(Digitally) Inking in the 21st Century

Michelle Annett
Autodesk Research and the University of Toronto

From tablets and smart watches, to laptops and cell phones, the ubiquity and mobility of contemporary computing devices has enabled us to consume and create content, anytime, anywhere. On mobile devices, a few swipes, taps, and pinches allows us to consume content at a moment's notice. When we need to create content, however, touch input is far from perfect. Composing an email, tweet, or snap is perfectly doable via touch input, but if we want to sketch a new product design, formulate mathematical proofs, or annotate a manuscript, touch input is ill equipped to support such activities. In these situations, the digital stylus, or pen, can come to the rescue.

First suggested by T.L. Dimond in 1957 as a method for toll switchboard operators to produce toll tickets quickly,[1] the stylus has been envisioned as the ideal computer peripheral. Because it shares in the design aesthetics of traditional pens and pencils, users can transfer a lifetime of learned and practiced behaviors into the digital world to reap the benefits of undo and redo, handwriting recognition, and stroke beautification. Unlike the mouse and keyboard, a stylus enables users to diagram, sketch, and annotate directly on the display screen. This allows users to ink with more fluidity and naturalness than with a mouse or keyboard, which utilize indirect interaction for input. When coupled with touch input, the stylus should enable users to simultaneously ink, manipulate the page, and switch between tools with ease.

Although the benefits of stylus input and the recent fanfare surrounding the Apple Pencil, Wacom Bamboo Smart products, and the Surface Pen may make it seem like the stylus has "arrived," it has yet to supplant traditional pen and paper or mouse and keyboard input. A quick scan around an elementary school, college campus, or industry boardroom will reveal that users work on tablets and smartphones, but few use a stylus with such devices. With more than 50 years of research and technological advancements behind it, and given its benefits over the mouse, keyboard, and touch input, why has the stylus yet to achieve universal adoption?

My thesis sought to understand the usability barriers and tensions that have prevented stylus input from gaining traction and reaching widespread adoption. It critically examined users' natural behaviors while they ink, and it identified the fundamental issues that plague today's stylus-enabled systems. The results of an initial exploratory study[2] uncovered the impact that device latency, unintended touch, and accuracy have on the stylus experience today; it also highlighted the importance of stylus design and stroke beautification to the digital inking experience. This article explores the systematic, empirical approach that I took to uncover the limits of human latency perception and gathers observations from natural inking behaviors to evaluate solutions to unintended touch.

Editor's Note

Michelle Annett received the 2014 Bill Buxton Best Canadian HCI (Human-Computer Interaction) Dissertation Award.

Initial Exploratory Study

Although pen computing has a long history, many researchers have focused on novel bimanual interaction techniques or exploited functionality such as pressure, tilt, azimuth, and hover.[3] There is, unfortunately, little empirical evidence about the problems users face while inking (writing and drawing) or the behavioral accommodations they make to overcome poor device and peripheral design. To bridge this gap, I conducted a user study to identify the behavioral, performance, and preferential differences between analog and digital technologies. To understand how various tablet and stylus properties affect user behavior, I asked 30 participants to transcribe mathematical

0272-1716/17/$33.00 © 2017 IEEE

content and sketch organic shapes using pen and paper (analog), an iPad with a (passive) capacitive stylus, and a Samsung Slate with an (active) stylus that included pressure input and an eraser.

The experiment revealed that participants wrote smallest on paper, slightly larger with the active stylus, and largest with the passive stylus. Many participants believed that the iPad was "incapable of detecting any strokes smaller than one-quarter inch so [they] had to write and draw much larger than normal." With the Slate, the nib's precision and the feedback provided about the nib's location before it touched the screen let participants write at sizes closer to that of paper. Still, paper was the most preferred medium, followed by the Slate and then the iPad.

Because paper has zero latency and a natural feel and texture, it provided direct contact with no parallax or delay and was easy for the palm to glide across. Although active stylus devices are optimized for inking, the Slate was rated lower than paper, but it was rated higher than the iPad. Participant comments included "the Slate was much faster than the iPad, but it was of course still slower than paper" and "the Slate didn't have the palm 'touchy' problems that the iPad did." Unsurprisingly, the iPad received poorer ratings than the Slate and paper, mainly because of the imprecise nature of the iPad stylus and its lack of pencil-like features. For example, participants said, "I couldn't see where I was writing because of the [passive iPad] squishy pen so I had to write bigger" and "without pressure sensitivity, the strokes looked awful."

The experiment also revealed that participants made a number of movement pattern modifications while using the digital devices. I identified novel hand-movement patterns, such as floating (the wrist, palm, and/or fingers were held aloft, above the writing surface) and dragging (the hand was placed on the media and dragged across the surface until it reached the end of the stroke being made, and then it was picked up and moved it to the next location). Participants were more likely to use a dragging movement on paper and the floating behavior on the digital devices. On paper, participants could slide their hands along the page because the friction between their hands and the surface was suitable, whereas with the iPad and Slate, it was too high so participants lifted their palms.

Overall, I found that many elements influenced behavior, performance, and preferences while inking (see Figure 1). Participants were most vocal about three primary features (stylus accuracy, device latency, and unintended touch) and two secondary features (stylus and device aesthetics and stroke beautification). The identification of these primary and secondary features, as well as the differences noted between the tablet devices, warranted further work and constituted the remainder of my thesis. This article explores device latency and unintended touch.

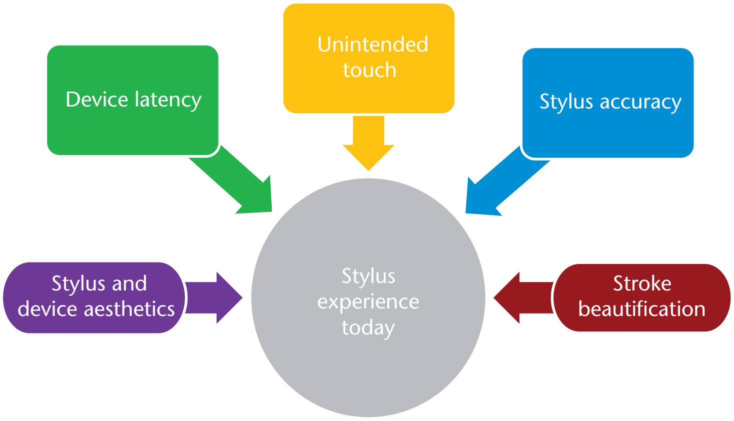

Figure 1. The five factors identified as most impacting the stylus experience. The major factors (device latency, unintended touch, and stylus accuracy) were explored within my thesis. This article explores device latency and unintended touch.

Latency Perception

One of the biggest complaints during our initial study was that participants were frustrated with the time it took for their ink to show up on the screen, especially when compared with the near perfect experience that of using pen and paper. With pen and paper, the delay (or latency) is essentially zero because ink instantly flows from the nib onto the page. On a digital device, however, data travels through a complex pipeline[4] that involves sampling input from the digitizer, filtering and reporting events to the operating system, event processing by the operating system, reporting and processing events by an application, and updating the display to provide feedback to the user. Each step in this pipeline contributes to the sluggish experience that many users encounter.

Because it is unlikely that stylus-enabled devices will ever achieve zero latency and we know little about the perception of latency, I set out to answer a simple question: What is the minimum latency that users can actually perceive? If we know the minimum perceivable latency, then we can work toward decreasing latency across the pipeline to achieve this target.

To determine the minimum perceivable latency and understand the factors that influence latency perception, I conducted a series of experiments.[5,6] Because commercial devices cannot achieve latencies below 50 milliseconds (ms), we developed a prototype high-performance stylus system (HPSS)

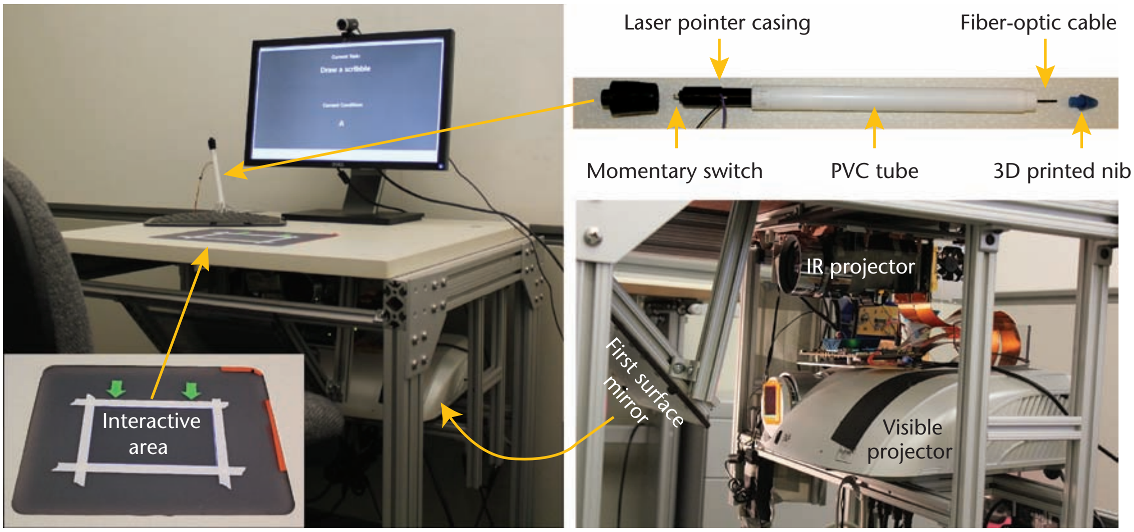

Figure 2. High-performance stylus system. The HPSS prototype consists of two high-speed projectors, a first surface mirror for rear projection, and a fiber-optic based stylus.

(see Figure 2).[6] The HPSS uses two Texas Instrument Discovery 4100 high-speed projector kits and a first-surface mirror for rear-projection onto a diffuse surface. The first projector rear-projects a series of grey-coded patterns to uniquely encode every pixel in the interaction area at 17,000 frames per second. A second projector refreshes at 23,000 binary frames per second and provides visual feedback to the user. For input, we designed a stylus that uses a 1-mm fiber optic cable mated to a photodiode to detect the grey-coded IR patterns. With this system, a minimum latency of 1.4 ms was possible when a square was displayed on the screen and 7 ms while ink was displayed (due to the filtering required to create smooth ink strokes).

Using an experimental paradigm from the psychophysics literature, which is similar to the experience one has while at the optometrist, we were able to measure the minimum latency that users could perceive. While administering an eye exam, optometrists will show you one possible prescription, quickly switch to another, and ask you which is better. Then, they will show you the option you just preferred, another slightly different one, and ask for your preference again. This process repeats itself until your responses converge toward your new prescription. A procedure similar to this, called the *Just Noticeable Difference Paradigm*,[7] was used with the HPSS to measure latency. For each trial, participants performed a task using a baseline latency—that is, 7 ms while inking and 1 ms while dragging a square. They then performed the task again with a test latency and indicated which of the two conditions had the least delay. An underlying algorithm attended to participants' response history to compute the next test latency and determine when participants had converged at their minimum perceivable latency.

The first experiment[6] asked users to perform two tasks: drag a square across the screen from left to right and back, and draw an oscillating line from the upper left corner of the screen down toward the lower right corner. During such tasks, many participants achieved a minimum perceivable latency of 2 ms while dragging the square, but because the system was unable to create smaller latencies, participants reached our floor of 1.4 ms. Even with these technical limitations, participants were able to distinguish between 1 and 2 ms of delay. For the line drawing task, however, the minimum latency participants could perceive was 40 ms—that is, they could not distinguish between the minimum latency of 7 ms and 40 ms. Because of the different baseline latencies, these results could not be compared, but they hint that the perception of delay could be influenced by the task being performed, the feedback users receive, what aspect of the feedback they attend to, and so on.

I conducted a second experiment[5] to understand if different tasks (such as drawing vertical lines, writing the word "party," and sketching a six-sided star) influenced latency perception. The results demonstrated that task did not influence perception (the minimum perceived latency across all tasks averaged 55 ms). I hypothesized that because there were many more possible points of reference to use when making latency judgments (compared with the oscillating line from the first experiment), participants attended to different visual cues. Thus, strategy and location may be involved in the detection of latency.

Inspired by the strategies used in this second experiment, a third experiment[5] manipulated the location of the digital ink, forcing attention away from the stylus nib toward other areas of the screen: under the nib and offset to the left of the nib. If participants focus on the relationship between the nib and ink to make latency judgments, offsetting the ink's location should decrease latency perception. The results showed that ink offset did not have a significant effect on participant's ability to perceive latency. A minimum latency of approximately 59 ms was found when no offset was present, approximately 50 ms with a small (6 mm) offset, and approximately 59 ms with a large (65 mm) offset. Because there was no significant difference between the offset conditions, participants likely did not use the distance between the nib and ink to make their judgments.

To understand how the motor system may influence the visual system during latency perception, I also included a condition where the participant's hand was obscured from view. The results indicated that participants were able to better perceive latency when they could view the stylus-wielding hand and stylus (approximately 59 ms) compared with not receiving this visual feedback (approximately 97 ms). The presence or absence of the hand and stylus thus appears to be an important referent.

Taken together, these experiments show that users can perceive substantially smaller latencies than those currently possible with systems today, and they highlight that many elements influence the perception of latency. (See the "Latency Perception Model" sidebar for more details.) If stylus-enabled systems could achieve latencies within the range of 40 to 50 ms, most users would not notice a delay while inking. This is likely achievable within the next few years, if spare milliseconds can be found across the whole data pipeline.

When it comes to tasks such as dragging a square, the fact that users could distinguish between 1 and 2 ms is worrisome. The HPSS allowed users to experience extremely low-latency stylus interaction. Bringing such experiences to future commercial systems will be all but impossible in the near future. Although it may be possible to find a few spare milliseconds here and there, freeing up 50, or even 30, milliseconds will require substantial development and innovation across all sensing, processing, and display subsystems.

Opportunities in Latency Research

Since my dissertation research was published, there has been a latency renaissance. Many companies have devoted great effort to reducing hardware and software latency, or at least to making consumers aware that latency differentiates their products. Within the last year, the marketing materials for products such as Apple's iPad and Microsoft's Surface Pro have explicitly touted the increased responsiveness of their systems and the fluid and natural experiences that their systems offer. Although these devices have not yet reached latency levels that users will not be able to perceive, it is encouraging that industry has acknowledged the importance of latency and devoted resources to improving the stylus experience for users.

Until latencies decrease to the point where users do not notice them, there are a few avenues to consider. Many applications have begun to use stroke prediction techniques to prerender strokes before they occur and then subtly adapt the prerendered strokes after the stylus has reached or passed a given point. This tricks users into believing that a device is more responsive that it actually is. Because the stylus's location reveals much about the intentionality of current and future interactions, application context and knowledge about hand postures or grip could also be useful to sampling subregions of the input sensor and redrawing targeted subregions of the display.

From a theoretical perspective, much is still unknown about latency. The latency perception model that I proposed provides a solid basis and potential blueprint for future work. (See the "Latency Perception Model" sidebar for more details.) By isolating each factor and examining its effects, we can better understand the roles of referents and responses. Is there a constant cost for having the referent and response in different modalities? Does each modality have its own cost when judging latency? The experiments I conducted provided some insight into these questions, but there is still much more to do.

Although my research uncovered the perceptual limits of latency, there is a difference between what users can perceive and what they will tolerate. It is likely that latency tolerance is higher than latency perception, as users outside laboratory environments will encounter environmental and task distractions. How much higher a user's latency tolerance is remains unknown and a fruitful area of research.

Unintended Touch

While inking, it is common to rest the palm, wrist, or forearm on the screen or page. When performed on a digital device, this behavior can unintentionally manipulate content or render unwanted ink. Although it may seem as if the solution is to filter all touch input while inking, there are many situations where the user actually wants some touch input to be detected and acted upon (for example, a swipe to the next page or a zoom in on an image). These competing desires present a dilemma: How do we differentiate between intended and unintended input or from input that is a byproduct of the skin resting on or grazing a display?

Within industry and academia, this problem is known as *palm rejection*.[8] But this term reduces the

Latency Perception Model

As my experiments demonstrated, the perception of latency is a multifaceted problem. In the third experiment,[1] the input action (task) and output response (feedback) were controlled, but the perception of latency changed. With the smaller offsets between the nib and line, participants could use visual information to make judgments, but once they could not see their hands, they had to use other information streams, perhaps auditory or tactile cues from the stylus that affected latency perception. Based on present work and prior literature, I developed a model that describes the processes underlying latency perception in stylus and touch interaction (see Figure A).

The input action can take many forms (such as touching a finger to the screen or pressing the stylus against the screen) and is invoked by the observer, another user, or an external system or device. The latency source handles the action once it occurs (for example, a sensor array, operating system, or application). This entity converts the input action into output responses and adds delay. These responses are most often visual (such as a dot, line, or simple shape), but they could manifest themselves via other modalities as well.

In addition to supplying the latency source with input, the input action also generates stimuli, or referents, that provide clues to the observer (such as a stylus barrel, fingernail, hand, or stylus nib). In prior work with tapping, the haptic sensation of the finger pad touching and moving along the surface produced an additional referent that assisted in the perception of latency.[2,3] Because the stylus naturally dampens the haptic sensations from the screen, it is likely that haptic referents play less of a role in stylus-based scenarios. Similar to output responses, there is spatial and temporal uncertainty about what influences the referents. The referent may also need not be a physical stimulus, but it may take the form of a cognitive initiation of an action (a mental "go" signal).

Once the referents or responses are available, the observer compares the original input action to them to determine the magnitude of latency generated by the latency source. The experimentation showed that during this comparison there are many contextual demands influencing the observer: judgment strategy, which plays a role in focusing or diverting attention from referents and responses; the location and amount of attention; and external environmental distractions. Tasks that require more attention and increased cognitive load (that is, inking versus dragging a square) seem to redirect resources away from perceiving the referents and responses, making latency judgments more difficult.

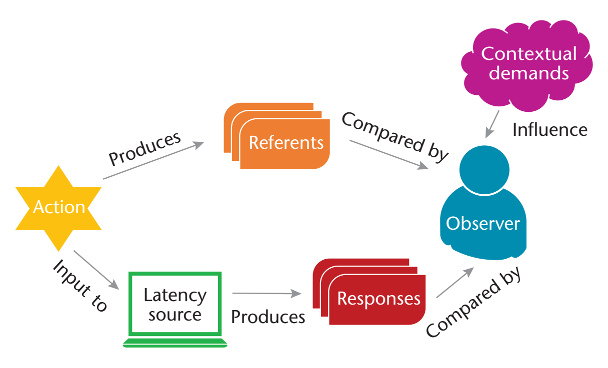

Figure A. The latency perception model detailing the role of the input action, the resulting referent stimuli, the latency source, and the output responses. The observer compares the referents and responses when perceiving latency and is likely influenced by contextual demands such as task requirements, loci of attention, environmental factors, and so on.

References

1. M. Annett et al., "How Low Should We Go? Understanding the Perception of Latency While Inking," *Proc. Graphics Interface*, 2014, pp. 167–174.
2. R. Jota et al., "How Fast Is Fast Enough? A Study of the Effects of Latency in Direct-Touch Pointing Tasks," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2013, pp. 2291–2300.
3. A. Ng et al., "Designing for Low-Latency Direct-Touch Input," *Proc. ACM Symp. User Interfaces and Software Technologies*, 2012, pp. 453–462.

challenge to identifying and rejecting the palm. My dissertation argued that *unintended touch* more appropriately describes this problem. On capacitive devices, a touch event is generated whenever the knuckles, wrist, palm, fingers, forearm, and so forth contact the screen (see Figure 3). Because these devices do not have dedicated stylus digitizers and their styli mimic the finger, stylus input is undistinguishable from finger-generated input. This means that stylus input needs to be disambiguated from touch input before it could be labeled as either intended or unintended. With devices that contain dedicated stylus digitizers (such as the Wacom Cintiq, Samsung Business Slates, and Surface Pro), distinguishing between the stylus and touch is not an issue, but disambiguating between unintended and intended touch remains. The user's skin can touch the screen at any time and in any location, so determining the intentionality of touch is still difficult.

To better understand skin input with stylus- and touch-enabled devices, my research team conducted an experiment with 18 participants.[9] Using a six-camera Optitrack motion capture system, three web cameras, a Sony Viao tablet (with NTrig digitizer), and an Ntrig stylus, we collected raw, unfiltered antenna magnitudes and motion-capture data from the stylus and tablet while participants were inking (see Figure 4). This dataset allowed me to identify how various hand postures and behaviors influence approaches to unintended touch.

A common suggestion to address unintended touch is to reject all touch input that is larger than a finger. I thus looked at the temporal growth of each touch event in the collected sample to determine when intentional touch was distinguishable from unintentional touch. The digitizer data demonstrated that touch input initially activated a very small number of sensors and did not grow to a differentiable size until approximately 33 ms after its initial activation (see Figure 5). Waiting for touch input to merge and stabilize before making a rejection decision is not feasible, especially when current latencies are easily perceptible by users.

Another common idea is to compare the shapes of each touch event and reject all touch input that does not conform to an oval or palm-shaped region. The digitizer data again demonstrated why such an approach was not sufficient. The shape of each touch event did not have a circular or oval form as one would expect; instead, it appeared as a straight line or irregular, jagged shape (see Figure 3). This evidence thus debunked two of the most common suggestions to overcome unintended touch.

In addition to these behavioral observations, my exploration also analyzed current and novel techniques to overcome unintended touch (such as rejecting touch input in fixed locations on the screen, creating a model of the wrist and forearm, rejecting all input inside this model, and using the stylus to dictate the regions where touch input is rejected). The details of each approach are outside the scope of this article, but the analysis generally showed that techniques that used larger areas of the screen where users could safely rest their hand outperformed techniques with smaller areas to rest the hand. This is because the larger areas better accommodated the variety of hand postures that users can exhibit. In addition, approaches that made use of the stylus location to dictate the bounds of such areas performed better than those that used generalized, restrictive data sources (such as the size or shape of the input) because the stylus location is a reliable indicator of where the fingers, palm, wrist, and forearm may touch the screen.

These results suggest that contextual information is crucial to overcoming unintended touch. If we consider a sketching application, where a user can zoom in and out via a pinching gesture, successful unintended touch techniques would need to be cognizant of the locations where the zooming could occur, the task-specific behaviors that may result (such as reorienting the screen), and the cues from the stylus that would provide in-

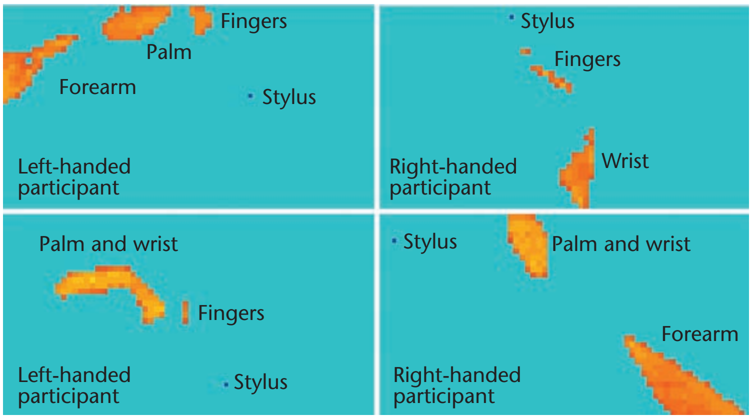

Figure 3. Example data provided by stylus and touch digitizers. The blue dot indicates the stylus location, orange represents the touch data, with the darker shades representing the higher levels of activation. All touch data in these images is unintended and should not generate touch events.

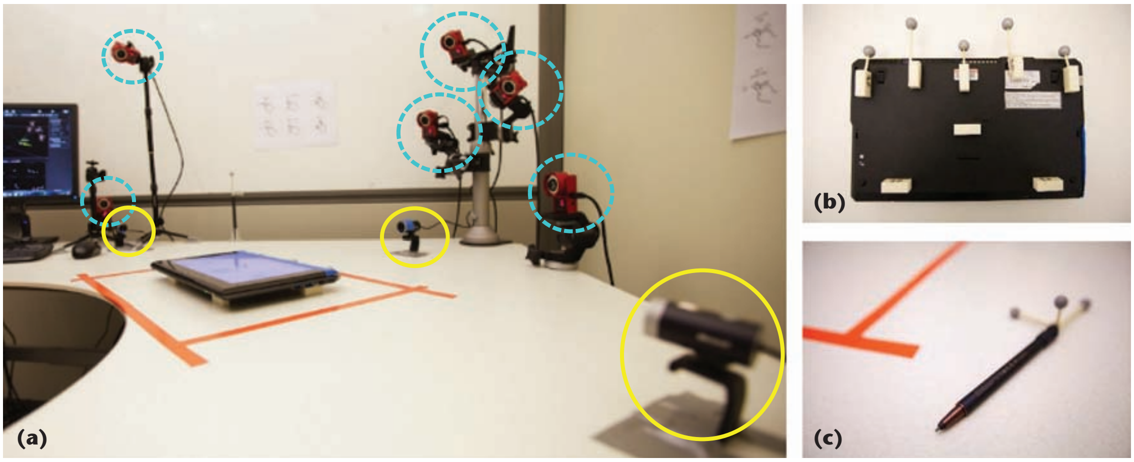

Figure 4. Experimental setup with a motion-capture system: (a) six-camera Optitrack system, (b) a Sony Viao tablet (with NTrig digitizer), and (c) an Ntrig stylus. The dashed blue circles identify the Optitrack cameras, and the solid yellow circles show three web cameras.

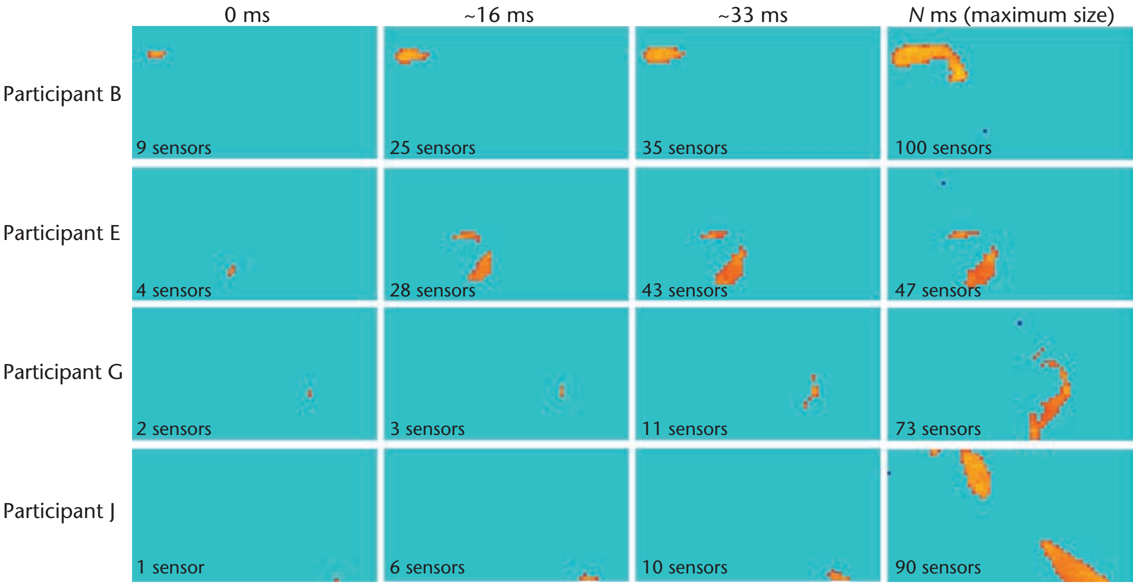

Figure 5. Examples of different touch events that initially activated a small number of sensors (denoted by shades of orange) and later grew to their full size. Note that the initial shape and direction of the touch input does not represent the eventual shape and orientation of the touch input.

formation about a user's intentions (for example, hover-based information can indicate where the hand will likely be placed on the screen). Although my study did not explicitly evaluate such a complex sketching scenario, it is easy to see how valuable the observation of natural behaviors can be to the design of unintended touch solutions.

Future Research in Unintended Touch

As my dissertation demonstrated, many elements contribute to unintended touch, making the problem, and the choice of an unintended touch solution, complex. Over the last two years, there has been an increase in the number of applications focused on unintended touch. A recent Windows update, for example, asked users to indicate which hand they write with, and it provides them with an option to turn touch input off while the stylus is in use. OneNote, SNote, and the Moleskine iPad app have similar options. Although turning touch input off is a bit of an admission that a solution to unintended touch has yet to be found, soliciting such feedback from users indicates that industry is at least aware that unintended touch is a problem.

So how do we actually solve unintended touch? First, I believe many other streams of data need to be harnessed. The pre-touch information that the stylus provides when in the hover state is more valuable than currently thought. If rejection processing and decisions can begin before the hand or stylus touches the screen, there will be more time to make informed rejection decisions. Mobile phones from Sony and Samsung already ship with support for such pre-touch capabilities (such as touch hover) for in-air gesturing. Integrating additional sensor information, such as a stylus roll or tilt,[3] or utilizing the grip on the barrel of the stylus[10] could also be useful. Developers would need to be mindful when integrating such information, as the use of additional sensors would increase monetary costs and latency and would require more sensor-fusion techniques.

Second, modifications to firmware or operating system APIs could be fruitful as well. If each touch input could provide developers with a "rejection confidence level," the processing to make rejection decisions could be distributed along the pipeline. Those applications concerned with touch and stylus input could perform further application-specific computations based on such a metric. Applications could indicate how intended they think touch events are by changing the opacity of strokes or using colored overlays. This could also raise awareness and prevent behaviors that cause stray markings or unintended navigation. The computation of a rejection confidence level could come from sensor magnitudes, stylus hover distances, device orientation, or any of the aforementioned additional data streams.

Lastly, until unintended touch is solved, it is imperative that the algorithms and metaphors that developers use are discoverable and easy to understand. On many occasions, users were quick to identify when an approach did not work for them, but they had difficulties determining why or how to adapt their behavior. This may lead many users to become frustrated with a device or application and feel that they, not the device, is inadequate. This should not be the way we design devices or solutions to support natural behavior.

The Future of Digital Ink

Although my dissertation has made many contributions to our understanding of stylus interaction, there are still many unanswered questions that remain ripe for exploration. My initial exploratory study identified the importance of device aesthetics and stroke beautification, but little attention is often given to these increasingly important factors. As the variety of implements compatible with tablets increases, understanding how users expect strokes to look will be crucial to providing a great user experience—for example, should a crayon-like stylus be able to draw brush and pencil strokes? In addition, with traditional pencils, repeated use leads to wear and subsequently different markings left by the lead. Determining if this natural decaying process is useful and how it could be integrated within the digital world will greatly increase the relationship users have with their content and the depth of content that is created.

Furthermore, questions also remain about the relationship between a device's texture, the nib's texture, and the user's haptic experience. With advances in material science and novel fabrication techniques, there will be a plethora of materials available to construct new stylus nibs and surface textures. As users continue to desire and expect digital experiences that mimic the physical world, it will be important to understand how user experience and performance change when displays or nibs are able to dynamically change texture, in addition to how such advancements will change the precision and accuracy of user input.

One area of research that my dissertation did not touch on is the relationship between the various factors. The systematic evaluation that was undertaken revealed a great deal about each factor, but it did not consider how the factors work in concert. For example, how is unintended touch improved when latency is decreased? And how does decreased latency influence accuracy? From a theoretical perspective, there is also still much to learn about the underlying psychology processes that govern interaction and how visual feedback and the motor system influence one's overall perception of a stylus-enabled device.

Outside academic research, there is still much to be done within the stylus ecosystem to create cohesive user experiences. As a result of competing corporate interests, there will always be diversity in the marketplace, with less expensive, limited styli-enabled devices and higher-end, fully functional devices. Such diversity will require designers to think carefully about the transition between styli with differing functionality, the utility of styli that use alternative input channels to provide pressure and tilt data, and stylus-enabled phone and smartwatch experiences. Only time will tell when we will be able to achieve stylus-based experiences that meet, or exceed, those we have come to love with pen and paper.

Acknowledgments

I thank all of my coauthors and colleagues at the University of Alberta and Microsoft Research who made my work possible. Special thanks to Walter F. Bischof for his guidance and advice throughout my master's and PhD programs. My dissertation research was supported by scholarships from Alberta Innovates Technology Futures, Google, the Government of Alberta, and the University of Alberta.

References

1. T.L. Dimond, "Devices for Reading Handwritten Characters," *Proc. Eastern Joint Computer Conf.*, 1957, pp. 232–237.
2. M. Annett et al., "The Pen Is Mightier: Understanding Stylus Behavior While Inking on Tablets," *Proc. Graphics Interface*, 2014, pp. 193–200.
3. Y. Xin, X. Bin, and X. Ren, "Natural User Profiles for the Pen: An Empirical Exploration of Pressure, Tilt, & Azimuth," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2012, pp. 801–804.
4. K. Hinckley and D. Wigdor, "Input Technologies and Techniques," *The Human Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, 3rd ed., J. Jacko, ed., CRC Press, 2012, pp. 99–132.
5. M. Annett et al., "How Low Should We Go? Understanding the Perception of Latency While Inking," *Proc. Graphics Interface*, 2014, pp. 167–174.
6. A. Ng et al., "In the Blink of an Eye: Investigating Latency Perception during Stylus Interaction," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2014, pp. 1103–1112.
7. M.M. Taylor and C.D. Creelman, "PEST: Efficient Estimates on Probability Functions," *J. Acoustical Soc. of Am.*, vol. 41, 1967, pp. 782–787.
8. R. Zeleznik et al., "Applications and Issues in Pen-Centric Computing," *IEEE MultiMedia*, vol. 15, no. 4, 2012, pp. 14–21.
9. M. Annett, A. Gupta, and W.F. Bischof, "Exploring and Understanding Unintended Touch during Direct Pen Interaction," *Trans. Computer Human Interaction*, vol. 21, no. 5, 2014, article no. 28.
10. H. Song et al., "Grips and Gestures on a Multi-touch Pen," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2011, pp. 1323–1332.

Michelle Annett *is an NSERC postdoctoral fellow at the University of Toronto and Autodesk Research. Contact her at mkannett@gmail.com.*

Contact department editor Jim Foley at foley@cc.gatech.edu.