Research

The overarching goal of my research is to bring the ease of use of pen and paper interactions to computer interfaces. Pen and paper are an extremely versatile tool used extensively by knowledge workers when sketching new concepts, exploring a design space by quickly sketching several variations, brainstorming during a meeting, or simply proofreading documents. Pen and paper interactions are rapid, fluid, and almost transparent to the user. Unfortunately, work captured on paper is often difficult to transfer back into the digital world where powerful computational resources are accessible. Through my work, I have demonstrated that the ease of use of pen and paper interactions and access to digital resources can be smoothly bridged. Blending hardware prototypes, software implementations, and empirical evaluation, I have proposed novel pen interfaces for a wide variety of digital surfaces, including: large wall-sized displays such as the Stanford Interactive Mural; portable notepad-sized systems such as tablet PCs; and digital pens such as the Anoto system which can record and process strokes made on a special pre-printed paper. My work establishes a unified framework for pen computing applicable to mixed paper-digital settings as well as a wide variety of future digital interactive surfaces such as "digital wall paper" or electronic paper (e-paper), an area I have been exploring since I moved to Cornell. My work has been mostly supported by grants from the NSF (about $2.5M in total including a CAREER award), has been published in first tier ACM conferences and journals, and had a significant impact (my H-index is 19[1]).

Pen interactions for digital displays

The standard desktop interfaces based on Windows, Icons, Menus, and Pointers (WIMP) have been extremely successful. Yet, they are not well adapted to direct pen interactions. Many WIMP interactions that were originally developed for the mouse are difficult to perform with a pen. A prime example is the double-click: while easy to perform with a mouse (since the pointer is stable), it presents a difficult task in pen-based interfaces. Other problems in adapting the WIMP paradigm to pen-based interfaces include occlusions (created by the user's hand when interacting directly on the screen), difficulties in using modifier keys (such as pressing "shift" to extend the current selection), and limited access to keyboard shortcuts, which are crucial for expert performance.

To address these problems, my research aims to design and study new interaction techniques that are well-adapted to the unique hardware configurations offered by pen-based interfaces. After my initial work on the wall-sized **Stanford Interactive Mural** (which has been cited 259 times and made its way into textbooks), my recent research has focused on portable platforms such as tablet PCs and digital pens.

Goal crossing interfaces

Just like checking a checkbox on paper, crossing a small target (or goal) on a screen is a very natural gesture in pen-based settings. This intuitive insight was confirmed by empirical studies which showed that crossing targets with a digital pen can be as efficient as pointing and clicking buttons with a mouse. Yet, it remained unclear if it was possible to design a complete interface solely relying on goal crossing. **CrossY**, a drawing application in which all interface elements (such as menus, scrollbars, and dialog boxes) use goal crossing as their primary interaction mechanism, was the first to demonstrate that this could be done. CrossY not only illustrated the feasibility of crossing as an interaction paradigm in a real life applications, it also provided initial feedback on the unique characteristics of crossing-based interfaces. We found that goal crossing offers designers more flexibility than the equivalent point-and-click interfaces, and encourages the fluid composition of commands—a feature unique to goal crossing interfaces. Further empirical evaluation showed that by relaxing the semantic of crossing interactions, crossing-based interfaces could be faster than point-and-click interactions. Published at UIST 2004, where it received the Best Paper Award, CrossY has been cited 104 times.

Efficient pen interfaces for expert users

Like traditional point-and-click interfaces, CrossY places command selection mechanisms (such as toolbars and menus) at the periphery of the work area. This implies that each command selection incurs

[1] Citation counts refer to Google Scholar reports.

the cost of a round-trip between the work area and the peripheral interface elements. Traditional desktop interfaces have alleviated this problem by relying extensively on the keyboard for expert interactions (for example, through keyboard shortcuts or modifier keys). However, this solution is not feasible for pen-based interfaces. In collaboration with Ken Hinckley at Microsoft Research, we have therefore developed **Scriboli**, an interaction test bed focusing on efficient command selection for expert users of pen-based interfaces. At the heart of our approach are two key features: First, a "command" button on the side of the display distinguishes between input intended as ink and input intended as commands. This offers a reliable way for Scriboli to identify user commands and is both easy to learn and efficient for expert users. Second, Scriboli uses a pigtail gesture to distinguish between the part of a command gesture selecting the arguments of a command (scope selection), and the part selecting the command itself. The pigtail gesture relies on motor programs similar to writing and offers a quick and reliable way to separate the scope selection and the command selection phase. We demonstrated experimentally that this approach is as fast as the best known alternative method. Scriboli can also handle complex, disconnected selections, reliably access a large set of commands, and it allows users to adjust command parameters (such as nudging the exact position of a moved object). This represents a significant improvement over previous pen-based systems. Published at CHI 2005, this work has been cited 112 times.

Towards a better understanding of bi-manual interactions

The techniques presented so far have only considered one-handed interactions. Of course, given the prevalence of two-handed activities in everyday life, it is natural to also consider two-handed interaction techniques for pen-based interfaces. As part of our exploration of pen interfaces, we conducted an in-depth analysis of the specific interface features representing strengths and weaknesses of two-handed techniques (e.g., Bier et al.'s ToolGlass). On the one hand, ToolGlass offers a speed advantage in some tasks, as it encourages users to select a command and adjust command parameters in the same stroke. On the other hand, the two-handed setting slows down users in some situations as it forces them to track both the tool selection palette and the primary cursor at the same time. Our work indicated that the advantages of merging command and parameter selection are not limited to two-handed settings, but can be realized in one-handed settings. In fact, we found that one one-handed technique, control menu was consistently faster than ToolGlass. Merging command and parameter selection later became an important feature of the Scriboli system described above. This work was published by ToCHI and cited 35 times.

Pen interactions for paper displays

Despite the advances in designing smaller and more powerful tablet-like computers, knowledge workers still rely heavily on paper. At the root of this apparent paradox are the different sets of affordances offered by paper (e.g., ease of navigation and annotation, high information density display) versus digital documents (e.g., ease of distribution, archiving, search). These differences encourage users to switch back and forth between the two media depending on their needs. Thus, it is important to consider ways to optimize the *cohabitation* between the two media. In our approach, based on digital pens, which can capture strokes made on special pre-printed paper, we capture all the annotations drawn on paper printouts and merge them back into their original documents. This brings the paper world and the digital world onto an equal footing: paper and computers are simply two different ways to interact with **Paper Augmented Digital Documents** (PADDs) during their life cycle. In the digital realm, PADDs offer all the digital affordances, but require the use of a computer to access them. In the paper realm, PADD printouts act as proxies of the corresponding digital document and can record marks and commands performed with a digital pen. Because PADD printouts only require a pen-sized computer to capture strokes, they offer all the affordances of paper. The PADD approach is well adapted to many activities that currently rely heavily on paper, such as proofreading, editing drafts, and annotating large format documents, such as blueprints. Published at UIST 2003, this work has been cited 144 times.

Pen-top interfaces for paper-based interactions

While the original PADD system focused on the capture and management of strokes made on printouts, it quickly became apparent that it could support a wider range of tasks if users could issue commands while interacting with printouts. To address this need, we developed **PapierCraft**, a pen-based marking interface designed specifically for passive media such as paper. Drawing from our experience on digital surfaces (such as Scriboli), PapierCraft is the first paper-based system to support active reading, including copying and pasting information from one PADD document to another, creating links between content found in two different paper documents, "stitching" two paper documents together, or even searching for a given word in a printout. Thus, PapierCraft combines the advantages of paper with those of digital annotation systems such as Microsoft OneNote. In its basic form, this interface does not require *active* feedback beyond the ink laid on the paper during pen interactions. Using modalities readily available on pen-top computers[2] such as tactile feedback, multi-color LED, and voice feedback, our paper-based interface can support features usually found on dynamic media such as command name discovery and easy error recovery. We showed empirically that this reduces the error rates for commands issued on paper to a level similar to those observed for dynamic media. In a follow-up study exploring the use of PapierCraft in proofreading task, we showed that participants perceived PapierCraft as combining the advantages of paper and digital media supporting our initial design goal. The PapierCraft system was described in a series of papers (UIST'05, UIST'06, ToCHI'08, ToCHI'12) together cited 157 times. Principles developed in PapierCraft were adapted for several systems including ButterflyNet, a field note management system for biologists developed in collaboration with Scott Klemmer's group at Stanford University (presented at CHI'06 and cited 135 times), and a paper interface for Classroom Presenter, an active learning system developed in collaboration with Richard Anderson's group at the University of Washington (presented at INTERACT'07 and cited 13 times).

Micro-projector augmented Pen interface

The PapierCraft system was designed for active reading tasks where most activities are carried out on standard A4 or Letter paper format. Another area in which paper is used extensively is architecture and construction. In that application area, large format (A0) sheets are used as references of what needs to be built and as a communication medium between the different parties at a construction site. Besides the obvious size differences, the use of paper is much more data intensive: a plumber, for example, might want to integrate her set of drawings with electrical information from the Building Design Database to check for possible incompatibility. In such settings, a standard pen-top interface with its limited feedback options becomes very tedious to use. In collaboration with AutoDesk, my group focused on two approaches to address this problem. With the **PenLight** system, we explored a configuration in which a small projector is included inside the barrel of the digital pen. With the **MouseLight** system, we explored how a location aware projector could be used in conjunction with a digital pen to implement the first version of the ToolGlass system for paper-based interactions. While MouseLight requires users to carry two pieces of equipment (the MouseLight and the digital pen), it has proven to be much more flexible than PenLight in terms of interface design. Both systems have been very well received. PenLight was presented at CHI'09 and has already been cited 24 times, while MouseLight was presented at CHI'10 where it was nominated for best paper award and has already been cited 12 times. We are now working on the final steps of a study comparing these two systems with more traditional PDA based approaches like the A-book system.

Pen interfaces for 3D models

As in the case of paper documents, physical models used in architecture and engineering are often annotated extensively to indicate needed changes or to correct errors. The production of physical models was once lengthy and expensive, but the advent of 3D printing technology has made this process cheaper and faster. This has changed established design practices by encouraging a faster iteration process between creating ideas and testing them in a concrete setting. As a result, the costs of transferring information from the tangible model to its digital equivalent are becoming more apparent. We demonstrated that the interaction techniques developed for paper documents are readily applicable to this nascent field. Developed in collaboration with Hod Lipson at Cornell, our **ModelCraft** system allows

[2] Pen-top computers are digital pens which can be programmed (such as the Fly pen computer introduced by LeapFrog)

users to capture annotations and edits (such as cuts or extrusions) made on 3D models using digital pens. This approach is highly scalable in terms of the number of objects being tracked and the number of pens being used and, in contrast to conventional tracking systems, it does not require additional tracking infrastructure. Through collaboration with the UMD School of Architecture, we established empirically that this approach represents a significant improvement over current practices during the early stages of the Architecture curriculum. The ModelCraft system was described in a series of three papers (UIST'06, INTERACT'07, ToCHI'09) together cited 26 times.

Moving beyond linear reading for ebook-readers

The advent of bi-stable display technology such as the E-ink system has started to change the way we read. E-ink is a very low power, reflective display technology making it possible to build light e-book readers like the Amazon Kindle. Recently, Amazon announced that it was selling more digital books than paper books, suggesting that the system is now good enough to support effective linear reading. At the same time, several studies have shown that the Kindle is currently inadequate for more active forms of reading, because it does not support effective navigation and annotation. Drawing on our understanding of digital and paper based interactions, we collaborated with Abi Sellen from Microsoft Research to develop **United Slates,** a new system to support active reading on ebook readers. Key to our system is the insight that the success of paper lies in the ease of annotations and the availability of multiple displays surface at any given time. In collaboration with an industrial partner, we developed a new ebook reader system extending our early work on dual-display systems (CHI'08 and CHI'09) toward a multi-slates system. On the hardware side, the system is lighter and thinner than the iPad yet includes a Wacom pen tracking system. On the software side, we developed the equivalent of a windows manager for a federation of slates. Using this system, users can easily transfer documents from one slate to the other, use several slates together to simplify navigation through multiple documents or link two slates to simplify creating and reviewing annotations. To alleviate the need for carrying multiple slates, a stacks manager lets users easily remap group of documents to slates to accommodate the current number of available slates. To simplify content creation, a laptop can join the federation of slates and call slate content directly on its screen. This greatly simplifies the creation of excerpts by avoiding a time consuming switch between a keyboard+touchpad based interface and a pen-based interface. A preliminary evaluation of the first version of the system received very positive feedback from participants. The description of this system is currently under review at ToCHI, In collaboration with Cornell libraries we are now preparing for a broader deployment to analyze the potential of multi-slate systems in support of active reading.

Toward very low power information appliances

Multi-slate systems such as **United Slates** will only be practically viable if each slate has an extended battery life. Of course, bi-stable display technology offers great energy savings, but we believe that even lower energy footprints can be achieved. Compared to conventional displays, the main energy draw in bi-stable displays comes from the processor, not the display itself. Thus, optimizing the *power efficiency* of human computer interaction techniques becomes increasingly important. In collaboration with Rajit Manohar at Cornell, we have started to explore this new area of research. Our first prototype, based on a highly asymmetric dual-processor design, demonstrated up to **3.2** times longer battery life for active reading tasks such as annotations. These initial results were obtained using a low frame rate E-Ink display. We are currently developing a new prototype using a high frame rate bi-stable display to explore the power efficiency of different types of interaction techniques for document navigation and command selection. This prototype will also gather user traces that will be used to tune a very low power asynchronous processor we are planning to manufacture in the near future. By combining a bi-stable display, a low power asynchronous processor, and customized interaction techniques that realize their respective power-savings potential, we are hoping to create a new class of information appliances that let users enjoy today's levels of interactivity at a fraction of the energy cost. We believe that a similar approaches could be applied to interaction techniques for laptop and desktop computers in order to significantly lower the carbon footprint of everyday computing.