Inking Outside the Box: How Context Sensing Affords More Natural Pen (and Touch) Computing

Ken Hinckley✉ and Bill Buxton

Abstract The authors were invited to present a reprise of a recently-published paper on *Sensing Techniques for Tablet + Stylus Interaction* at the WIPTTE 2014 Workshop. The talk took the original contribution as a point of departure, because for the WIPTTE venue we felt that the most important role of the work was to illuminate and help the audience understand more deeply the interaction modalities of pen and touch – as well as their use in tandem. And in the process the authors felt like they came to understood the topic more deeply as well, hence the paper that follows.

One of the premises of the talk was that even a concept as seemingly straightforward as 'touch' – not to mention pen + touch, used together in complementary roles – is perhaps not as well understood as we might think it is.

In particular, we argue that beyond the standard idiom of touch (and multi-touch) interaction on touchscreens, there are many aspects of 'touch' that are rarely considered (much less actually sensed) by existing devices and interaction designs. We show how this surrounding context of manual activity – how the user is holding the tablet, how the user is gripping the pen, and how each device is oriented and moving relative to the other – have the potential to considerably enrich interaction with tablets, and thereby to re-define what we conceive of as 'natural' interaction with pen and touch.

1 Introduction

Pens seem like such a relic of the twentieth century, and even more antiquated epochs long before that. Isn't it 2015? Why are we still talking about pens?

Well, by way of simple illustration I thought it would be fun to visit the supply room of one of the largest technology companies on the planet. Surely one can go in there and find all manner of technological contrivances. Since it's just down the hall from my office it was easy to snap a photo. Surely it should stacked to the ceiling with fancy robot arms and 3D printers, or perhaps bins chock full of the legendary flux capacitor.

Sadly, this was not to be the case:

Ken Hinckley✉
Microsoft Research, One Microsoft Way, Redmond, Washington, USA, e-mail: kenh@microsoft.com

Bill Buxton
Microsoft Research, Redmond, Washington, USA, e-mail: bibuxton@microsoft.com

Fig. 1 The supply room of one of the largest technology corporations on the planet. Contrary to what one might expect, close inspection reveals the contents are quite pedestrian—and well stocked with writing instruments of every description. *Photo © Ken Hinckley, 2014. Used by permission. All rights reserved.*

But it does reveal something about what actually goes on when knowledge workers retreat to their cubicles and try to think outside the box.

The first thing they do is grab a pen!

Now, while the above illustration is a slightly tongue-in-cheek way to make the point, it is only just slightly that. Because the research bears this out.

The work of students and creative professionals typically involves heavy doses of reading. And this is not just ordinary recreational reading, but rather a very particular way that people have of digging deeply into documents and source texts. Such deep, purposeful engagement with content (often multiple pieces of content) is known in the literature as *active reading*.

What people are trying to do in active reading is to distill and crystallize knowledge from diverse sources.

The reader's tasks are typically complex, open-ended, ill-defined, and intellectually challenging.

The reader may annotate, mark-up pages, take notes, sketch out ideas and connections, or formulate summaries and responses based on what they've read.

And active reading, perhaps more than anything else, is characterized by reading side-by-side with writing. Such tasks involve reading in combination with writing and typically span multiple documents (or working surfaces) as well, such as the canonical yellow writing pad used to jot down notes while reviewing a manuscript.

A great entre to this literature is Sellen & Harper's classic book, The Myth of the Paperless Office [19], but many studies of this activity have been published [3,8,15,17,18,20].

One of the key viewpoints that has emerged in recent years is that although these types of working patterns carry over to electronic document work as well, there are clearly significant opportunities to improve on current practice.

Another emerging viewpoint is the recognition that, in contrast to the pen-only devices of the past, if we have a screen that supports *both a stylus* ***and*** *touch*, then we have something very powerful indeed.

Certainly much more expressive than a device with touch alone.

A device that can not only sense the human hand, but also our *pre-existing skill for dexterous manipulation of a mechanical intermediary*. A tool that has blazed a trail for many centuries with significant impact on science, education, and human intellect. A tool that lends itself to freeform expression and the creation of content that lives beyond the confines of any single data-type, ruled business form that one must fill in, or the strictures of a particular application.

The pen. The perfect tool for inking outside the box.

Tools can elevate human skill and we should not lose sight of this in our fervor to pursue 'natural' interfaces. Natural interaction does not require bare-handed interaction, and therefore a touchscreen, in and of itself, does not a natural user interface make.

2 Human Skilled Manipulation

Another way of looking at this is to back away from the technology and look again to the underlying human habits and behaviors.

But this time rather than considering work-practices in general, let's zoom in to the very particular details of the manual activities involved.

In fact let's completely eschew technology for the moment and recruit my six-year-old daughter to the effort:

Fig. 2 The first author's six-year-old daughter demonstrates the power of crayon and touch technology—through skilled use of both hands. *Photo © Ken Hinckley, 2014. Used by permission. All rights reserved.*

Now, of course she's using a crayon—a type of 'pen.'

But what we also see in this real-life example is that she's using 'touch' as well.

Not only is her preferred hand partially contacting and resting on the paper while holding the crayon.

But also, prior to coloring this particular area of the page, she has positioned and oriented the sheet of paper as well, using her nonpreferred hand to optimize its placement for the action of the preferred hand.

A number of important lessons for pen and touch technology can be drawn from this simple illustration:

1. The nonpreferred hand manipulates the workspace with 'touch.'

2. The preferred hand articulates strokes with the crayon, which has an effect that is distinct from the fingers on the page. This stands in stark contrast to existing practice in many tablet applications, even today, where either a finger or the digital pen can be used interchangeably to leave marks on the page. In our view this is neither natural nor an effective use of pen and touch as distinct interaction modalities with unique affordances.

3. While the nonpreferred hand is *intentionally* positioning and orienting the page, at the same time the preferred hand may be *unintentionally* contacting the work surface. Here, it has no effect, but in a charcoal sketch (for example) this might produce undesired smudges—or it might be employed intentionally by a skilled artist to soften and blend pencil lines previously left on the page. As such this hints that the very concept of 'palm rejection' on pen and touch displays is a misnomer—all touches are potentially valid, and can be put to good use by applications, so long as their nature can be understood and interpreted appropriately to the task and context.

Since one of the major themes of this paper is the ways in which close scrutiny of manual behavior—how people grip and manipulate objects such as tablets and pens—can inform interaction design, let's further unpack the first of the three statements above.

When we say *the nonpreferred hand manipulates the workspace with touch,* what form does this take? What properties characterize this manipulation?

To answer this question and reveal the *oft seen, but seldom noticed* role of the nonpreferred hand in everyday interaction, the first author undertook an experiment in his Ph.D. thesis [7] that was inspired by a classic study of handwriting conducted by Yves Guiard [4].

In my variation of Guiard's experiment, the task required subjects to draw a 'perfect' circle that passed through tic marks at 90-degree intervals.

One diligent participant in the study produced the hand-drafted circles shown below, on the left-hand side of Fig. 3.

What the same participant *actually* drew, however, was the jumbled mess on the right.

Which is also a perfect circle.

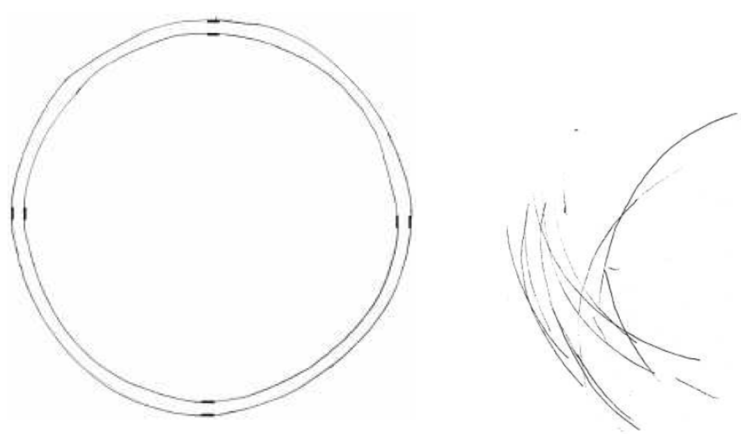

Fig. 3 When a circle is not a circle. **Left**: Figure drawn by a participant. **Right**: The strokes left behind on the desk blotter, which contained a hidden piece of carbon paper. The tightly clustered strokes demonstrate that the participant positioned and oriented the sheet of paper to make it easier, biomechanically, to draw the arcs comprising the full circle [7].

There's a trick here, of course, which depends on your frame of reference.

The circles on the left show the strokes *relative to the sheet of paper.* This is the finished work product.

By contrast the strokes shown on the right are the strokes *relative to the desk surface itself.*

These were captured by placing a blotter on the work surface. The pen strokes on the paper passed through and were also recorded on the blotter underneath, which surreptitiously included a hidden layer of carbon paper.

And as the user shifted and oriented the drawing, the resulting impressions reveal where on the desk the user actually drew the individual strokes comprising the circle.

What these crowded pen-stroke impressions reveal, then, is that the nonpreferred hand rotated the page to dynamically adjust the frame-of-reference to suit the action of the preferred hand: biomechanically it is far easier to draw smooth arcs in certain hand directions than in others.

This illustrates decisively that drawing a figure such as this circle—or as Guiard originally demonstrated, handwriting a page on dictation—is not, in fact a one-handed activity, but rather is the joint product of the activity both hands:

The nonpreferred hand positioning and orienting the underlying page;

Plus that of the preferred hand, which performs the actual micro-metric movements of the pen tip itself.

This simple illustration is ripe with lessons for interaction design in the context of pen and touch—and beyond.

It underscores that not only are there two hands, but that they play two distinct roles that are complementary.

And it further emphasizes that the same philosophy of thought can be applied to the *input modality* of the pen, as opposed to the input modality of touch, even though either hand can be used to make contact with the screen.

Guiard demonstrated that it is not helpful to ask, "Which hand is best, left or right?"

The correct question, rather, is one which frames the hands in a cooperative and complementary viewpoint:

What is the logic of the division of labor between the hands?

That is, which hand, left or right, should be used when assigning bimanual tasks to the preferred and nonpreferred hands?

Likewise, the question we should ask ourselves is not, "which input modality is best, pen or touch?" — but rather the following, which follows naturally from viewing pen and touch as cooperative and complementary input modalities:

What is the logic of the division of labor between pen and touch in interaction design?

3 Of 'touch' screens and pens: A tale of two modalities?

The question above has informed much of our own recent work exploring pen and touch interaction, such as our explorations of simultaneous pen and touch interaction [12,14], or our explorations of how pen and touch can support more informal ways of working with electronic documents [8].

But to really delve more deeply into the issues this question raises, we must turn a more critical eye towards two the interaction modalities that dwell at the heart of the WIPTTE moniker—the Workshop on the Impact of *Pen and Touch* Technology on Education—and challenge ourselves to think more deeply about what those two little words really mean.

Pen and touch.

Because they feel familiar and well-understood.

But neither of these terms are as well understood as you might think.

In particular, we want to raise some pointed questions about 'touch.' It's a term we often take for granted in the context of interaction with tablets.

A term rife with double-meanings and unintended consequences.

The foregoing discussion showed how a question as simple as "Which hand do you write (or draw circles) with?" is in fact a loaded and ill-posed one.

First off because it assumes that you write (draw) with only one hand, which is often not the case.

And second because inherent in that question is a worldview that one hand is "better" than the other, and that only our choice of task assignment to the preferred hand is what matters. Whereas in real life both hands play critical and complementary roles.

Taking this new perspective into account, and applying its lessons to pen and touch as interaction modalities, we can see hints of some fresh interpretations of 'touch' already:

Is the 'touch' articulated with one hand or two?

Is the touch made by the left hand or the right, or perhaps even by both in combination?

Going even further, we can take the scenario illustrated for 'crayon and touch' in Fig. 2 and carry it over to the digital context.

Here, if we look closely enough at manual interaction with "pen and touch" tablets, further nuances and ambiguities rise to the fore:

Fig. 4 What input modality is this? At first glance one might dismiss this as a straightforward illustration of a 'touch' interaction – a user about to put his index finger to the touchscreen. But a more careful look at the *oft seen, but seldom noticed* details of how the participant is using the tablet, and the stylus, with both hands suggests that the answer is not so simple. *Photo © Ken Hinckley, 2014. Used by permission. All rights reserved.*

The kneejerk reaction is that the participant illustrated above is about to place his index finger on the touchscreen, and this is completely correct. However, if we subject it to a discriminating eye for detail, a number of observations can be noted from the moment in time captured by the photograph:

1. **Touch: the obvious interpretation,** reflects nothing more than imminent intentional contact with the touchscreen, with a single finger—the index finger of the preferred hand.

2. **The intentional, but unavoidable, grip of the nonpreferred hand** plays another role here, namely to hold up and skillfully orient the tablet. And to do so, the nonpreferred hand clearly must grip the edge and a portion of the back surface of the device. This is 'touch' as well, albeit in a manner that is traditionally not sensed by existing tablets.

3. **The grasp of the pen**—although secondary to the interaction, with the stylus stowed in the preferred hand by palming—presents us with a third form of touch, albeit one again not traditionally sensed by electronic pens. This behavior of holding the pen at the ready hints that it was recently employed, and furthermore that the user anticipates writing with it again soon. In other words, it informs and situates the context of the interaction; and it distinguishes the impending touchscreen interaction as one performed with the pen-in-hand, as opposed to one performed bare-handed, or by the hand that is not holding the pen.

4. **The incidental, non-prehensile contact of the user's lap** with the bottom edge of the tablet shows yet another manifestation of 'touch' which helps to support the weight of the device. This aspect of touch is perhaps not one of much relevance for direct manipulation, but it does tell us something potentially important about the context of use: the tablet is partially resting on the user's lap, and the user's interactions will be necessarily constrained and encumbered by this situation. Although we leave this point as an exercise for future work, a cleverly designed application could very well take this information into account so as to better afford and accommodate this manner of interaction for its users.

5. **Inadvertent contact:** What about that pesky thumb of the hand gripping the tablet—not to mention the palm of the preferred hand? Accidental touch

is a fifth form of touch that could be occurring (or about to occur) in this photograph. The thumb of the left hand rests on the front surface of the tablet, and as such rests perilously close to the touchscreen. What if the tablet had only a very thin screen bezel, none at all, or perhaps even one that curves around the edges of the device, as has recently come into commercial practice with mobile phones? As well, if the user is not careful, his knuckles or the palm of his hand could brush against the screen the next time he tries to write with that pen-at-the-ready. If the virtue of a touchscreen is that all you have to do is touch something to activate it, this fifth manifestation of touch shows that it has a dark side as well: all you have to do is touch something (by accident), and it might activate! Palm rejection and other related problems of incidental contact with touchscreens all stem from this inherent property of touch.

6. **Pen plus touch, and much more.** Note that we have not even begun to discuss more advanced forms of interaction, such as using the pen and touch in combination [1,6,14,22], or the possibilities of sensing how the pen is oriented—and the tablet is tilted—as the user interacts with the devices in various ways. These, too, are aspects of touch—and pen—that could greatly enrich each modality, as well as their use in tandem.

Collectively, the six perspectives enumerated above show that there is much more to touch interaction than intentional contact with the touchscreen. And while wide-ranging, we make no claim that these six represent an exhaustive list, as we will discuss in more depth shortly in relation to point #3.

These perspectives also show that a holistic consideration of touch is not even limited to one implement: when writing on a tablet equipped with a pen, users necessarily must handle both devices. In this sense a pen is no more a peripheral to a tablet than a pencil is to a sheet of paper – each has an existence, and role, independent of the other, and we could just as easily say that the tablet is a peripheral to the pen.

Finally, and perhaps most importantly, the six perspectives identified above show that not all touches are intentional, or even desired. The user may not even be consciously aware that their hand or some other part of their body has come into contact with the tablet. Or the user may forget that they are even holding the pen at all.

Yet if we think about the interaction from the perspective of the tablet, which can only see what is going on in the very limited plane of the touchscreen, it should be clear that much of the information it needs to fully understand touch is impoverished, or missing altogether.

The same goes for its awareness of what the user is doing with the pen.

As such, there is great potential for additional sensors to augment the context of touch—broadly considered, across all six perspectives—such as to enable more intelligent, more nuanced, and more empowering pen and touch interactions in the future.

But first, by way of example, bear with us for a brief digression that goes deeper into one of these perspectives, namely how people hold the pen.

4 A brief look at pen grips during pen + touch interaction

We just rattled off six ways to look at touch. One of our key points is that the surrounding user behaviors are often nuanced, dependent on task, and influenced by what the user has just done or what the user expects to do next.

To demonstrate this more concretely, we conducted an observational study of how people hold pens while working on electronic tablets with touchscreens. Full details of this study are reported elsewhere [11], but what matters is this:

Some thirty variations of grip and poses of the hand resulted.

For example, people often stow the pen in their hand while engaged in other tasks. This takes two primary forms: tucking the pen between the fingers; or palming the pen with their fingers wrapped around the barrel. Some people even exhibit both behaviors, depending on what exactly they are doing.

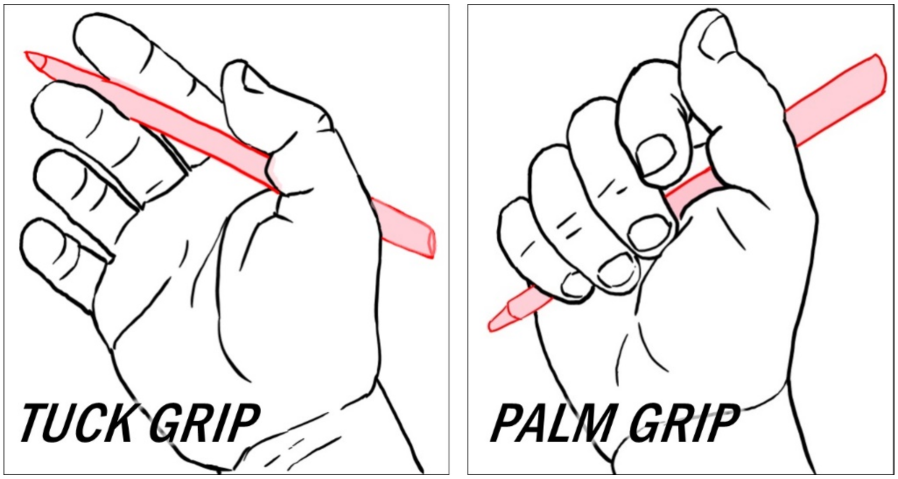

Fig. 5 Two common grips for stowing the pen. *Graphic © Ken Hinckley, 2014. Used by permission. All rights reserved.*

Taking this simple observation as a starting point, we can then go even deeper. From either grip, users will also extend their fingers to tap the screen, pinch to zoom, and perform other touch gestures. But certain common behaviors (such as tapping the screen with the middle finger, for example) were only ever observed from the *Tuck* form of the grip, and never from the *Palm* variation:

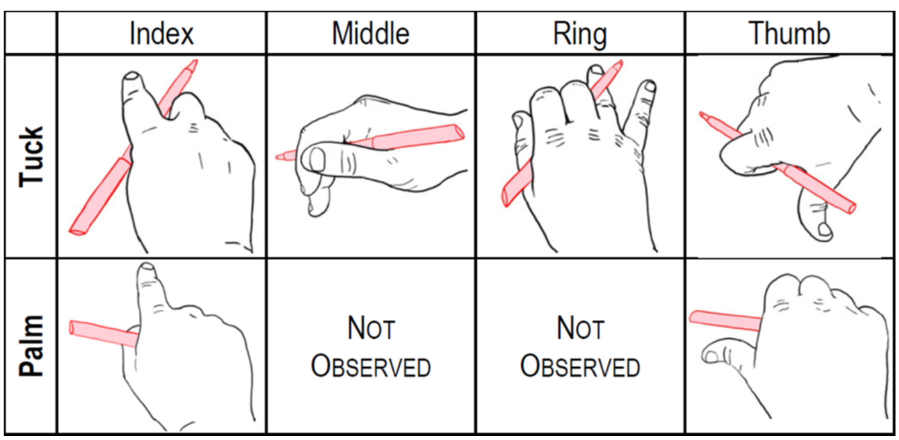

Fig. 6 Ways that users stow the pen while reaching out to touch the screen. We can get a sense of the richness of the user behaviors underlying even something as simple as gripping the pen by looking at how users extend a finger to touch the screen. We observed that the behavior exhibited depends on how the user prefers to stow the pen ('palm' vs. 'tuck') as well as the particular finger brought to bear on the task—which further derives from the particular manner in which that user grips the pen while writing. *Graphic © Ken Hinckley, 2014. Used by permission. All rights reserved.*

Hopefully this gives a sense of the richness of the human manual behaviors surrounding pen and touch interaction—as well as how much they can reveal about what the user is doing and how they are likely to interact with a device.

But to be clear, our goal in the technical exploration which follows was never to recognize all of these grips. Rather, they served as a source of inspiration, while also representing variations in user behavior that our recognition techniques and interaction designs had to accommodate and otherwise take into account.

6 A “hard” look at the context of pen and touch

By now our worldview should be clear. Our goal is to understand touch more deeply, as well as the many ways that people use the pen, so that we can make the combination of pen and touch on tablets richer, simpler, and more satisfying.

Indeed, many of the shortcomings of current pen and touch experiences—such as accidental contact of the hand with the screen while writing—can be viewed not as “user error” but rather as a result of the system’s lack of awareness of what is really going on.

And by extending the tablet’s aura of awareness beyond the confines of the screen itself, we can take some initial steps towards making experiences with technology feel more natural and complete.

Our hardware platform consists of two primary components, tablet and pen, each of which is augmented by a similar array of sensors.

Fig. 7 Augmented stylus prototype in action. The current prototype is tethered, and somewhat bulky, but not so much so that it precludes interactions such as touching the screen while the pen is tucked between the fingers. *Photo © Ken Hinckley, 2014. Used by permission. All rights reserved.*

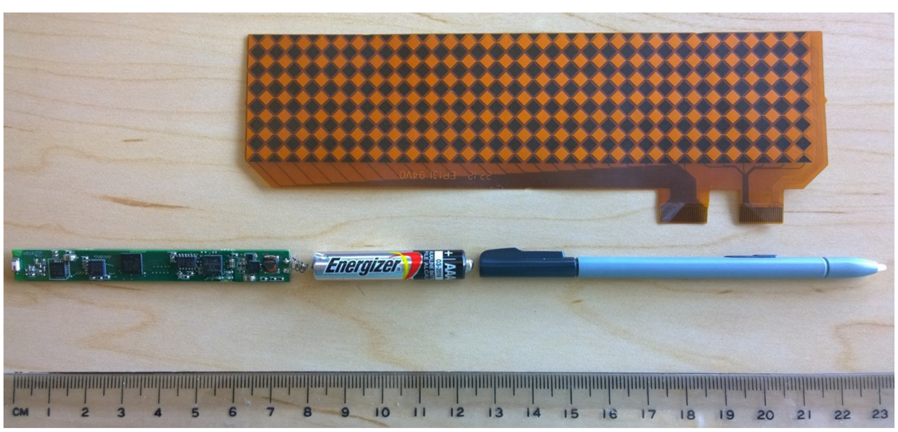

Fig. 8 The augmented stylus, unpacked. The green circuit board on the left contains all augmented sensors. The quad-A battery provides power, but due to the delicate nature of our mechanical prototype it is difficult to change the battery, and hence we reverted to a tether for power (as seen above in Fig. 7). The electronic pen sensed by the tablet digitizer is the Slim Pen from Wacom (MP200). The grip sensing is achieved by a flexible capacitive grid printed on Kapton, which we wrap around the barrel of the assembled stylus. *Photo © Ken Hinckley, 2014. Used by permission. All rights reserved.*

The pen contains an accelerometer, gyroscope, and magnetometer. These represent the standard trio of sensors for detecting inertial motion. Each sensor has three degrees of freedom, such that inertial movement of the device can be fully characterized in three dimensions. We can determine only orientation from these sensors, not 3D position, but still this is enough to infer quite a bit about what is going on.

The entire barrel of the pen is furthermore wrapped in a flexible capacitive grid that provides full multi-touch sensing of the user's grip on the pen. To protect the sensor and provide mechanical stability, we currently sheathe the entire pen in shrink wrap, which adds some bulk and accounts for the garish orange color of our present prototype, but this simply a symptom of our early research prototype that could be done away with in a more meticulously crafted and engineered design.

The tablet itself is outfitted with a similar array of sensors, which provide us with full motion sensing and grip sensing along most of the outer surfaces of the tablet (only the approximately ½ inch wide outer bezel of the front surface—around the screen—is insensitive to touch).

Together, these sensor augmentations therefore give us a complete picture of how the pen is oriented relative to the screen of the tablet, as well as how the user is gripping both devices. Collectively these reveal a great deal about what the user is currently doing (or not doing, as the case may be).

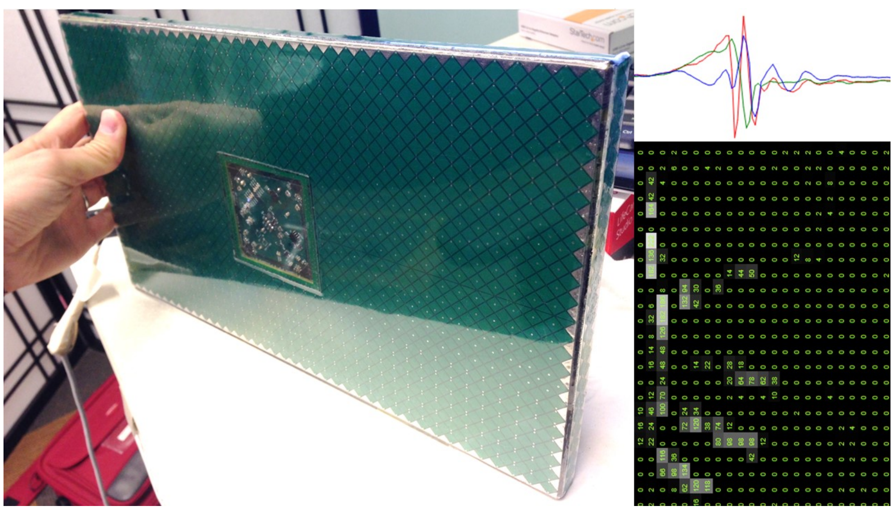

Fig. 9 The back surface of the grip-sensing tablet case. The tablet contains essentially the same circuitry as the pen, except that the grip sensing is integrated into a case which fits over the tablet itself. Grip can be sensed on all edges and the entire back, save for a roughly two square inch area that contains the circuitry. On the right side of the figure, examples of the motion (top) and grip (bottom) sensor data are provided.

Clearly, we went to some effort to construct this Frankensteinian vision of tablet computing. But it was with a clear sense of purpose to address the problems and missing nuances brought on by the lack of awareness suffered by present-day pen and touch interfaces. As such this effort, although a research endeavor, was completely consistent with the belief system and vision of pen computing set forth above—and which we believe will come to practical fruition in the (hopefully near) future, much as sensors have come to pervade everyday experiences with mobile devices and smartphones in the last 10 years [13].

In the following sections we now illustrate some of the ways in which such a sensing platform can address these problems, as well as to bring out greater expressiveness—which is really just greater respect for human manual *and bi-manual* skill, whether with the bare hand or through a mechanical intermediary—in pen and touch experiences for tablet computers.

7 Sensing techniques for stylus + touch interaction

Given a new sensor—which is essentially a new input modality—often the first instinct for an interaction designer is to focus on the new types of gestures that modality can enable:

The new types of *intentional* gestures that one can make when the device is at the *foreground* of the user's attention.

However, while granting that new foreground gestures can have significant value, at least at the outset here we would like to focus on a different way of thinking about input to computers [2], in part because this way of thinking about interaction with technology is so often neglected by designers.

As we have hinted at in the six perspectives outlined above, when we think of 'touch' in its broadest sense, many of the problems and missing pieces result from the activity that surrounds the intentional touch gestures themselves. These are the less often considered aspects of 'touch' that may not even reach the user's conscious attention—the qualities of the interaction and manual habits in the *background.*

The missing context.

Context sensing, then, gives computers greater awareness of what is going on in the background, and the design of background interactions therefore seeks ways to leverage this additional context to automatically enhance and adapt the user experience.

What background sensing really boils down to is the following:

> *There is tremendous potential to resolve ambiguity using sensors rather than foisting complexity on the user to establish the missing context.*

We can see this perspective in the way that modern "point and shoot" cameras work, as one concrete example. The user attends to pointing the camera, while the camera senses the distance to the subject, the illumination levels, and many other properties. When the user "shoots" he need not be burdened by the manual configuration of these many settings.

Another example is the automatic screen rotation [13] now commonplace in mobile phones. The user simply holds the device in the desired posture, while the sensor detects the movement and automatically rotates the screen to the correct orientation.

Here we have adopted this perspective and used it to yield insights—and to design new experiences—for pen and touch interaction.

For example, we can employ our sensors to distinguish some types of intentional versus unintentional touch in the context of palm rejection. Incidental palm contact is inherent to the act of writing. Said another way, the act of writing is a hidden piece of context that could help to resolve the potential ambiguity of palm contact with the touchscreen.

And we can break this down further into a number of telling details that together can lead us towards a more satisfactory solution to this problem:

1. **Palm contact while writing typically results from the hand holding the pen**—that is, from the preferred hand—and not the other hand, which may be intentionally holding or moving objects, or manipulating the page itself.

2. **To write the user must hold the pen in a writing grip—as opposed to palming the pen, or tucking it between his fingers.** Hence the type of grip is a major clue as to whether the user is about to start writing—and whether a touch articulated by the preferred hand is consistent with intentional contact with the screen, or merely incident to the act of writing itself.

3. **The pen must be held at a certain orientation to write**. The orientation of the pen is not an unambiguous signal in and of itself, but it complements the information revealed by grip and provides another clue to user activity.

4. **Preferred hand motion and pen motion are necessarily correlated.** When the user moves, gestures, or touches the screen while the pen is held

in hand, the pen moves too. When the hand makes contact with the screen, if we have good enough sensors and we maintain a high sampling rate, this subtle discontinuity in the movement can also be picked up by the motion sensors.

Specifically, when the user rests his or her palm on the screen, we can sense that the pen is held in a writing grip. And we can also sense the resulting “bump” signal produced by the hard-contact force of the palm with the display.

This then lets us associate the hand contact with the hand holding the pen, because the new touch reported by the touchscreen occurs *at the same time* as the bump signal sensed through the pen.

Meanwhile, an intentional bare-handed touch with the nonpreferred hand produces no such signal, since the pen doesn’t move in a manner consistent with the nonpreferred hand’s motion.

In this way we can screen out many incidental palm contacts while simultaneously allowing full bimanual interaction with both hands.

And we can still allow *intentional* preferred-hand touches, even while the pen is held tucked between the fingers, because such a grip (and the orientation of the pen) is not consistent with the user placing a hand on the screen to write.

We do not claim this approach is perfect—certain signals can still fool our straightforward signal-processing software, particularly if the palm contact is very light—but it does show how the palm-rejection problem can be re-cast as one of insufficient context.

And just as important, this approach allows us to *permit intentional touch* instead of solving the problem by discounting touches coincident with pen interaction altogether—a measure operating systems, device firmware, and applications often resort to. If we took such an approach, it would preclude the entire space of simultaneous pen + touch interaction.

However, these inferences potentially can allow pen and touch interfaces to do far more than handle incidental touch better. The four telling details listed above make clear that we can now distinguish touches with the pen-at-hand from other bare-handed touches, such as those produced by the nonpreferred hand.

Our prototype therefore employs this distinction to support fresh nuances of expression, such as:

Fig. 10 Tapping the screen with the pen-at-hand brings up the pen tools. *Photo © Ken Hinckley, 2014. Used by permission. All rights reserved.*

Tap the screen with the pen in hand to bring up the pen tools. —A menu bar with different pen colors and thicknesses, as well as other pen tool modes such as lasso selection, the highlighter tool, an eraser tool, and so forth comes up when the user taps the screen with the preferred hand while the pen is tucked. Our reasoning was that it was only logical to bring up the tools associated with the pen when tapping the screen in this manner. The pen tucked between the fingers indicates that the user intends to use the pen again soon. This gives a simple and intuitive way to access all the tools, modes, and settings associated with the pen. Tapping with the (bare) nonpreferred hand, by contrast, produces an "ordinary" tap.

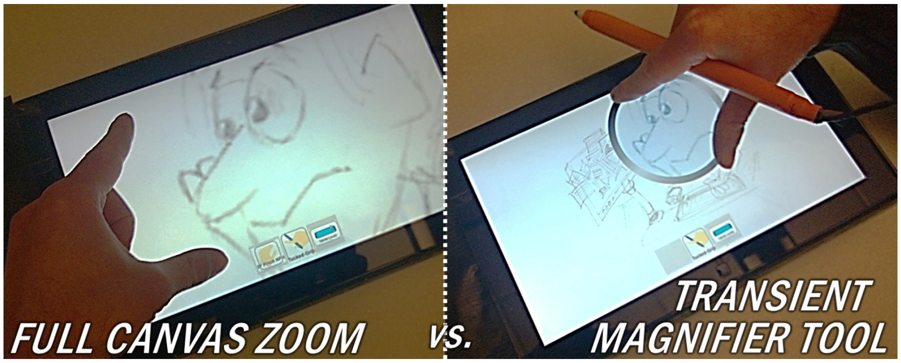

Fig. 11 Pinching with the nonpreferred hand zooms the full canvas, while a preferred hand pinch with the pen-in-hand brings up the transient magnifier tool. In this way we support a new class of multi-touch gesture which incorporates the context of *which had produces the touch* as well as *whether (and how) the pen is being gripped* at the time of the touchscreen contact.

Nonpreferred hand: Full canvas zoom. As is the standard idiom in touch interfaces, the nonpreferred hand can be used to pinch to zoom (Fig. 11, left). However, in this case, the user can pinch-to-zoom even if the palm (of the preferred hand) is resting on the screen preparatory to writing—because our sensors let us distinguish which touch is produced by which hand. Furthermore, this lets us entertain a design where a nonpreferred-hand pinch has a different function than one articulated by the preferred hand, as explored in the following.

Preferred hand: Transient magnifier tool. Touching down with two fingers while the pen is tucked-in-hand brings up the magnifier tool (Fig. 11, right). This provides localized zooming in one area of the canvas, which is idea for detail work with the pen—without disrupting the 'surround' of the workspace. Hence performing a pinch while the pen is at hand brings up another tool uniquely suited for the pen (much like the pen tools of Fig. 10).

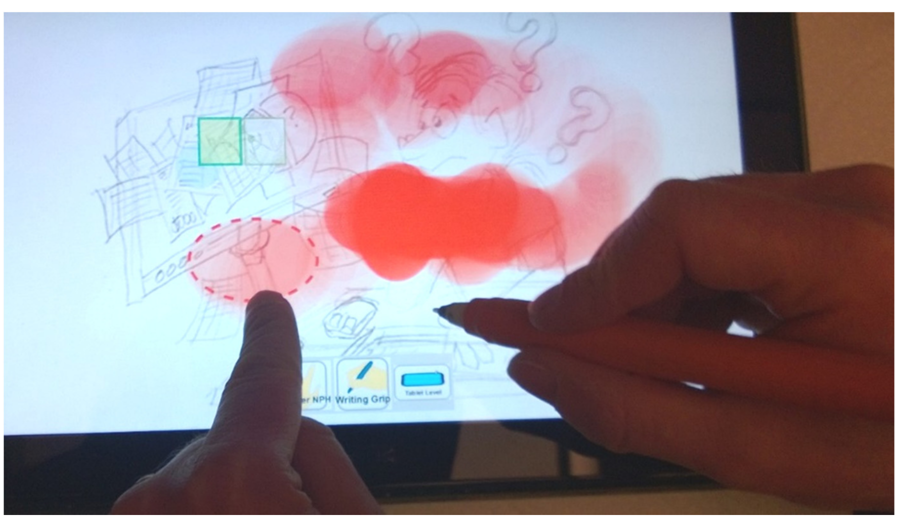

Fig. 12 The airbrush tool. A touch from the nonpreferred hand indicates where to airbrush, while the preferred hand orients the pen in 3D to control the conic section of the resulting spray. Note that the pen does not have to stay "within range" of the traditional ~1cm sensing range of the tablet digitizer during this interaction, because our augmented sensors provide the 3D orientation information.

Drafting tools suited to the writing grip. When the user is actively working with the pen, they often maintain a writing grip. We therefore use this to offer various drafting tools when the user taps the screen with the nonpreferred hand *while the pen is held at the ready in the writing grip*. For example, we provide a Compass tool (for drawing arcs) and an Airbrush tool in this mode (Fig. 12). The Airbrush tool, in particular, is interesting because it uses the additional degrees-of-freedom afforded by our pen's orientation sensing capabilities to provide a natural airbrush capability that allows the user to adjust the conic section of the spray that results by tilting the pen.

8 Tablet grip sensing

In the early going, we motivated our work by illustrating the many ways that users can 'touch' a tablet computer that are not respected as such by traditional touchscreen interaction. However, in the examples given above we primarily focused on interaction scenarios with the pen-in-hand, often in cases where both hands were interacting with the front of the device.

In part this is simply because even with fairly limited sensing we can do a lot. And because the techniques illustrated above do not necessarily depend on tablet grip sensing to realize.

But having said that, we believe the techniques above could be improved by more fully integrating the context of the current *tablet* grip into the interaction. That we have not yet done so and demonstrated more fully why it could valuable is more a symptom of limited time (and in some cases, perhaps limited imagination) than an inherent limitation of what is possible the sensing modality. And indeed we hope that others will pick up this line of work and fill in the gaps, look for additional or alternative mappings for how these nuances of pen and touch can be leveraged in user interfaces, and generally advance this area in many other ways.

Nonetheless we have explored many other techniques that do rely on the tablet grip sensing, some of which appear in our earlier work that was the foundation for this paper [11], and in another forthcoming paper (accepted, final revision in progress [21]).

For example, we mentioned that "pesky thumb" on the front of the screen. One simple remediation we have experimented with is to discount thumb contact that occurs incident to gripping the tablet, such as when first picking it up (Fig. 13).

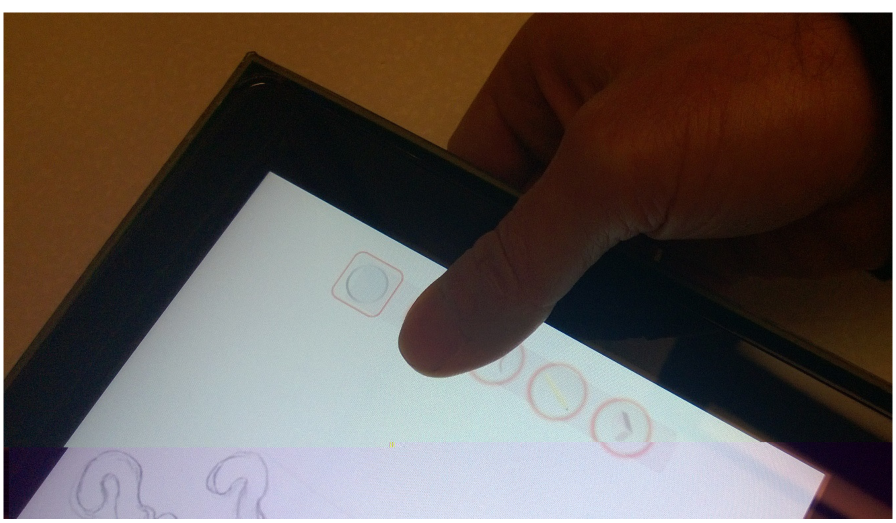

Fig. 13 Discounting thumb contact when initially gripping the device. The hand grip sensed on the back of the device tells us that the touch, immediately above on the front of the screen and near the edge, is necessarily associated with the thumb of the same hand. Controls near the edge of the screen initially appear semi-transparent and cannot be triggered by such incidental touches from the thumb.

Another example is to sense when the user walks away while holding the tablet—but not the pen. This can be sensed in a straightforward manner by sensing the gait pattern of walking, while also noting that the pen is neither held (gripped) nor moving [9]. The user can then be reminded not to forget their pen. Since in pen computing it is always a matter of *when*, not *if*, the user will lose the pen, such a simple technique could be of considerable value to users.

A third example is to sense common ways of holding or using the tablet. For example, we observed that while reading books or paper documents, users pick up the material with both hands and angle it closer to themselves when deeply engaged in reading. This can be used to support an immersive reading mode, which automatically emphasizes the text while removing background distractions and ancillary screen controls, when the user grips their tablet in this manner.

A fourth example is to sense when the user passes their tablet to another user. People exhibit very distinctive patterns of grip and orienting the device when handing an object to another person. We have explored various techniques to sense this and to use it to support alternative semantics of "sharing" digital content with another that are much more akin to physical sharing and therefore also very different than emailing an attachment, or placing a file in the cloud, for example.

Finally, we have explored how grip sensing can help a user to more effectively refer to and cross-reference information across multiple devices. Tablets and e-readers are increasingly being used as companion devices—often with multiple devices to support facile work with multiple documents and information sources, as repeatedly observed in natural information work with both paper and electronic reference materials [3,15,17,19]. Sensing which devices the user is holding, therefore, and how they are gripping them therefore can potentially be a crucial piece of missing context in managing multi-device interaction [3,5,15,16].

9 Summary, Conclusion, and Future Work

We have come a long way in this discussion. We have demonstrated how a keen eye for the *oft seen, but seldom noticed* role of manual (and bi-manual!) interaction in general – and in the behaviors manifested by users with pen and touch tablets in particular – lead to at least six perspectives of what it means to 'touch' a device. And there are likely many more beyond that.

We have argued that many of these missing perspectives in touch interaction are a symptom of missing context, we set out to build sensor-augmented stylus and tablet devices to redress this. The prototype devices support inertial motion sensing as well as capacitive grip sensing, but certainly other sensors could be envisioned to accompany these modalities in the future. But even with a limited palette of additional sensors, we showed how we can detect a number of salient elements of context: how the user is gripping the pen, which hand is contacting the display, and what incidental motions are imparted to the devices during these interactions. We then presented techniques which illustrated how these sensors and contextual inferences allowed us to support richer, more natural, and more nuanced pen and touch interaction.

While the directions explored are hopefully provocative, and have been explored through implementing prototypes and by conducting observational studies as well as preliminary technological evaluations with test users, many of them remain speculative in the sense that their potential is not yet fully clear. Some of the techniques we propose may ultimately find broad acceptance, while others are likely to remain laboratory curiosities for the foreseeable future. And the truth of the matter is that – without the benefit of hindsight – it is often difficult to tell which is which. However, we take heart in some of our own early work on mobile sensing [2,10,13], which in some ways was indeed ahead of its time, but in other ways felt painfully similar to the present project. *Certainly such sensors are too expensive,* or so we heard from practically-minded people, *the power they demand too much, their interpretation too uncertain and too taxing of limited computational cycles.*

But here we are in 2015, and many of the techniques we proposed have come to fruition. And even if not always in the exact way we implemented it and anticipated it coming along, in several instances it was awfully close. We hope, and expect, that the same will be true of the techniques we have explored here for pen and touch, for stylus and tablet, some fifteen years from now – if not much sooner.

In the meantime we will continue to strive to go much further, and see how the new capabilities and nuances afforded by rich sensing capabilities may allow us to design interaction technologies, techniques, and experiences in new ways. Our goal remains to address latent and unmet needs – for freeform input, for more human ways of informally organizing information, and in the broadest sense for *inking outside the box* . Ultimately, we believe that such advances can help students, information workers, creative professionals, and educators to work together and to bring their efforts to fruition in a much deeper and more natural way.

Acknowledgements The work underlying this paper is the result of many years of study and prototyping by a large number of people. In particular we would like to acknowledge the efforts of Michel Pahud, Hrvoje Benko, Pourang Irani, Marcel Gavriliu, Francois Guimbretiere, Xiang 'Anthony' Chen, Fabrice Matulic, and Andy Wilson on our Stylus + Tablet Sensing Techniques paper [11] which prompted our invited presentation at WIPTTE (and hence, the present manuscript).

References

1. Brandl, P., Forlines, C., Wigdor, D., Haller, M., Shen, C. Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces. *AVI'08*. 2008.
2. Buxton, W. Integrating the Periphery and Context: A New Taxonomy of Telematics. *Proceedings of Graphics Interface '95*. 1995. Quebec City, Quebec, Canada.
3. Chen, N., Guimbretiere, F., Sellen, A., Designing a Multi-Slate Reading Environment to Support Active Reading Activities. TOCHI, 2012. **19**(3): p. Article 18 (35 pp.).
4. Guiard, Y., Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. Journal of Motor Behavior, 1987. **19**(4): p. 486-517.
5. Hamilton, P., Wigdor, D.J. Conductor: enabling and understanding cross-device interaction. *CHI '14*. 2014.
6. Hamilton, W., Kerne, A., Robbins, T. High-performance pen + touch modality interactions: a real-time strategy game eSports context, *UIST '12*. 2012.
7. Hinckley, K., *Haptic issues for virtual manipulation*, in *Department of Computer Science*. 1997, University of Virginia: Virginia.
8. Hinckley, K., Bi, X., Pahud, M., Buxton, B. Informal Information Gathering Techniques for Active Reading. *CHI 2012*. 2012.
9. Hinckley, K., Chen, X.A., Benko, H. Motion and context sensing techniques for pen computing. *Proceedings of Graphics Interface 2013 (GI '13)*. 2013: Canadian Information Processing Society.
10. Hinckley, K., Horvitz, E. Toward more sensitive mobile phones. *ACM UIST Symposium on User Interface Sofware and Technology*. 2001. New York, NY: ACM.
11. Hinckley, K., et al. Sensing Techniques for Tablet+Stylus Interaction. *UIST'14*. 2014.
12. Hinckley, K., Pahud, M., Buxton, B. Direct Display Interaction via Simultaneous Pen + Multi-touch Input. *Society for Information Display (SID) Symposium Digest of Technical Papers*. 2010.
13. Hinckley, K., Pierce, J., Sinclair, M., Horvitz, E. Sensing techniques for mobile interaction. *ACM UIST Symposium on User Interface Software and Technology*. 2000. New York, NY: ACM.
14. Hinckley, K., et al. Pen + Touch = New Tools. *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. 2010. New York, New York, USA: ACM.
15. Hong, M., Piper, A.M., Weibel, N., Olberding, S., Hollan, J. Microanalysis of active reading behavior to inform design of interactive desktop workspaces. *ITS '12*. 2012.
16. Marquardt, N., Hinckley, K., Greenberg, S. Cross-Device Interaction via Micro-mobility and F-formations. *UIST '12*. 2012.

17. Morris, M.R., Brush, A.J.B., Meyers, B.R. Reading Revisited: Evaluating the Usability of Digital Display Surfaces for Active Reading Tasks. *Proc. Tabletop'07*. 2007.
18. Schilit, B.N., Golovchinsky, G., Price, M.N. Beyond paper: supporting active reading with free form digital ink annotations. *ACM CHI Conference on Human Factors in Computing Systems*. 1998: ACM.
19. Sellen, A.J., Harper, H.R., *The myth of the paperless office*. 2002, Cambridge, MA: MIT Press.
20. Tashman, C., Edwards, W.K. Active Reading and Its Discontents: The Situations, Problems and Ideas of Readers. *CHI'11*. 2011.
21. Yoon, D., et al. Sensing Tablet Grasp + Micro-mobility for Active Reading. *UIST'14*. 2014.
22. Zeleznik, R., Bragdon, A., Adeputra, F., Ko, H.-S. Hands-on math: a page-based multi-touch and pen desktop for technical work and problem solving. *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. 2010. New York, New York, USA: ACM.