tree-Maps: a space-filling approach to the visualization of hierarchical information structures, in *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc. pp. 152-59, 1999.

24 Johnson, W., H. Jellinek, L. K. Jr., *et al.* Bridging the Paper and Electronic Worlds: The Paper User Interface. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 507–12, 1993.

25 Kelley, J. F. An Iterative Design Methodology for User-Friendly Natural Language Office Information Applications. *ACM Transactions on Office Information Systems* **2**(1). pp. 26–41, 1984.

26 Kerouac, J., *On the Road*. Reprint ed: Penguin Books. 307 pp. 1991.

27 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

28 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.

29 Klemmer, S. R., M. W. Newman, R. Farrell, *et al.* The Designers' Outpost: A Tangible Interface for Collaborative Web Site Design. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 1–10, 2001.

30 Klemmer, S. R., A. K. Sinha, J. Chen, *et al.* SUEDE: A Wizard of Oz Prototyping Tool for Speech User Interfaces. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 1–10, 2000.

31 Lange, B. M., M. A. Jones, and J. L. Meyers. Insight Lab: an Immersive Team Environment Linking Paper, Displays, and Data. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 550–57, 1998.

32 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.

33 Ljungstrand, P., J. Redström, and L. E. Holmquist. WebStickers: Using Physical Tokens to Access, Manage, and Share Bookmarks to the Web. *DARE: Designing Augmented Reality Environments*: ACM Press. pp. 23–31, 2000.

34 Mackay, W. E., A.-L. Fayard, L. Frobert, and L. Médini. Reinventing the Familiar: Exploring an Augmented Reality Design Space for Air Traffic Control. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 558–65, 1998.

35 McGee, D. R., P. R. Cohen, R. M. Wesson, and S. Horman. Comparing Paper and Tangible, Multimodal Tools. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 407–14, 2002.

36 Moran, T. P., E. Saund, W. V. Melle, *et al.* Design And Technology For Collaborage: Collaborative Collages Of Information On Physical Walls. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 197–206, 1999.

37 Myers, B., S. E. Hudson, and R. Pausch. Past, Present, and Future of User Interface Software Tools. *ACM*

Transactions on Computer-Human Interaction **7**(1). pp. 3–28, 2000.
38 Nelson, L., S. Ichimura, E. R. Pedersen, and L. Adams. Palette: A Paper Interface for Giving Presentations. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 354–61, 1999.
39 Nokia, *Nokia SU-1B Digital Pen*. http://www.nokia.com
40 OQO, *model 01+*, 2006. http://www.oqo.com
41 Pering, T., R. Ballagas, and R. Want. Spontaneous Marriages of Mobile Devices and Interactive Spaces. *Communications of the ACM* **48**(9). pp. 53-59, 2005.
42 Phan, D., L. Xiao, R. Yeh, *et al.* Flow Map Layout. *InfoViz: IEEE Symposium on Information Visualization*: IEEE Computer Society. pp. 29-38, 2005.
43 Prante, T., C. Rocker, N. Streitz, *et al.*, Hello. Wall—Beyond Ambient Displays., in *Video and Adjunct Proceedings of UBICOMP Conference.* 2003.
44 Raskar, R., P. Beardsley, J. v. Baar, *et al.* RFIG Lamps: Interacting With a Self-Describing World Via Photosensing Wireless Tags and Projectors. *SIGGRAPH: ACM Conference on Computer Graphics and Interactive Techniques*. pp. 406-15, 2004.
45 Rettig, M. Prototyping for tiny fingers, *Communications of the ACM*, vol. 37(4): pp. 21-27, 1994.
46 Schilit, B. N. and U. Sengupta. Device Ensembles. *Computer* **37**(12). pp. 56–64, 2004.
47 Sellen, A. J. and R. H. R. Harper, *The Myth of the Paperless Office*. 1st ed: MIT Press. 242 pp. 2001.
48 Shaer, O., N. Leland, E. H. Calvillo-Gamez, and R. J. K. Jacob. The TAC Paradigm: Specifying Tangible User Interfaces. *Personal and Ubiquitous Computing* **8**(5). pp. 359-69, 2004.
49 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.
50 Tufte, E. R., *The Visual Display of Quantitative Information*. 2 ed: Graphics Press. 197 pp. 2001.
51 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.
52 Underkoffler, J. and H. Ishii. Urp: A Luminous-Tangible Workbench For Urban Planning And Design. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 386–93, 1999.
53 Vogel, D. and R. Balakrishnan. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. *UIST: ACM Symposium on User Interface Software and Technology*. pp. 137–46, 2004.
54 Ware, C., *Information Visualization: Perception for Design*: Morgan Kaufmann. 486 pp. 2004.
55 Weiser, M. The Computer for the 21st Century. *Scientific American*. pp. 94–104, 1991.
56 Wellner, P. Interacting With Paper on the DigitalDesk, *Communications of the ACM*, vol. 36(7): pp. 87–96, 1993.
57 Yahoo! Inc., *flickr*. http://flickr.com/
58 Yee, K.-P. Peephole Displays: Pen Interaction on Spatially Aware Handheld Computers. *CHI: ACM Conference on Human Factors in Computing Systems*. pp. 1–8, 2003.
59 Yeh, R. B., C. Liao, S. R. Klemmer, *et al.* ButterflyNet: A Mobile Capture and Access System for Field Biology Research. *CHI: ACM Conference on Human Factors in Computing Systems*, 2006.