Research

Pen and paper are extremely versatile tools and they are used extensively by knowledge workers when sketching new concepts, exploring a design space by quickly drafting multiple iterations, brainstorming during a meeting, or simply proofreading documents. Pen and paper interactions are rapid, fluid, and almost transparent to the user. My research aims to design interfaces which integrate the ease of use and expressiveness of pen and paper with the digital and computational affordances that have become fundamental to modern knowledge workers. Blending hardware prototypes, software implementations, and empirical evaluations, I have proposed novel pen interfaces for a wide variety of digital surfaces, including large wall-sized displays such as the Stanford Interactive Mural; notepad computers such as the Microsoft Surface; and digital pens such as the Anoto system which can record and process strokes made on a special pre-printed paper. More recently, I have extended my work into the 3-dimensional space to explore how the nascent 3D printing revolution can support reflective design practices. The goal is to provide novices and expert alike with 3D printed physical instantiations of their digital designs throughout the creative process. In combination, these lines of research build towards a future in which the boundaries between physical and digital world and between human and computer increasingly blur and disappear.

PEN INTERACTIONS FOR DIGITAL DISPLAYS

My earliest work focused on pen interactions for digital displays. My thesis introduced the wall-sized **Stanford Interactive Mural** (UIST'01) and described **MilleFeuille**, an interaction toolkit for WireGL, one of the first distributed OpenGL rendering systems. Later I focused on slates (then called TabletPCs) and explored a suite of interaction techniques specifically designed for pen-based interactions. With **CrossY** (UIST'04 – Best Paper Award), we demonstrated that it was possible to design a drawing application in which all interface elements (such as menus, scrollbars, and dialog boxes) use goal crossing as their primary interaction mechanism. 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.

As part of our exploration of pen interfaces, we also 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). Through systematic analysis of user performance, we demonstrated that while ToolGlass improves user performance by merging command selection and parameter selection in the same stroke, it can also slow users down because it forces them to track both the tool selection palette and the primary cursor at the same time. Our work (ToCHI'05) 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.

My ongoing efforts to optimize command selection for pen interfaces culminated in **Scriboli** (in collaboration with Ken Hinckley at Microsoft Research, CHI'05), an interaction test bed focusing on efficient command selection for expert users of pen-based interfaces. 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 of a stroke. With the recent revival of pen-enabled devices such as the iPad pro and the Microsoft surface, such insights are likely to have a renewed impact on the field.

DIGITAL SUPPORT FOR ACTIVE READING

Over the past 15 years, my students and I have focused on bridging the gap between digital and paper-based interactions for knowledge workers. During that time span, we have witnessed amazing technological progress including ever-increasing miniaturization along with the advent of "always on" connectivity, which offers access to almost unlimited computational resources from even the smallest of devices. At each step along the way, we have demonstrated how emerging technological advances can be leveraged to make annotation interfaces ever more fluid and transparent to reviewers of digital documents.

Paper-based interactions

In the early 2000s, knowledge workers still relied heavily on paper printouts to edit iterative drafts of their digital documents. In doing so, they capitalized on the unique 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). Thus, users switched back and forth between the two media depending on their needs. Since digital tablets were still too bulky and too low resolution to offer both sets of affordances, we focused on supporting *cohabitation* between paper-based and digital documents. Specifically, we leveraged the potential of digital pens, which can capture strokes made on special pre-printed paper, to capture annotations drawn on paper printouts and merge them back into their original digital documents. This brought the paper world and the digital world onto equal footing: Paper and computers were simply two different ways to interact with **Paper Augmented Digital Documents** (PADDs, UIST'03) during their life cycle. In the digital realm, PADDs offered all the digital affordances, but required the use of a computer to access them. In the paper realm, PADD printouts acted as proxies of the corresponding digital document and could record marks and commands performed with a digital pen. Because PADD printouts only required a pen-sized computer to capture strokes, they offered all the affordances of paper. The PADD approach was well adapted to many activities that relied heavily on paper, such as proofreading, editing drafts, and annotating large-format documents, such as blueprints.

With the basic cohabitation mechanism in place, we then introduced **PapierCraft** (UIST'05, ToCHI'08), a pen-based marking interface designed specifically for passive media such as paper. Drawing from our experience on pen-based commands for digital surfaces (such as CrossY and Scriboli), PapierCraft was the first paper-based system to support a wide range of commands 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 combined 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. When extended with modalities readily available on pen-top computers[1] such as tactile feedback, multi-color LEDs, and voice feedback, PapierCraft could support features usually found on dynamic media such as command name discovery and easy error recovery (UIST '06). We showed empirically that this reduces the error rates for commands issued on paper to a level similar to those observed for dynamic media (ToCHI'12). Principles developed in PapierCraft were adapted for several additional systems including **ButterflyNet,** a field note management system for biologists (in collaboration with Scott Klemmer at Stanford, CHI'06**),** and a paper interface for Classroom Presenter, an active learning system (in collaboration with Richard Anderson at UW, INTERACT'07).

In collaboration with AutoDesk, we also explored how the original PapierCraft system could be extended to large format prints used by building contractors in the field. Because of the larger format, the standard pen-top interface with its limited feedback options becomes very tedious to use in such settings. We addressed this limitation by augmenting interactions with a projector located either in the barrel of the pen (**PenLight**, CHI'09) or on a separate mouse-like device (**MouseLight**, CHI'10 - nominated for best paper award). Mouselight implemented the first version of the bimanual 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.

Moving beyond linear reading for ebook-readers

In the mid to late 2000s, the advent of bi-stable display technologies such as the E-ink system 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. Such pre-iPad devices were convincing enough for

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

substantial numbers of readers to initiate a transition from paper to digital books. However, studies showed that early E-ink devices remained inadequate for more active forms of reading, because they did 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** (ToCHI'12 - Best of 2012, Computing Reviews)**,** a new system to support active reading on ebook readers. Key to the United Slates design is the insight that the success of paper lies in the ease of annotation and the availability of multiple display surfaces at any given time. In collaboration with an industrial partner, we designed and built 60+ prototypes of a new ebook reader system extending our early work on dual-display systems (CHI'08) toward a multi-slates system. On the hardware side, the system was lighter and thinner than the first generation of iPads, yet it included 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 could easily transfer documents from one slate to the other, use several slates in conjunction to simplify navigation through multiple documents, or link two slates to facilitate the creation and review of annotations. To simplify content creation, a laptop could join the federation of slates and call slate content directly on its screen. A deployment of this system demonstrated the importance of having several separate display surfaces to recreate the affordances of paper (CHI'13).

A richer annotation system for collaborative review.

In many workplaces, collaborators iteratively review and exchange feedback as a core practice to improve the quality of written documents. Effective communication of edits, questions, and comments is central to the evolution of a document and plays a key role in maintaining group dynamics. Working face-to-face (F2F) has many advantages for these kinds of collaborative processes. F2F collaborators enjoy a shared context in which they can verbally explain details and gesture over documents, often taking notes as they do so. Different modes of communication are interleaved seamlessly and often support each other. Unfortunately, many of these channels are not transcribed using digital annotation systems such as ink markup, textual annotations, and emails. To address this problem, we developed **RichReview** (UIST'14), a document annotation system that brings the richness and expressivity of F2F discussion to asynchronous collaboration scenarios. The system is based on a novel interface allowing users to freely interleave and combine freeform inking (best for small edits), voice (best for high level comments), and deictic gestures (to support speech) as they review a document. To avoid crowded margins, comments are inserted within the flow of the text using the **TextTearing** technique (UIST'13). When RichReview replays annotations, it uses time-synchronization between voice, ink, and gesture streams to let users quickly reach a point of interest in any of these modalities and replay them in a synchronized manner. Because the system balances ease of production and consumption, it makes it possible to create rich 'conversations' around the document, despite the fact that the discussing parties are neither collocated nor working on the document at the same time. This system proved very efficient in delivering instructors' feedback on term papers (CSCW'16). In fact, the system was so well received by students, that they told us they preferred this form of delivery to standard office hours. We further extended the RichReview system by developing **TypeTalker** (CSCW'17), a new interface to quickly edit voice transcripts and diminish speakers' self-consciousness about leaving voice comments. In TypeTalker voice is always transcribed to text which can be quickly corrected or revised using standard textual methods, before being translated back to a synthetic voice which provides anonymity to the speaker. The system maintains synchronization between voice and gestures and preserves speech timing to create a more natural delivery. In combination, the Richreview system and its subsequent extensions were able to leverage the latest technological advances (slim, compact pen-based tablets and advanced, cloud-based speech recognition) to deliver a new annotation experience.

A MORE REFLECTIVE APPROACH TO 3D MODEL DESIGN

Over the past decade, we have seen a rapid push towards the democratization of 3D printing. People have drawn a parallel between this situation and the rise of desktop publishing in the 1980's: The Apple Macintosh and its companion laser printer let users produce near print-shop quality documents, often at a fraction of the cost. Yet, it is proving far more difficult to create complex 3D models that can leverage the full potential of the new 3D printers. No matter which program is used to create a digital model, the common 3D printing work cycle (designing on a computer, printing, iterating) requires a high level of skill for all but the simplest models and is hampered by the isolation of the design from the real object due to the slow production process. This stands in sharp contrast with traditional craft activities, such as clay coiling, in which design and building can occur at the same time. The intimate interaction between the designer and the material at hand establishes a constant reflective feedback loop promoting faster convergence towards a satisfactory design as described by Schön in his book "The Reflective Practitioner".

Capturing Annotating on 3D models

As 3D printing became widely available, designers started to annotate their newly printed models with ink annotations reflecting the required changes. Yet, as in the case of paper documents, it proved very difficult to seamlessly transfer this information back into the original digital model. With **ModelCraft** (in collaboration with Hod Lipson at Cornell, UIST'06) we demonstrated that the interaction techniques developed for paper documents were readily applicable to this nascent field. The ModelCraft system allows users to transfer annotations and edits (such as cuts or extrusions) that are made on 3D models using digital pens directly into the original 3D model. 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. Like the PADD system (and many other systems we designed in that space) ModelCraft empowers the users pick the medium most appropriate to at the task at hand.

Interactive 3D printing

More recently, my group developed of series of prototypes exploring how the design process can ultimately be merged with the 3D printing process. Our first system, **D-Coil** (CHI'15), used a handheld, computer-controlled wax gun to assist users in creating accurate geometry. D-Coil allowed users with no CAD experience to create complex digital models using a coiled wax model as the primary feedback. D-Coil fostered a constant dialog between the user, her tool, and the wax proxy of the digital model. However, because the printing element was handheld, it remained fairly slow.

To address this problem, we set out to speed up the creation of low-fidelity models using fused deposition modeling, the standard for low-cost 3D printing. First, we developed **WirePrint** (in collaboration with Patrick Baudish and Stefanie Mueller at HPI, UIST'14), which accelerates the printing process by printing a simplified 'skeleton' of the model. Because the printing is laid out as plastic wire (as opposed to layer by layer), it increases printing speed up to 10 times. It is also relatively simple to implement on off-the-shelf printers as exemplified by a recent instantiation of WirePrint on Cura, an open-source 3D printing software. We further extended this approach by developing **On the Fly Print** (CHI'16), a system which prints a physical 3D model concurrently with the design of the digital model. To achieve this goal, we further sped up the WirePrint process via water-mist cooling and added a two degrees of freedom (2DOF) printing platform to allow for incremental printing. Using On the Fly Print, users can touch and handle the model at any time during the design process, fostering a reflective interaction with the shape being designed. The addition of a 2DOF platform to a standard 3D printer proved very valuable in other areas such as patching 3D prints (with Patrick Baudish's group at HPI, UIST'15), printing arbitrary meshes (with Steve Marschner's group at Cornell, SigGraph'16) and printing working electromagnetic devices (with Scott Hudson at Disney Research, UIST'16).

We are in the process of further extending On the Fly Print by using a human-friendly robot arm to perform printing inside a shared design space. In essence, this approach will translate the benefits of iterative 2D sketching (mentioned by Schön) into the 3D realm.