Manual Deskterity: An Exploration of Simultaneous Pen + Touch Direct Input

Ken Hinckley[1]
kenh@microsoft.com

Koji Yatani[1,3]
koji@dgp.toronto.edu

Michel Pahud[1]
mpahud@microsoft.com

Nicole Coddington[2]
nicolec@microsoft.com

Jenny Rodenhouse[2]
jrodenh@microsoft.com

Andy Wilson[1]
awilson@microsoft.com

Hrvoje Benko[1]
benko@microsoft.com

Bill Buxton[1]
bibuxton@microsoft.com

[1]Microsoft Research
[2]Microsoft Corporation
One Microsoft Way
Redmond, WA 98052 USA

[3]University of Toronto
Dept. of Computer Science
40 St. George St. Rm. 5170
Toronto, ON, M5S 2E4, CANADA

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.
ACM 978-1-60558-930-5/10/04.

Abstract
Manual Deskterity is a prototype digital drafting table that supports both pen and touch input. We explore a division of labor between pen and touch that flows from natural human skill and differentiation of roles of the hands. We also explore the simultaneous use of pen and touch to support novel compound gestures.

Keywords
Pen, touch, gestures, tabletop, tablets, bimanual input

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: Input

General Terms
Human Factors, Design

Introduction
We are witnessing a shift towards systems employing direct manual input where the user interacts directly with the display, rather than indirectly, as with the mouse and cursor of traditional GUI's. This has renewed interest in both pen and touch input, in form factors ranging from hand-helds, slates, desk-tops, table-tops, and wall displays. The iPhone, Tablet PC, Wacom Cintiq, Microsoft Surface, and Smartboard are, respectively, examples of each. Neither multi-touch nor pen input are new, but few systems explore their use in conjunction [3,27,28]. The presence of both modalities may alter our perspective on multi-touch input; the same can be said for pen gestures. Hence simultaneous *pen + touch* is a nascent topic in need of further study.

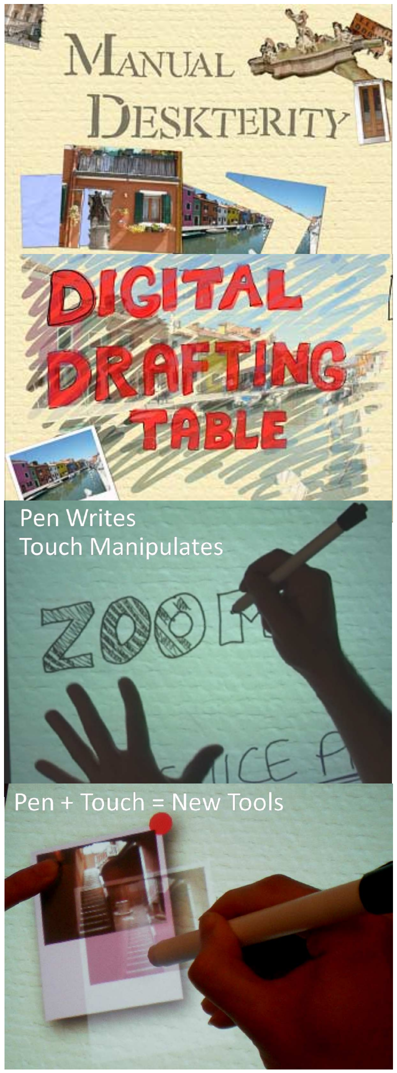

figure 1. *Manual Deskterity* prototype: The pen writes, touch manipulates, and pen+touch yields new tools.

This argues for a holistic approach rather than focused evaluation of individual techniques [12], even though this is often not rewarded by CHI. Our experience is that trying many ideas [4]— some good, some bad, and some intentionally chosen to highlight conflicting conventions or thorny design decisions rather than hide them [11]— is an excellent way to draw out nuances and gain insights into novel input modalities (e.g. [17,20,21]). Our systems-oriented approach offers a realistic perspective of how combined pen and touch input influences UI design issues and trade-offs.

The result is *Manual Deskterity*, a scrapbooking application inspired by how designers work with design boards and notebooks [4,14]—plus our experiences with related prototypes [15,18]. We advocate a division of labor between pen and touch: *the pen writes, touch manipulates, and the combination of pen+touch yields new tools*. This articulates how our system interprets unimodal pen, unimodal touch, and multimodal pen + touch inputs, respectively. We contribute novel pen + touch gestures, while also raising, by way of examples, design questions that probe how the roles of pen and touch should be differentiated (or not) in UI design.

Related Work

Many current direct input systems employ only one of touch or pen input. Yet an earlier generation of devices, such as the Bell-South/IBM Simon smartphone (1993), the Psion Series 5 PDA (1995), and the Palm Pilot (1996), supported use of either pen or touch. Part of what limited these earlier systems was that the technology could not differentiate pen contact from finger contact. Emerging dual-mode digitizers distinguish pen and touch [9], but existing drivers do not yet support the two simultaneously.

Several research efforts explore the combination of pen and touch. Yee [28] uses single-touch + pen input to support panning a canvas while drawing with the pen. Wu [27] describes two combined pen and touch gestures. Brandl [3] explores bimanual pen + multi-touch techniques that assign the pen to the preferred hand and touch to the nonpreferred hand. We tease apart the factors of (1) *pen vs. touch,* (2) *preferred vs. nonpreferred hand assignment,* and (3) *unimanual vs. bimanual interaction*. For example, we consider unimanual cases where the user interleaves pen and touch interactions with the preferred hand, and we explore a wider vocabulary of novel pen + touch gestures that afford compound transactions.

Cohen discusses the complementary role of natural language and pen gestures [7]; he treats multimodal input with a probabilistic approach [8]. We instead treat pen+ touch input in a manner that affords *deterministic* state-machine-driven GUI's [6]. Also, because pen and touch are both *manual* input modalities, the nuances of how the two complement one another are more subtle, and we must overcome a longer legacy of designs that have treated pen or touch interchangeably.

Guiard [13] observes that the hands cooperate to accomplish tasks, so the question is not "Which hand is better?" but rather "What is the logic of the division of labor between the hands?" Likewise, in our research we ask: *What is the logic of the division of labor between pen and touch in interface design?* Guiard observes that the nonpreferred hand frames the action of the preferred hand. Our bimanual pen+touch gestures build on this: pen gestures act upon an object, and are phrased together by muscular tension [5] from the user's nonpreferred-hand fingers held on the object.

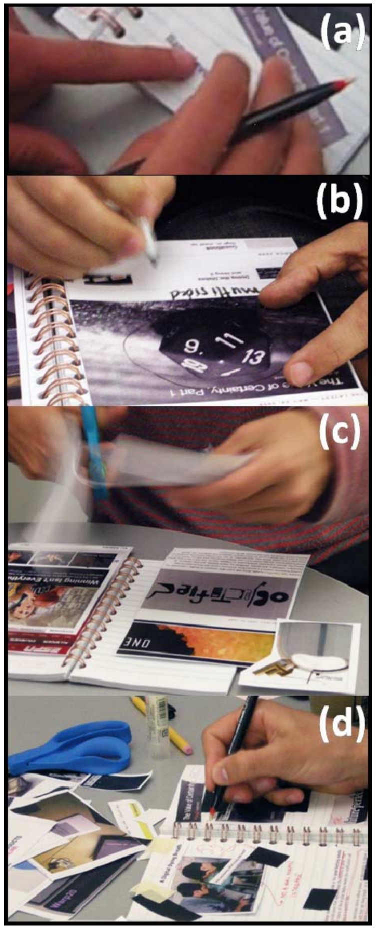

figure 2. *User behaviors observed during design study*. **(a)** Users tuck the pen between fingers while manipulating items. **(b)** Thumb and forefinger grasp an item while writing about it. **(c)** Cut scraps fall onto the work surface. **(d)** Users often pull tools and new content onto the notebook from above.

Frisch describes a user-elicited collection of touch and pen gestures [10], and reports that users often treat pen and touch interchangeably. But as a result, the user-defined gesture set contains many ambiguities. Which pen, touch, or pen + touch gestures should a system support (or not support), and why? In our experience, user-defined gestures are insightful but must be taken with a grain of salt because users have difficulty envisioning how they would employ new modalities of which they have little or no experience.

Indeed, our exploration of *Manual Deskterity* convinces us that if each input modality offers complete coverage of all possible interactions, it quickly robs the combination of pen and touch of much of its vigor. Differentiating between pen and touch, rather than treating them interchangeably, offers a consistent and rich *designed* input vocabulary. Nonetheless we build our gestures on a vocabulary of natural occurring bimanual actions, as shown in the following study.

Design Study Using a Paper Notebook

We conducted an observational study to gain insight into how people naturally work with pens, tools, and pieces of paper. We asked each participant to illustrate ideas for a hypothetical short film by pasting and annotating clippings in a paper notebook. To simulate a slate computer where the user could move between pages, we provided a small paper notebook as the authoring space. We provided users with pens, tape, scissors, and 20 sheets of inspirational materials.

Eight people participated in the study. We looked for patterns in how users gestured and structured their working space. We observed behaviors (B1-B9) that informed specific gestures and features in our system:

B1. Participants tucked the pen between the fingers of their preferred hand when interleaving writing and moving clippings (Fig. 2a). People were remarkably adept at interleaving pen and touch in this manner. As a result we consider unimanual multi-touch gestures performed while the pen is tucked.

B2. Participants temporarily held clippings in place with one finger of the nonpreferred hand (Fig. 2a).

B3. Participants exhibited a strong tendency to hold a clipping with their nonpreferred hand while writing about it with the pen (Fig. 2b).

B4. A common hand posture was to frame a clipping with thumb and index finger (Fig. 2b) while writing about it. This seemed to help users *mentally focus on a source object and reference annotations to it*.

B5. Participants used only parts of the inspirational materials. They cut a sheet while holding it in their nonpreferred hand, above the notebook. The unwanted part fell onto the page (Fig. 2c). Users occasionally adopted these scraps into their work.

B6. Participants arranged the workspace with the notebook proximal to their body, while reaching above it to access tools and materials (Fig. 2d).

B7. Piling clippings was a common behavior. Users formed piles of "interesting" items while holding the remaining items in the nonpreferred hand.

B8. While not a common behavior, a couple of people did employ clippings as a constraint for the pen, to draw a border around an item (Fig. 3e).

B9. Tearing sheets of paper was a bimanual behavior performed with the fingers (Fig. 3f).

The behavioral observations above form a valuable contribution. By starting from a suitable task context with physical objects, our approach elicits a naturally

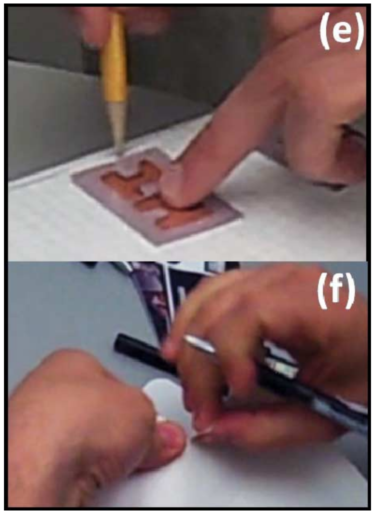

figure 3. *Additional Behaviors observed during design study*.
(e) Drawing a border around a clipping by holding it with the nonpreferred hand while using its edge to constrain the path followed by the pen held in the preferred hand. **(f)** Tearing a page by anchoring it with a thumb while pulling it away using fingers of the opposite hand. Note that in this example, the pen is again tucked between the fingers of the preferred hand, demonstrating how prevalent this was with our participants during a variety of manual activities.

occurring set of unimanual and bimanual behaviors from users, with both pen and bare-hand manipulation, that exhibits a clear differentiation between the roles of pen and touch. Terrenghi's design study notes a similar finding: physical manipulation of real-world clippings on a tabletop yielded rich bimanual interactions, whereas corresponding digital ones did not [25]. These stand in contrast to the results of user-elicited gestures [10], which might otherwise lead us to treat pen and touch interchangeably. We feel it would be mistaken to do so.

We should also emphasize here that our intent is not to mimic the specific actions required to work with physical paper. Pen + touch gestures should go beyond physical paper, but the best foundation for such gestures likely lies in behaviors that people already exhibit when working with pen, paper, and clippings.

Implementation

We use Microsoft Surface for our prototype. The pen uses an infrared LED, activated during surface contact via a tip switch. The pen is much brighter than hand/finger contacts, so we can robustly identify the pen as the brightest spot in the image. The software is written in C# with WPF and the Microsoft Surface SDK.

A potential limitation of pen + touch input is the so-called "palm rejection" problem: the user may rest his hand on the screen while writing, potentially leading to unintended operations. Brandl [3] makes no mention of this issue, even though he shows a black glove on the user's pen hand to prevent the digitizer from reacting to palm contact. We treat touches with a large contact area as incidental, which is sufficient for prototyping pen + touch techniques. Robust handling of incidental contacts remains an important problem for future work.

Application Scope and Motivation

Manual Deskterity is in some respects a "toy" system intended primarily as a research vehicle to explore pen + touch, which we believe has many potential applications. Nonetheless we emphasize practical functionality well-suited to "idea collection," note-taking, and mark-up functionalities that have been well documented by previous research (e.g. [15,18,23]). For example, several papers note the importance of writing, annotation, selecting, copying, arranging, and aggregating objects both in digital [1,27] and physical contexts [2,19,26]. Our design study suggests additional behaviors of interest in the context of pen + touch, such as holding items while acting on them with the pen (B2, B3, B4), cutting and tearing operations (B5, B9), creating objects from the space above the notebook (B6), and employing clippings as a constraint for the pen (observation B8). These formed the basis of the features we elected to explore using pen, touch, and pen+touch interactions.

Core Tasks: Pen Writes, Touch Manipulates

The core interactions of *Manual Deskterity* are driven by multi-touch interactions including zooming, flipping pages, moving objects, selecting objects, and creating new objects such as digital post-it notes. The other central task is writing; here, only the pen produces ink strokes (although in some contexts the finger "smears" colors, as discussed later). Notwithstanding the exception of finger-painting, for these core tasks *the pen writes, and touch manipulates*, period. This makes the entire canvas, and objects on the canvas, available for immediate annotation with the pen, while pan/zoom, page navigation, and object manipulation are also immediately available via touch, without any explicit mode switches.

figure 4. *Selection and Context Menu*. **(Top)** Selection feedback consists of a pink highlight and a large drop shadow to visually "float" the object. Deselected objects have only a thin drop-shadow. **(Bottom)** Context menus use a radial (marking) menu. The user touches pen or finger to the red dot and strokes towards the desired command. This represents one situation where it is necessary to support both pen and touch interchangeably to match user expectations, even though elsewhere the system sticks to a differentiation of roles between pen and touch.

Selecting and Manipulating Objects

Our prototype scrapbooking application includes objects such as photos and post-it notes that the user arranges on the canvas via direct manipulation (single touch drags, while multi-touch rotates and scales). The user defines a collective scope (i.e., multiple object selection) via finger-taps that incrementally add individual objects to the current selection. Dragging an item that is already selected drags it *and all other selected items* while maintaining their relative spatial relationship. Dragging an item that is *not* already selected drags *only that item*; this enables multi-touch dragging of items to multiple different locations.

Context menus. On selection, a radial menu appears at the upper right corner of objects (Fig. 4). Initially, our radial menus required use of the pen, but users uniformly expect this menu to be operable via touch as well. As a high level principle we still advocate differentiation of roles: *the pen writes, and touch manipulates*. But acting on the radial menu represents a limited spatially-multiplexed context where pen and touch should indeed be treated interchangeably.

Creating new objects. In our design study we observed that users bring in new materials from above (B6). We incorporate this observation into our interface design via a finger-activated *bezel menu* that builds on Bezel Swipe [24]. The user performs a continuous finger-drag that crosses the top screen bezel and onto the canvas to *create and position* objects of various types, such as digital post-its (see video). The user can then annotate the new object immediately.

Summary. By supporting core manipulative tasks with touch, and with touch only, *Manual Deskterity* enables users to fluidly interleave annotation and other secondary tasks. This approach also supports *graceful degradation to one-handed usage*, which we believe is an important property to afford mobile pen+touch form-factors when they become available in the future. Nonetheless two-handed interaction is encouraged when the usage context makes it suitable (e.g. working at a desk, rather than while mobile). The key is that ***the mode switch is in the user's hand***: he can work one-handed and flip between pen and touch by rapidly tucking the pen between fingers, or he can work with two hands by performing most touch operations with the nonpreferred hand, while writing and annotating with the pen in the preferred hand.

Pen + Touch Yields New Tools

Next, we explored how the expanded input vocabulary afforded by the combination of pen and touch can give new tools. Our design keeps the primitives for pen+touch operations simple. The richness of the gestures arises from how the primitives are combined. For the pen the primitives we use are tap, drag-off, crossing, or drawing a stroke. For multi-touch, we employ single-finger tap, single-finger hold (i.e. a tap with a long duration, as seen in our design study, behavior B2), holding with thumb and forefinger (B4), and crossing. We do not implement all combinations, but rather support a sufficiently rich set of operations, with semantics that map well to our application domain, to illustrate the expressiveness of our approach.

The pen+touch techniques described below all use fingers of the nonpreferred hand hold an item, while the pen acts in reference to the item. This builds on the tendency that we observed for users to hold clippings with the nonpreferred hand while making pen markings

figure 5. *Stacks of items formed via the Stapler pen+touch gesture. The user staples items by (1) tap-selecting a series of items, (2) holding a finger on the representative item for the stack, and then (3) tapping the pen on that item while continuing to hold it. The representative item appears on the top of the resulting stack.*

figure 6. *Intricate X-Acto cuts formed by one user during an informal evaluation of our system.*

in reference to them (B3), and it also corresponds well to Guiard's principle that the nonpreferred hand sets the frame of reference for the preferred hand [13]. Thus, we construct gestures that allow for non-physical digital effects, yet remain grounded in people's naturally occurring behaviors with physical paper.

Stapler: Grouping Items into a Stack

To support piling (design study B7) in an intuitive way that reduces the manual effort required to drag widely scattered items into piles, we support stapling items into a stack (Fig. 5). The user can finger-tap-select a number of items, and then staple all of them together by holding an item and tapping it with the pen. The item that the user holds appears on the top of the stack, thus promoting it to represent the entire stack.

Tap-selecting a series of items and stapling them together enables quick tidying of a messy work surface into a few piles (see video). This transaction separates the identification of the items to stack from the decision of which item should become the representative item on the top of the resulting stack. Performing the pen+touch gesture on the representative item keeps the user's attention focused on it, at exactly the moment the user makes this decision, which seems to correspond well with users' mental model of the task.

X-acto Knife: Cutting Items (and Tearing Items)

The user can turn the pen into an X-acto knife by holding an object and fully crossing it with the pen. That is, the pen stroke starts outside the object, crosses through the interior of the object, and finishes on the exterior of the object. The pen stroke within the item can follow any path, allowing intricate cuts if desired (Fig. 6). When the pen exits the item, both the cut and the scrap piece appear on the page, following the real-world action where we observed scraps falling onto the work surface (design study B5).

To probe the semantics of pen vs. touch in analogous gestures, we also implemented *tearing* items by holding an object with a single finger, and then crossing the item with another finger (study, B9). This tears the item along the line connecting the entry and exit points of the finger. The operation is similar to cutting, but produces a different visual affordance in the scraps. This technique demonstrates how touch can sometimes be used to sneak a nuance of expression into a transaction, by applying a different look or different default command parameters. On the other hand, in our system this precludes using the touch gesture for a different command, such as layering. Should the semantics of analogous pen+touch vs. touch+touch gestures be similar or contrasting? This remains an open design question raised by this example.

Carbon Copy: Drag-Off with the Pen

The user copies an object by holding with a finger (B2) and then "peeling off" a copy with the pen (Fig. 7). This gesture is similar to a *copy* gesture identified by Frisch [10], but here we identify the interaction pragmatics as well as why this gesture differs from a *Copy* command. While we have so far been somewhat critical of the user-elicited gestures methodology, in this case we seen an example of how the methodology can yield fertile ground for suggesting plausible gestures (so long as we keep in mind that users are not designers).

Once the pen drags away by a minimum distance, a semi-transparent copy of the object appears attached to the pen. As the pen continues to drag, the object

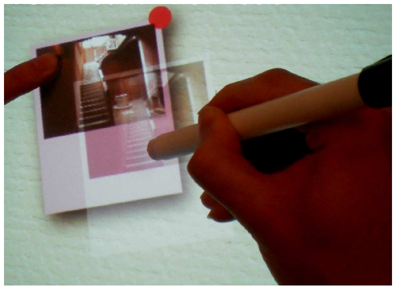

figure 7. Duplicating an item by holding it with a finger of the non-preferred hand while dragging away with the pen.

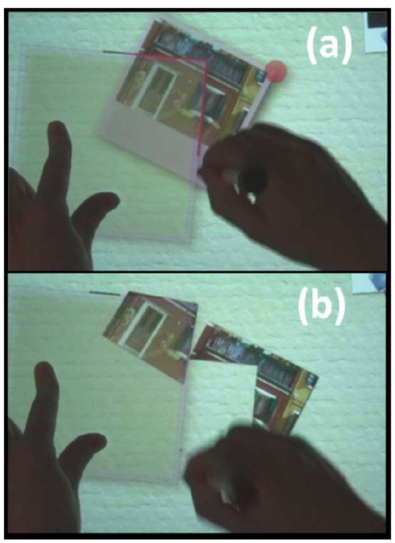

figure 8. Composing the Ruler tool with the X-acto tool. The user can cut along a straightedge, including around corners, as shown here, to produce interesting cutting effects.

becomes opaque. The user may then proceed to drag the object to "paste" it at the desired location. Any annotations on top of the object are also copied.

The properties of this approach differ from using the *Copy* command of the object's context menu (Fig. 4, bottom). *Drag-off with the pen* phrases together the entire transaction (select, copy, and drag to final position) into a single cognitive chunk [5] via the muscular tension of the nonpreferred hand holding down the original item. We found that this corresponds well to users' mental model of duplicating items– they not only want to copy the item, but also place the duplicate at a particular location. By contrast, *Copy* from the context menu divides the select-copy-position transaction into multiple steps, enabling one-handed copying at the cost of more syntactical complexity [5].

Holding an item and peeling off a copy with the pen is a good example of a pen+touch technique that ostensibly violates the principle that *the pen writes and touch manipulates* because here the pen drags the copy. However, those principles apply to pen *or* touch as unimodal inputs. The transaction is consistent with the principle that we use to guide multimodal gestures: *pen + touch yields new tools.* The gesture feels natural and effective because it is grounded in people's naturally occurring behaviors with physical paper, such as holding an item and making pen strokes in reference to it (Guiard [13], and design study B3).

Ruler: Using an Object as a Straightedge

The user can employ an object as a straightedge by holding down the object with the thumb and index finger, like the framing gesture observed in our design study (B4, Fig. 2b). The user can then stroke along the object with the pen stroke constrained to its border, as inspired by observation B8.

When two fingers come down on an object, an animation starts that increases the transparency of the item and adds a dotted line around its border. Informal test users who tried early versions of the system suggested that items should become mostly-transparent in this manner so that users can see the relationship of the straightedge to the underlying surround of other objects and strokes.

The *Ruler* uses content as its own tool, introducing a subtle duality between content and tool into our system. We could add a dedicated ruler object, but this would necessitate acquiring the ruler before drawing a straightedge. Using an object as its own straightedge also supports other sketching techniques, such as adding a drop shadow or "outer glow" to an item. To be clear, however, we are not arguing that the system should not include a dedicated ruler tool; rather, we are arguing that using an object as its own straightedge has some interesting interaction design properties that a dedicated ruler tool does not.

Composition of the Straightedge with Cutting

To illustrate how our interaction design allows multiple atomic interactions to be composed together into higher-level input phrases, we implemented X-acto cutting along straightedges (Fig. 8). To do this, the user first finger-tap-selects the photo to cut, and then uses the two-finger thumb and index finger grasp on an overlapping item to establish a straightedge. The user can then stroke the pen across the selected photo along the straightedge to cut it in a straight line. The user can even cut around the corner of the masking

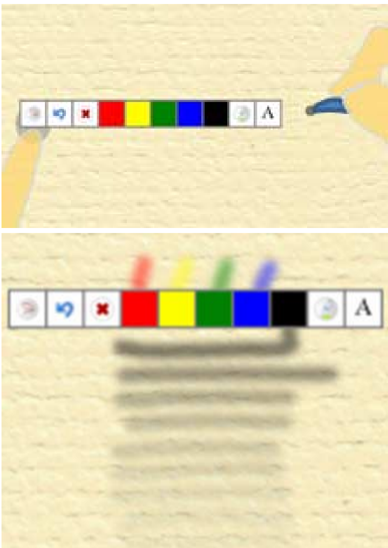

figure 9. (top) Calling up additional commands via the Finger Shadow. **(bottom)** Finger painting from the color pots. Each successive finger paint stroke is progressively fainter, until "all the ink is rubbed off" of the user's finger. In addition to mimicking physical media such as charcoal sketching, this facilitates a lightweight transition out of the finger-painting mode.

object (Fig. 7). This composition of techniques is possible because *tapping* provides a way to define a "collective" scope of more than one object, and because *holding* phrases together multiple pen and touch inputs into a single "chunk" that the system can interpret as a unitary command [5].

Brush & Stamp: Using an Object as Its Own Tool

Our system includes a couple of techniques, which we call stamping and brushing, for producing creative effects by letting the user employ content on the page as its own tool. Due to space constraints, these are illustrated only in the accompanying video figure.

Finger Shadow & Finger Painting

We have also experimented with the *Finger Shadow* (Fig. 9, top), which uses a finger-tap on the canvas to bring up in-place commands (see video figure). This accesses additional commands, including color pots.

In early demonstrations of our system we observed that most people expect touching the color pots to enable smearing of colors with the finger. To probe this issue further, in this context we intentionally break our design rule for unimodal inputs that *the pen writes and touch manipulates.* To prevent finger painting from becoming a heavyweight mode, each successive finger paint stroke appears fainter until "all the ink has been rubbed off" the user's finger (Fig. 9, bottom). This also mimics how artists naturally work with some physical media, such as charcoal sketching.

With finger painting, we have to face a genuine design dilemma: when ink remains on one's finger, should it be possible to finger paint on top of objects? In our system, the answer is no. Touching an object to select or move it (or touching other controls, such as the bezel menu, or flipping pages) always takes precedence over finger painting. Otherwise, a more explicit means to "stop" finger painting and return to the default behavior that *touch manipulates* would have to be introduced, and we did not wish to do so at this point.

Summary and Discussion

Let us reflect briefly on how we use touch to hold items, in combination with one or more pen strokes that act in reference to the item that the user is holding. We have seen that this approach enforces a strong notion of phrasing [5] in the resulting interface design: there are no persistent modes, but rather tool use is always tied to holding an object with the nonpreferred hand. There is no possibility of getting stuck in a tool mode with the pen, nor is there ever a question of how to switch back to the default action of drawing ink strokes on the page. Once the user releases objects, drawing on the page with the pen *always* leaves ink strokes.

This approach is akin to nonpreferred-hand mode switching [16,22], but requires no physical button and thus readily scales to a plurality of modes [25]. It also offers an additional advantage: object selection is integrated with the mode switch itself when the user's hand touches down on an object on the display. Hence a unique design property of pen+touch is the facility with which it can support modes and tools specific to particular objects on the screen.

Informal Usability Evaluation

In addition to our design study, and observations from informal demonstrations, we also conducted a usability evaluation with seven professional designers. Due to space constraints we note only a few observations here.

Users found the core operations of writing and manipulating objects via touch to be largely obvious. Dragging new items from the top bezel was particularly well received. The general pattern of our combined pen + touch gestures resonated with users, but users did have to be told how to articulate the gestures when first encountering them. After the study, one designer commented that "the way it works is just like the way I already work in my notebook!" Another user commented that "I wouldn't have guessed the gestures work that way, but once I tried it, it felt pretty natural."

Conclusion and Future Work

Our work is motivated by a desire to extend use of pen and touch, including their use in concert, in order to enable users to take better advantage of each. Inherent in this statement is our belief that each has its own strengths and weaknesses, and that these are largely complementary. Yet, what should be the division of labor between pen and touch input when both are available? What is the design vocabulary afforded by combining pen and touch? What are some potential limitations? Such questions provide a sense of the space that we set out to explore in this study. Our purpose has not been to advocate any particular design; rather, to conduct a quick informal probe that yields insights into how interface designs can effectively combine pen and touch input.

The fluency, flow and engagement afforded by our system (despite the relatively primitive and sometimes arbitrary tools tested), is an encouraging existence proof that the path which we set out on is worth probing more deeply and systematically. While we do not claim to have created a revolutionary new graphic design package, our system offers many examples and insights that should be of considerable interest to the other researchers and designers who are now actively exploring and seeking to characterize the design space of pen+touch techniques.

We have observed that users, even with limited experience, are able to engage with the system. This demonstrates that pen + touch input on a direct display can afford a design where the pen writes, touch manipulates, and the combination of pen and touch yields new tools. A number of simple pen or touch actions (holding, tapping, dragging, crossing) can be composed in a phrase structure, delineated by the muscular tension of holding an object or state, that affords a rich design space of interactions. All this can be done without using a button on the pen, keys for non-preferred hand mode switching [16,22], or permanently visible UI that detracts from the user's focus on the workspace [20].

We have seen that if *pen+touch yields new tools*, implicitly this means that in some contexts we will violate the first two principles: the pen does not *always* write, nor does touch *always* manipulate. Our explorations have convinced us that if a system strictly limits itself so that the pen ONLY writes, and touch ONLY manipulates, this leads to a simple and consistent but artificially crippled system. Hence the design trade-offs we explore are intricate for good reason. Furthermore, we could have presented one or two of the techniques in this paper, with formal studies, but chose not to do so because in our view the primary significance of pen + touch accrues from the rich design space of compound gestures that is afforded, rather than from the individual, atomic techniques themselves. For example, for any one of our simultaneous pen+touch techniques, one could likely devise a pure multi-touch technique that works nearly as well. But we believe it would be difficult to design a system supporting the same breadth of operations without resorting to increased modes, on-screen widgets, and a larger set of arbitrary and/or more complex touch gestures.

In addition to proving that our approach scales from a demo application to a full-blown one, it also remains to demonstrate how these techniques serve to enhance the effectiveness and user experience of other applications, such as working with a spread-sheet, for example. We believe that they can and will, but that is a long-shot from actually doing so. As well, despite our hopes for these techniques on other form factors, the fact remains that we have not yet done those tests. There is still work to do.

References

[1] Agarawala, A., Balakrishnan, R. *Keepin' It Real: Pushing the Desktop Metaphor with Physics, Piles and the Pen*. CHI 2006.
[2] Bondarenko, O. and Janssen, R. *Documents at Hand: Learning from Paper to Improve Digital Technologies*. CHI 2005.
[3] Brandl, P., Forlines, C., Wigdor, D., Haller, M. and Shen, C. *Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces*. AVI '08 *Conf. on Advanced Visual interfaces*. p. 154-61.
[4] Buxton, B., *Sketching User Experiences: Getting the Design Right and the Right Design*. 2007, Morgan Kaufman.
[5] Buxton, W. *Chunking and Phrasing and the Design of Human-Computer Dialogues*. IFIP Information Processing `86.
[6] Buxton, W. *A three-state model of graphical input*. Proc. INTERACT'90. p. 449-56.
[7] Cohen, P. *The role of natural language in a multimodal interface*. UIST'92. p. 143-49.
[8] Cohen, P., et al. *QuickSet: Multimodal Interaction for Distributed Applications*. ACM Multimedial 97.
[9] Engelhard, L. *Native Dual Mode Digitizers: Supporting Pen, Touch and Multi-Touch Inputs in One Device on any LCD*. Society for Information Display SID 08 Digest.
[10] Frisch, M., Heydekorn, J. and Dachselt, R. *Investigating Multi-Touch and Pen Gestures for Diagram Editing on Interactive Surfaces.* ITS '09 Conf. on Interactive Tabletops and Surfaces.
[11] Greenberg, S., Boyle, M. and LaBerge, J., *PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal.* Personal Technologies, 1999. **3**(1): p. 54-64.
[12] Greenberg, S. and Buxton, B. *Usability evaluation considered harmful (some of the time)*. CHI '08. p. 111-20.
[13] Guiard, Y., *Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model.* The Journal of Motor Behavior, 1987. **19**(4): p. 486-517.
[14] Hartmann, B. *Pictionaire: Supporting Collaborative Design Work by Integrating Physical and Digital Artifacts*. CSCW 2010.
[15] Hinckley, K., et al. *Codex: a dual screen tablet computer*. CHI '09. p. 1933-42.
[16] Hinckley, K., Guimbretiere, F., Baudisch, P., Sarin, R., Agrawala, M. and Cutrell, E. *The Springboard: Multiple Modes in One Spring-Loaded Control*. CHI 2006.
[17] Hinckley, K., Pierce, J., Sinclair, M. and Horvitz, E. *Sensing Techniques for Mobile Interaction*. UIST 2000. p. 91-100.
[18] Hinckley, K., et al. *InkSeine: In Situ Search for Active Note Taking*. CHI 2007. p. 251-60.
[19] Kidd, A. *The marks are on the knowledge worker*. CHI '94.
[20] Kurtenbach, G. and Buxton, W. *Issues in Combining Marking and Direct Manipulation Techniques*. UIST'91. p. 137-44.
[21] Kurtenbach, G., Fitzmaurice, G., Baudel, T. and Buxton, B. *The Design of a GUI Paradigm based on Tablets, Two-hands, and Transparency*. CHI'97, p. 35-42.
[22] Li, Y., Hinckley, K., Guan, Z., Landay, J., *Experimental Analysis of Mode Switching Techniques in Pen-based User Interfaces*. CHI 2005. p. 461-70.
[23] Marshall, C. and Bly, S. *Saving and Using Encountered Information: Implications for Electronic Periodicals*. CHI 2005.
[24] Roth, V. and Turner, T. *Bezel swipe: conflict-free scrolling and multiple selection on mobile touch screen devices*. CHI 2009.
[25] Ruiz, J. and Lank, E. *A study on the scalability of non-preferred hand mode manipulation*. ICMI '07 *Conf. on Multimodal interfaces*. p. 170-77.
[26] Terrenghi, L., Kirk, D., Sellen, A. and Izadi, S. *Affordances for manipulation of physical versus digital media on interactive surfaces*. CHI '07. p. 1157-66.
[27] Wu, M., et al., *Gesture Registration, Relaxation, and Reuse for Multi-Point Direct-Touch Surfaces*, *IEEE Intl Workshop on Horizontal Interactive Human-Computer Systems*. 2006.
[28] Yee, K.-P. *Two-Handed Interaction on a Tablet Display*. CHI 2004.