

Living with a Tabletop: Analysis and Observations of Long Term Office Use of a Multi-Touch Table

Daniel Wigdor, Gerald Penn, Kathy Ryall, Alan Esenther, Chia Shen

TR2007-076 December 2007

Abstract

Multi-touch tabletops have been the focus of significant recent study but, to date, few devices have moved from prototype to installed use. In this paper, we present observation and analysis of a subject who has used a direct-touch tabletop as his primary computing environment for the past 13 months, driving all manner of applications in a standard MS Windows environment. We present the results of three research instruments: a structured interview with the user, an analysis of touch and click locations when operating in desktop and tabletop modes over several days, and linguistic analysis of email composition over several months. From the product of these instruments we then report on several open avenues for research including physical parameters, hardware limitations, touch vs. click in the WIMP, and text entry techniques.

Tabletop 2007

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Living with a Tabletop: Analysis and Observations of Long Term Office Use of a Multi-Touch Table

Daniel Wigdor^{1,2}, Gerald Penn², Kathy Ryall¹, Alan Esenther¹ Chia Shen¹

¹*Mitsubishi Electric
Research Labs*

{ryall | esenther | shen} @ merl.com

²*Department of Computer Science,
University of Toronto*

dwigdor@dgp.toronto.edu, gpenn@cs.toronto.edu

Abstract

Multi-touch tabletops have been the focus of significant recent study but, to date, few devices have moved from prototype to installed use. In this paper, we present observation and analysis of a subject who has used a direct-touch tabletop as his primary computing environment for the past 13 months, driving all manner of applications in a standard MS Windows environment. We present the results of three research instruments: a structured interview with the user, an analysis of touch and click locations when operating in desktop and tabletop modes over several days, and linguistic analysis of email composition over several months. From the product of these instruments we then report on several open avenues for research, including physical parameters, hardware limitations, touch vs. click in the WIMP, and text entry techniques.

1. Introduction

Horizontal, direct-touch tabletops, which overlay large display and input surfaces, have recently been the focus of much study. Although a great number of experiments have been conducted which examine various aspects of tabletops, the majority of these experiments are conducted in a lab or similar setting, and require participants to perform some task over the course of several minutes or hours ([4][6][7][9][10][12][14]). Although scientifically valid for answering specific research questions, these efforts have limited abilities to predict patterns and desires for *long term* users of direct-touch tabletops.

In the present work, we describe the results of the study of an executive who has been using a direct-touch tabletop in place of an office computer for the last 13 months. Because of the length, setting, and tasks performed by this user, a great deal of “in the wild” experience is reflected in his responses. In addition, we perform a pair of analyses in order to learn more about his use of the table: first, we report the locations and frequencies of touch events, and compare it to logs of his use of a traditional pointing device, in order to extend and validate previous results suggesting that touch table use differs from mouse use in this measure. Second, we report the results of a computational linguistic analysis of email messages

sent over the 13 month period, comparing those composed on the tabletop and those typed on a regular keyboard.

A touch table as a primary office system is outside its typically described use. It is our hope, however, that the insights gleaned by studying this user, who has chosen to use the table in this way for his work, will be helpful to the community. It is our aim, in conducting this research, to inform the design of general problems, rather than those encountered only in this type of use.

1.1 Participant

The participant, *AB*, is a marketing executive at a local research lab. The tasks performed on the table are every day office tasks, and are limited to common applications – very little custom software is included in his setup (Figure 1). *AB*’s use of the tabletop is motivated primarily by his work: his role is the marketing and sale of the touch table. The system driving the table is *AB*’s laptop, which he also uses on the road and at home, not connected to a touch table. This pairing of input devices has allowed us to perform some simple comparative statistics on table and desktop use.



Figure 1. Our participant working at the touch table in his office. The table is his primary computer for everyday office tasks.

1.2 Research Instruments

In order to gain insight in to AB's usage of his touch table, we employed three research instruments. We now describe each in turn.

1.2.1 Mouse vs. touch use analysis

As we have described, AB uses a touch table for everyday computing tasks while in the office. While travelling on business or working at home, he uses traditional interfaces to control the same computer. In order to compare his pointing activity, we instrumented his laptop computer with two pieces of logging software: one to log the location of mouse clicks, and one to record touches on the touch table. This software recorded all such events over a three week period.

1.2.2 Email linguistic analysis

When AB composes email on the touch table, he uses an on-screen soft keyboard, while email written away from the office is typed with a traditional keyboard. In order to examine differences in email composition, we performed a comparative linguistic analysis on all of AB's outgoing email sent during the 13 month period he used his table as his primary office computer.

1.2.3 Interview

Finally, we conducted an extensive interview with AB, in which we sought to understand how his usage of the tabletop office system has evolved over the past 13 months.

1.3 Limitations

There are several limitations to the results described in this paper, each of which must be considered before generalizing our results to other designs. Despite these limitations, however, we are confident that the results described in the present work will be of use to researchers and designers.

First, as we have described, the present work describes the observation and analysis of a single user. Although AB has been using the touch table for every day computing for an extended period of time, this paper described only his experiences. Second, because we will be examining AB's use of a touch table for every day office tasks, nearly all of the experiences we will report on will be of single user applications. And, finally, because the tabletop is setup to operate as a "normal" desktop system, most of the interaction with the system was done by mapping the tabletop input to a single point, emulating a mouse.

Despite these limitations, many of the results we will describe offer useful insights for researchers and designers of interactive tabletop systems.

2. Related work

Although in its infancy, a number of researchers have reported results in the tabletop domain which are relevant to the present work. Generally, we divide these in to two categories: those which relate to touch locations, to inform on our mouse versus touch analysis, and those which report results from observation and contextual inquiry, and therefore relate to our interview. We are unaware of any use of linguistics for the comparative analysis of text entry devices.

2.1 Touch Location

Several results describe factors which might influence touch location. Scott and various co-authors have described issues in territoriality (summarised in [12]) which suggest that individuals working in groups will establish a personal working area directly in front of them. This is confirmed in by Ryall et al. who examined the interaction of group and table size on performance of tasks around a tabletop, and found that group members tended to focus their interaction (touches) directly in front of them. Finally, Ringel Morris et al. describe the results of a comparative study, which found that, when given the choice, users prefer UI that each group member have private copies of controls close to them on the table, rather than centralised, shared controls.

Although compelling, each of these previous works does not directly inform on the present work, since each examined the use of a tabletop in a group, rather than for private use by an individual.

2.2 Observation and Contextual Inquiry

In [12], Scott reported observations of a group of individuals interacting around a traditional table, describing the portioning of the workplace for individual and group tasks. In [9], Rogers et al. describe the observed use of finger talk, whereby users would gesture for both manipulation and conversation. Ryall et al. reported on observations of tabletop users in a multitude of environments, offering several recommendations and guidelines [11]. Both Tang et al. [13] and later Kruger et al. [6] described the use of orientation and spatial positioning of on-table objects for the delimiting of personal and group working spaces.

Although informative, each of the research efforts described above has been limited in their scope due to the constraints of time and scope. In this paper, we describe the results of interview and analysis of work performed on the tabletop over more than a year, and by a user who has adopted the tabletop as his primary workspace.

3. Environment

In this section, we describe the physical environment, the tasks being routinely performed, and the software and hardware employed.

3.1 Physical Space

AB's office features a regular desk, complete with traditional computer apparatus (monitor, keyboard, and mouse), as well as a large *DiamondTouch* touch table [2], with a touch surface measuring approximately 87x64cm. The keyboard, mouse, and monitor are attached to the same system as the tabletop, with the monitor displaying the same content. The desktop setup is arranged on a desk behind the tabletop, such that the monitor cannot be seen while working at the table.

The DiamondTouch is mounted on a sloped surface, creating a drafting table style of interaction. As shown in Figure 2, the height of this was set so as to allow AB to use it while standing (left) or seated (right).



Figure 2. The touch table is oriented as a drafting table. Its height is set for both standing and sitting.

3.2 Application software

Although special circumstances may arise, AB uses the tabletop primarily for everyday office software. The list of regularly used applications includes word processing (MS Word), spreadsheets (MS Excel), presentation software (MS PowerPoint), email (Mozilla Thunderbird), and the web (Mozilla Firefox).

3.3 Touch table software

All of the application software used by AB is unmodified for use on a tabletop. To enable direct-touch input to drive these applications, two tools are required: mouse emulation, and a soft keyboard.

3.3.1 Mouse input

Because AB's primary use of the DiamondTouch is in applications designed for the mouse, he uses a software utility, *DTMouse*, which maps multi-touch input to mouse events [3]. Although input is ultimately delivered to a single pixel when it is converted to mouse actions, a set of gestures on the tabletop is used to differentiate between mouse actions (see Table 1).

Table 1. The mapping of multi-touch gestures to single-pixel mouse events in *DTMouse*.

<i>Mouse Event</i>	<i>Touch Action</i>
Left mouse: Depress Drag Release	Single finger: Touch table Slide along table Lift from table
Right mouse: Depress Drag Release	Touch table with finger, then: Tap second finger Slide first finger Lift first finger from table
Middle mouse: Depress Drag Release	Touch table with finger, then: Double-Tap second finger Slide first finger Lift first finger from table
Wheel Wheel up Wheel down	Place closed-fist on table, then: Slide fist up Slide fist down

In addition to this mapping, *DTMouse* can be used in a *precise selection* mode, in which the user places any two fingers on the table simultaneously. The cursor then relocates to the centre point between those fingers, without sending any button events. Touching a third finger to the table depresses the left mouse button, and removing it from the table releases the button. In effect, this mode allows for two enhancements over the simpler operation described above: first, it allows the pointer position to be set without depressing any buttons; second, it allows the user to view the position of the cursor more precisely, since the finger does not occlude the display [3].

3.3.2 Soft keyboard

To enable text entry from the DiamondTouch, AB makes use of the soft keyboard built-in to Microsoft Windows. The keyboard generates text when the system pointer is used to click soft buttons arranged in a QWERTY layout. Interaction with the keyboard is facilitated using *DTMouse* to simulate mouse clicks. As such, speed of text entry is limited by the need to have only one finger in contact with the keyboard at a time, similar to tapping text entry with a stylus.

In the following sections, we sought to gain insight into AB's use of his tabletop system for every day office tasks. In order to do this, we employed three research instruments. The first two, a touch location analysis and linguistic analysis of email, relied on a comparison between his use of a laptop while away from the office and the touch table while in his office. In this section, we describe these instruments and discuss their results.

4. Touch vs. mouse logs

As previously described, the DiamondTouch table in AB's office is driven from his laptop computer. When not in the office, he uses the same computer with its built-in keyboard and track-pad. As we have previously described, a number of researchers have reported results which suggest that, when working in groups on a tabletop, individuals tend to organize their work spaces such that their efforts are focused in an area immediately in front of them, perhaps even closer than arm's reach might require [7],[10],[12]. None of these researchers examined whether this pattern was repeated for individuals, allowing for possible confound with fatigue [14].

We note that AB's touch table is sufficiently small that all points are within arms reach while it is in use. We were interested to learn whether AB, who has had ample experience to perfect any strategies that might be employed to mitigate fatigue, would follow a usage pattern similar to that described in previous work, which attributed it to effects of working in a group.

4.1 Method

In order to perform this comparison, we instrumented AB's laptop with logging software which recorded clicks and touch locations for a three week period. In both tabletop and laptop mode, AB's screen is set to 1024 x 768 pixels. The software was not capable of detecting the target applications for input.

Although it is usually the case that AB uses laptop controls when away from the office, and the touch table while at the office, our study period included several days that the laptop controls were used in the office while the touch table was unavailable.

4.2 Results

The logs contained 7034 mouse clicks and 47,972 touch events. This asymmetry was expected, since AB uses the mouse primarily while out of the office, and because multiple touch events are required to generate many types of mouse actions.

To simplify our analysis, touch and click locations were organized in to 10 x 10 pixel bins, and the results normalised. The results are shown Figure 3.




Figure 3. Click locations (left) and touch locations (right), organized in to bins. Red level indicates frequency.

4.3 Analysis

Initially, the data are quite striking: it is very apparent that a greater number of touch events is being generated towards the bottom of the screen when using the tabletop than when using the laptop. What became apparent upon further analysis, however, is that the large cluster of events in this region was due almost entirely to the placement of the soft keyboard. Because our logging method did not provide us with a means classifying touch targets by destinations, we are unable to separate for analysis those touches used for typing from the rest of the data. As a result, we are unable to attribute the entirety of this cluster to the keyboard. What is apparent in the data, however, is that if this region is removed entirely from both sets, the remaining areas have nearly identical frequencies.

This result appears to support the hypothesis that the workspace partitioning described in the previous work ([7],[10],[12]) is indeed due to the presence of additional group members, and refutes our suggestion that this might be equally attributable to fatigue.

Although promising to inform on these issues, two significant factors limit our ability to generalise the current result. First, as previously described, the study is limited to a single participant in an uncontrolled environment. Second, an equally reasonable explanation of the difference between our single user and the results reported a user working in a group has to do with the software executing on the table during our logging. Because it was written for use with the mouse, interface features are situated without regard for minimizing reach. A quick glance at offers some support for this explanation: dense clicks can be seen at the top (menu bar), bottom (tool bar), and right side of the screen (scroll bar). Although AB might wish to organize the desktop to better suit a tabletop, he is limited by the available software.

Despite these limitations and potential confounds, the results of the study are still compelling. They suggest the confirmation of past results, and demonstrate that a tabletop is sufficient as an input device for everyday office tasks.

5. Email Linguistic Analysis

Our second research instrument is a linguistic analysis of email sent by AB over the 13 month period that the touch table was used as his primary office computer. Because AB has chosen to use a soft keyboard as his input device for text entry while working on the table, we suspected that a comparative analysis of messages composed while working on that device and those composed on the laptop keyboard would yield significant differences. Such differences might arise due to difficulties with using the table and/or soft keyboard, such as during the composition of long email messages, or due to a user such as AB's belief that touch tables are more appropriate for sending certain kinds of emails.

5.1 Method

We retrieved from AB a log of outgoing messages for the study period (13 months). Messages were separated from each other, and their headers were removed using a MIME parser. From these messages, we removed those which contained repeated, standardized text (such as might be sent in reply to a sales enquiry). We also removed from analysis quoted text from each message, to ensure that the analyzed text in each message was only that which was newly composed. We then classified each message as having been composed using either the soft keyboard on the touch table, or the laptop keyboard. This was usually possible because AB's email client is configured to attach a particular footer to messages composed using the touch table.

A large number of features were computed for each message, aided in part by the Brill Tagger [1]. These measures included rates (per 1000 tokens) of:

- 1) *first-person pronominals*
- 2) *second-person pronominals*
- 3) *third-person pronominals*
- 4) *conjunct words*
- 5) *present-tense verbs*
- 6) *past-tense verbs*
- 7) *"private" verbs*
- 8) *"public" verbs*
- 9) *verbs of saying*
- 10) *nouns*
- 11) *nouns that start with 'q' or 'Q'*
- 12) *adverbs*
- 13) *"that" complementizers*
- 14) *non-existential "there"*
- 15) *split infinitives, and*
- 16) *technical vocabulary*
- 17) *word-instance/sentence ratio*
- 18) *average word length*
- 19) *word type/word instance ratio*
- 20) *total message length in characters*
- 21) *total message length in whitespace-delimited strings, and*
- 22) *total message length in lines.*

The last three being determined by the Unix *wc* command. *Word instance* refers here to a single instance of a word, whereas *word type* refers to a particular kind of word. *The meeting takes place on the 31st*, for example has seven word instances, but six word types, since the type *the* occurs in two instances.

With the exception of the three total message length features, these features are typical in computational linguistics for detecting the genre of a text document. Genre refers to a mode of language use with conventions that native speakers of the language implicitly agree upon. Instructions on tax forms, romance novels and personal letters, for example, are all genres between which these features are known to vary considerably. It is also well-known that these features are sensitive to differences between spoken and written language. To our knowledge, no previous study of their variation across input devices or input modalities has ever been conducted.

5.2 Results

We were able to classify all but 83 of the original 2376 email messages as having been written on either the tabletop or the computer keyboard. Some mailers attach quotations in such a way that the distinguishing footer that we used for classification was ambiguously attached either to the new message or to a quote of an earlier message. This was the case for all 83 of the emails that we forced discarded. We removed all of the quoted text contained in the emails, which removed a further 16 messages from consideration (messages which were forwarded email without any additional text). This left us with 1124 messages composed on the tabletop and 1153 messages written with the keyboard.

We performed a comparison of means between the two input devices for each of the word types. Of the 22 features considered, only six were found to have significantly different values on the tabletop keyboard.

The use of second-person pronominals (eg: *you*, *your*) was significantly more frequent for email composed on the tabletop than for email composed on the keyboard ($F_{1,2275} = 7.38$, $p = 0.007$, with mean frequencies of 31.3 and 27.8 per 1000 words for tabletop and keyboard messages respectively). Similarly, the frequency of technical vocabulary was significantly higher on the tabletop ($F_{1,2275} = 4.62$, $p = 0.032$, with means of 23.2 vs. 20.5 for tabletop and keyboard input respectively). On the other hand, the use of *that* complementizers was more frequent in keyboard email than in touch table email ($F_{1,2275} = 4.95$, $p = 0.026$, 10.4 vs. 11.9 touch table and keyboard). Total message length in both characters ($p = 0.008$) and whitespace-delimited strings ($p = 0.009$) was significantly longer in email composed on the tabletop, but total message length in lines, which is simply a count of newline characters, was significantly less ($p = 0.000$).

5.3 Analysis

Of particular interest is that the total message length was significantly different in every dimension, but no significant difference was found in either the number of characters per word, or the number of tokens (words) per sentence. It is clear that, although it has been previously reported that soft keyboards are insufficient for text entry on a touch table [11], AB shows no sign of being incapacitated by the table in composing his messages. Quite to the contrary, his messages are longer. Also of interest is that both frequent use of second-person pronominals and infrequent use of *that* complementizers are associated with less formal genres of written language, although several other features with the same association showed no significant difference. Frequent use of technical vocabulary, while also explainable by variation in topic, is associated with more formal genres, on the other hand. Topic and genre are not the same, although not statistically independent either.

Once again, we are unable to generalise the results of this to all users. In addition to the reasons previously stated, the results here point to a similarity in genre and length statistics between tabletop and physical keyboard use, but it is possible that AB's language use might happen to fall in to a genre which is particularly robust to differences introduced by a change in input methods, where others might not be.

Despite these limitations, the results point to the reasonable conclusion that the use of a soft keyboard has not limited AB's ability to perform everyday office tasks. Indeed, in most respects, the style of his language, seems more or less completely unaffected by the input device, and the lengths of his messages are longer.

6. Interview

In this section, we present the results of our interview with AB. As often as possible, the aim of the interview was to glean insights that are generally applicable to direct-touch tabletops, rather than the particular hardware, software, or domain applications that the participant may have chosen.

6.1 Dual use

AB's tabletop system is the only horizontal surface at that side of his office. He has reported that it is often called upon to play a dual role. In addition to its traditional use as a computing device, it is also often used as furniture, serving as both coffee table and conference table for small meetings, in extreme cases being deactivated so-as not to project imagery on to objects on the table. In fact, AB reported that the angle at which his table was placed was chosen because it was the maximum possible angle before objects would begin to slide across its surface.

Because the DiamondTouch table is debris tolerant, objects placed on the table are not processed as touches by the system [2], making it ideal for this dual use. Designers of systems based on other technologies, however, must consider this dual use. Approaches might include the occlusion of a shelf in the physical design of the table, and mechanisms to rapidly disable arbitrary portions of the touch surface.

6.2 Ergonomic considerations

Several ergonomic factors were considered in the design of the setup in AB's office for the tabletop.

First, the display was placed on a slant, making it resemble a drafting table. This was done to allow both easier reach and easier viewing of pixels at the top of the display. Interestingly, the angle of the display was set such that it was as steep as possible, while ensuring that drinks placed on the table (see "dual use", above) would not slide down its surface.

Second, the height of the table was set so that it could be operated while either sitting or standing.

Third, AB had his choice of two different DiamondTouch tables, one with a diagonal measurement of 81 cm and another measuring 107 cm. Despite its use as a single-user desk, he opted for the larger table. He explained that, although this would mean having to reach farther for targets at the top of the screen, he preferred the larger field of view provided by the larger device. This, despite that the display on the DiamondTouch table is provided by an overhead projector, and so display resolution is independent of table size.

AB reports that he does not experience arm fatigue while working with the large touch surface, in contrast to previous study results which suggest otherwise [14]. This finding suggests that the field of view / fatigue trade-off is heavily skewed towards field of view.

The final ergonomic consideration is that, because the large touch surface provides a large display area, AB reports that privacy is sometimes an issue. Therefore, on some occasions he elects to work away from the touch table so as to limit this exposure.

6.3 Space management

As has previously been reported in the domain of group work, it is essential that designers provide mechanisms for easy repositioning of screen content to allow for dynamic space management by users [7], [13]. Previously, however, this was reported as necessary to facilitate the definition of shared and private spaces and to allow the use of shared resources. AB pointed out to us, however, the need for space management tools is also required to prevent the occlusion of one object while working on another.

6.4 Transition and expertise

As might be expected, AB's transition to the touch table for everyday use was a difficult one. It was initially mired by the transition to direct-touch and familiarity with the text input device. Despite these difficulties, he has persisted, and is now what can reasonably be described as an expert user (as evidence, recall that our touch analysis data, which covered only 3 weeks of use, included nearly 50,000 touch events). As such, many of the issues described in the following sections may relate to novices, experts, or both.

6.5 Bimanual interaction

Although AB describes himself as "severely right handed," he has consciously trained himself to perform touching the table with both his right and left hands. As we have previously described, the DTMouse software employed on AB's tabletop converts multi-touch input to mouse events. As such, each input event is, in effect, unimanual. AB, therefore, uses two hands in order to perform tasks which could otherwise be performed using just one hand. These include selecting objects using the closer hand and performing faster drag operations by starting the drag with one hand and completing it with the other.

In both cases, bimanual actions are utilised to save time through spatial rather than role specialisation. This suggests that, when designing interactions for the non-dominant hand, symmetric roles should be considered along side traditional, asymmetric roles [5].

6.6 Touch precision

Previous results explain that selection of small targets using a direct-touch interface is made difficult by two factors: the occlusion of the target while pointing, and the arbitrary mapping of large touch area to a more precise election point [8]. As has also been reported, this problem is exacerbated on horizontal surfaces, since near and distant touches present different fingertip shapes to the input device [4] (mitigated somewhat in AB's case by the slanting of his table). As we have described, the DT Mouse software provides a high-precision mode to aid with selection of small targets.

Perhaps surprisingly, AB reports that, after extensive experience, he no longer finds it necessary to use the high-precision mode to select small UI components, including window borders for resizing, which are only 4 pixels wide. This suggests that the issues with precise selection described in the literature may well be overcome with additional practice. Interestingly, AB *does* persist in using the precise selection mode when selecting blocks of text.

The distinction between the UI widget and the text-selection cases is the availability of visual feedback:

the effects of selecting UI elements, such as the resizable border of a window, are verifiable by viewing areas of the screen not occluded by the finger, whereas precise selection of text is not. This suggests that designers seeking to overcome the issue of precision might be well served by first addressing occlusion.

6.7 Gestures and direct-touch

As we have previously described, the DTMouse software maps multi-touch input to the tabletop to a simulated mouse device. Although the simulated device operates on a single-point, the input to DTMouse does not: right and middle buttons are simulated with the tapping of a second finger on the table, and mouse wheel scroll events are generated by dragging a closed fist on the table surface.

Of interest to designers is AB's facile and immediate adoption of this mix of the tabletop as both an input and display device, without always interpreting touches in a direct-touch way, in contradiction of the results suggested by previous work by Potter et al. [8].

For example, to send a right-click to the table, AB first selects the point for the click by touching it with a finger (direct-touch). The right-click is then generated when a second finger, usually his thumb, is tapped on the table. Interestingly, the location of this tap is not an input – even if it is directly on-to a UI component, it will not generate a direct-touch event for that location. Even more striking is the indirect mapping of scrolling from touch to display: dragging a closed fist downward causes a simulated downward scroll of the mouse wheel. The consequence of this, in most applications, is that on-screen content moves *up*, in precisely the opposite direction of the gesture.

When these inconsistencies in the input metaphor were pointed out to AB, it was clear from his reaction that he had no trouble conceiving of the tabletop simultaneously as a direct and indirect device.

6.8 Text Entry

As AB put it, "a reasonable person would not use an on-screen keyboard". He pointed out to us that a number of solutions are available for text entry to a direct-touch table, including installation of a keyboard drawer. Despite these solutions, however, AB persists in his use of the soft keyboard for text entry. He attributes this decision to his desire to maintain the use of direct-touch. And, as we have seen, AB's use of the soft keyboard has had little to no impact on the content of his email messages, suggesting he is not allowing the limitations to impact is overall performance.

7. Conclusions and Future Work

In this paper, we have presented a case study of AB, a long-term user of a direct-touch multi-point touch table. We have presented the results of three research instruments, two of which revealed few differences in behaviour between tabletop and mouse/keyboard operation of a WIMP interface. A structured interview with AB provided several insights into direct-touch interfaces.

Based on the results of these instruments, several avenues for research remain open. First, it is important to remember that, because of the limited sample size and potential confounds, none of the results reported here can be safely generalised. Rather than provide authoritative results, our aim is to present insight in to potential research directions. Therefore, validation of our statistical results through larger studies is the first clear open research question.

Next, each of the sections of the interview suggest clear open questions. AB's dual use of his tabletop is supported in part by a combination of debris tolerance and projector muting. The design of software and of physical apparatus to further support this is an open question.

The interview also pointed to several ergonomic considerations: it is clear that AB prefers to angle his table towards him. However, this may limit it serving as a dual-use surface. Additionally, optimal height, and size, and mechanisms to support private work are all open questions.

The alternative use of bimanual input to minimize work, rather than to allow for more complex manipulation, is demonstrated aptly by our participant, as is the potential for learning to overcome the limitations in precision of direct-touch interfaces. Also of interest is AB's ease of mixing direct-touch and gesture/postures on the table. This suggests the need to closely examine the tradeoffs and design of these interactions to suite a particular task.

Finally, text entry is a clear open question. As AB points out, the use of a keyboard on the tabletop is one solution, but there may well exist others. Thoughtful study is required.

Given the multitude of open questions pointed to by this work, it is clear that it represents a first step, rather than a comprehensive result. Further study in all manner of tabletops for single users is warranted.

8. Acknowledgements

We offer thanks to Clifton Forlines for assistance with statistical analyses, and great thanks to AB, our participant.

9. References

- [1] Brill, Eric (1995). Transformation-Based Error-Driven Learning and Natural Language Processing: A Case Study in Part-of-Speech Tagging. *Computation Linguistics*, 21(4):543-565.
- [2] Dietz, P. and Leigh, D. (2001). DiamondTouch: a multi-user touch technology. *Proc. UIST*. p. 219-226.
- [3] Esenther, A., Ryall, K., 2006. Fluid DTMouse: Better Mouse Support for Touch-Based Interactions. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI 2006)*, p. 112- 115.
- [4] Forlines, C., Shen, C., Wigdor, D., Balakrishnan, R. (2007). Direct-touch vs. mouse input for tabletop displays. *Proceedings of CHI 2007*. p. 647-656
- [5] Guiard, Y. (1987). Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *The Journal of Motor Behavior* 19(4). p. 486-517.
- [6] Kruger, R., Carpendale, S., Scott, S. and Greenberg, S. (2003): 'How people use orientation on tables: comprehension, coordination and communication', *Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work*. Sanibel Island, Florida, USA, ACM Press, 2003, pp: 369-378.
- [7] Morris, M. R., Paepcke, A., Winograd, T., and Stamberger, J. (2006). TeamTag: exploring centralized versus replicated controls for co-located tabletop groupware. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. p. 1273 – 1282.
- [8] Potter, R., Weldon, L., and Shneiderman, B., 1988. Improving the accuracy of touch screens: an experimental evaluation of three strategies. *Proceedings of the ACM CHI*. p. 27-32.
- [9] Rogers, Y., Hazlewood, W., Blevis, E. and Lim, Y. (2004). Finger talk: collaborative decision-making using talk and fingertip interaction around a tabletop display. *Extended abstracts of CHI*. p. 1271-1274.
- [10] Ryall, K.; Forlines, C.; Shen, C.; Ringel-Morris, M. (2004) Exploring the Effects of Group Size and Table Size on Interactions with Tabletop Shared-Display Groupware. *Proceedings of CSCW*. p. 284-293.
- [11] Ryall, K., Morris, M., Everitt, K., Forlines, C., and Shen, C. (2006). Experiences with and observations of direct-touch tabletops. *Proceedings of IEEE TableTop the International Workshop on Horizontal Interactive Human Computer Systems*. pp. 89-96.
- [12] Scott, Stacey D. (2005). *Territoriality in Collaborative Tabletop Workspaces*. Ph.D. Dissertation, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada,
- [13] Tang, J.C. (1991). Findings from observational studies of collaborative work. *International Journal of man-Machine Studies*, 34. p. 143-160.
- [14] Wigdor, D., Leigh, D., Forlines, C., Shipman, S., Barnwell, J., Balakrishnan, R., Shen, C. (2006). Under the table interaction. *Proceedings of the 2006 ACM conference on User Interface Software and Technology (UIST 2006)*. p. 259-268