Research Statement

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 pen-sized systems such as digital pens 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 am starting to explore. My work has been published in first tier ACM conferences and journals, and has been supported by grants from the NSF (including a CAREER award) and Microsoft Research (through a grant to HCIL at the University of Maryland).

Pen interaction 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 on the increasingly prevalent digital boards, tablet-sized computers, or personal digital assistants. 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 with more than 150 citations[1] is the 15th most cited UIST paper, and was described in a textbook), I have recently focused on portable platforms such as tablet PCs.

Goal crossing interface

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. My group’s work on *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 can be done. CrossY not only illustrated the feasibility of crossing as an interaction paradigm in a real life application, 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

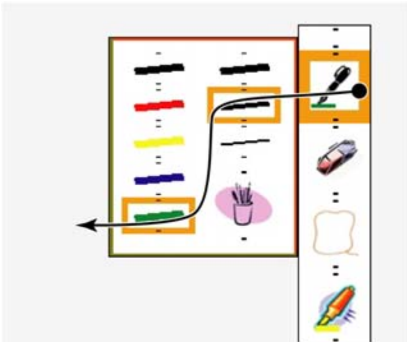

Figure 1: Crossing-based interfaces allow users to select a pen tool, stroke width, and color in a single stroke

[1] Citation counts refer to Google Scholar reports.

to goal crossing interfaces. For example, to select a pen tool and set the line thickness and color, users first cross through the pen icon and then through the desired width and color in a single gesture (as shown in Figure 1). Published at UIST 2004, where it received the Best Paper Award, CrossY has already been cited 35 times. Building on this earlier work, we are currently working on a systematic evaluation of the layout parameters that influence the performance of crossing-based interfaces.

Efficient pen interface 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 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 (Figure 2). 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. In it simplest form, Scriboli lets users select a target and a command in a single gesture (Figure 2). 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 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 already been cited 29 times.

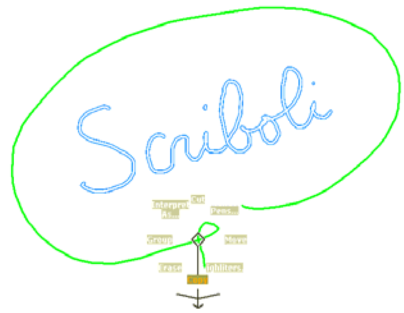

Figure 2: Copying ink in Scriboli: After pressing the command button, users lasso the target, draw a pigtail to switch to command selection, and select the "copy" option.

Pen interaction for static 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. We believe that this will be the case for the foreseeable future until improved technologies such as electronic paper take hold. 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 already been cited 53 times.

Pen-top interface 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. Principles developed in PapierCraft were adapted for several systems including ButterflyNet, a field note management system for biologists (in collaboration with Scott Klemmer's group at Stanford University), and a paper interface for Classroom Presenter, an active learning system (in collaboration with Richard Anderson's group at the University of Washington).

Pen-top interface 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 cost of transferring information from the tangible model to its digital equivalent is becoming more apparent. We demonstrated that the interaction techniques developed for paper documents are readily applicable to this nascent field. Our *ModelCraft* system allows 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.

Empirical evaluation

The techniques presented so far have only considered one-handed interactions. Of course, given the prevalence of two-handed activities in everyday life, it would have seemed 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 as well. 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.

One of the main difficulties in conducting any empirical evaluation in HCI is the intrinsic variability of human interaction data which may mask important general trends. For example, participants may not follow the instructions correctly, or it may be difficult to identify the most prevalent patterns of use. With

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

the large number of interaction samples gathered in today's experiments it is often impractical to reliably identify these patterns through manual scoring. *ExperiScope*, a new empirical data analysis tool (developed in collaboration with Ken Hinckley at Microsoft Research), addresses this problem by clustering typical interaction patterns and visualizing them graphically to simplify their comparison. For example, ExperiScope proved valuable in the analysis of the complex datasets we collected during a study of spring loaded command invocation. In particular, our tool made it easy to establish that in some contexts, selecting a command at the border of the screen can be faster than using a pop-up menu system, a very surprising finding. In addition to its utility for seasoned experimenters, we plan to explore the potential of ExperiScope to introduce students to experimental research

Toward a convergence of paper and digital interaction

A new generation of displays based on bi-stable display technology promises to deliver e-paper with a high refresh rate. These displays provide a unique blend of affordances: like Tablet PCs, they can provide direct feedback; like paper, they do not consume energy to maintain their content. They also can be very thin and flexible, and they could be built in very large formats. With these new displays, it will be possible to design pen computing environments in which several thin, flexible, high-resolution slates are connected wirelessly to a small computing hub. As with paper documents, users will be able to easily spread out several displays and enjoy direct access to computational resources on all of them. Users will have access to the fluidity of today's paper interactions in a fully interactive setting. I believe such a new display technology has the potential to revolutionize pen-computing and fluid interaction interfaces, and I will focus my efforts in the near future on demonstrating its potential.