The performance of hand postures in front- and back-of-device interaction for mobile computing

Jacob O. Wobbrock[a,*], Brad A. Myers[b], Htet Htet Aung[c]

[a]*The Information School, University of Washington, Box 352840, Seattle, WA 98195, USA*
[b]*Human-Computer Interaction Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA*
[c]*Siemens Medical Solutions, Soarian Clinicals, 51 Valley Stream Parkway, Malvern, PA 19355, USA*

Abstract

Three studies of different mobile-device hand postures are presented. The first study measures the performance of postures in Fitts' law tasks using one and two hands, thumbs and index fingers, horizontal and vertical movements, and front- and back-of-device interaction. Results indicate that the index finger performs well on both the front *and* the back of the device, and that thumb performance on the front of the device is generally worse. Fitts' law models are created and serve as a basis for comparisons. The second study examines the orientation of shapes on the front and back of a mobile device. It shows that participants' expectations of visual feedback for finger movements on the back of a device *reverse* the direction of their finger movements to favor a "transparent device" orientation. The third study examines letter-like gestures made on the front and back of a device. It confirms the performance of the index finger on the front of the device, while showing limitations in the ability for the index finger on the back to perform complex gestures. Taken together, these results provide an empirical foundation upon which new mobile interaction designs can be based. A set of design implications and recommendations are given based directly on the findings presented.
© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

Mobile devices have been shrinking dramatically over recent years. Their small size has made them sufficiently portable to be used "on the go" while standing, walking, and riding. As a result of devices' nomadic use, it has become important to many designers to support mobile interaction without the need of a stylus. Examples include techniques for finger-driven touch screens (Karlson et al., 2005; Vogel and Baudisch, 2007), capacitive sensing keypads (Rekimoto et al., 2003), and text entry joysticks (Wobbrock et al., 2007), all of which allow mobile devices to be controlled by fingers or thumbs, not styli. But although designers have begun to refine finger- and thumb-based interactions, the shrinking of devices has made these designs even more challenging, and devising new successful devices and interactions is difficult. For example, designers are currently exploring the use of the *backs* of mobile devices for interaction (Hiraoka et al., 2003; Sugimoto and Hiroki, 2006; Wigdor et al., 2007; Wobbrock et al., 2007), which has potential benefits in that a user's hands do not obscure targets or suffer from poor pointing precision (Vogel and Baudisch, 2007). But will such interactions be successful? How will a finger perform on the back of a device when the hand to which it belongs is also holding the device? How does an index finger on the back of a device compare to the thumb on the front (Fig. 1)? Although extensive prototyping and experimentation can help answer these questions, a foundational empirical basis would provide insights into which designs are likely to be best, and why.

To this end, we present three studies of hand-posture performance with a mobile device. The first study examined one- and two-hand grips, thumbs and index fingers, front- and back-of-device interaction, and horizontal and vertical

---

*Corresponding author. Tel.: +1 206 616 2541; fax: +1 206 616 3152.
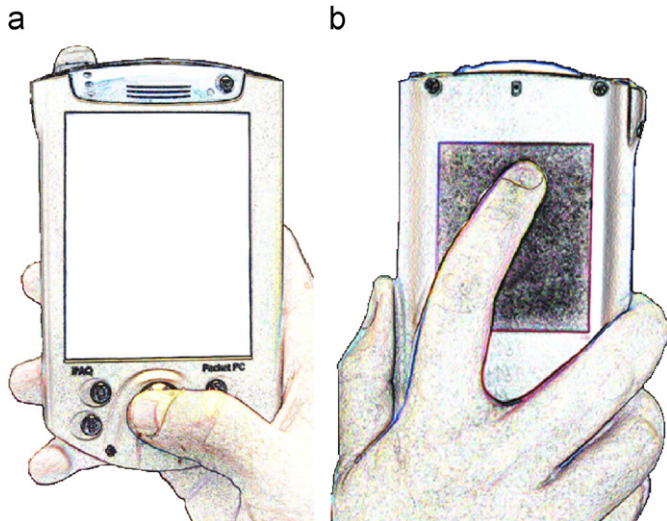E-mail address:* wobbrock@u.washington.edu (J.O. Wobbrock).

a   b

Fig. 1. (a) Typical one-hand interaction with a Pocket PC using the thumb on the front of the device. (b) A hypothetical device showing an index finger on a touchpad on the back. Among other things, the current study examines the performance of a finger on such a device.

movements in Fitts' reciprocal pointing tasks designed to elicit the throughput of these postures. An interesting finding was that the index finger on the back of a device outperforms the thumb on the front of a device in horizontal movements.

The second study was designed to investigate how people interpret orientation on the fronts and backs of mobile devices. In this study, the orientations of letter shapes were captured on the front and back of a device with and without visual feedback. We found that the mere *expectation* of visual feedback actually reverses the orientation with which users write letters on the back of a device. Specifically, their orientation shifts from a motor-driven one to a visually-driven one.

The third study examined gesture-making on the front and back of a mobile device. The study had participants enter characters from the EdgeWrite alphabet (Wobbrock et al., 2003) while holding the device with the hand postures used in the first study. Results confirmed the superior performance of the index finger on the front of the device, and showed that letters can indeed be made by an index finger on the *back* of a device. However, performance of the index finger on the back was relatively worse for complex gestures than it was in simple moving and pointing tasks from the first study, indicating the types of tasks for which a back-of-device index finger is best suited.

Taken together, the results from these studies serve as a basis for mobile interaction design. Designers can benefit from these findings when creating new input and interaction techniques involving one or two hands, thumbs or index fingers, front- or back-of-device interactions, and horizontal or vertical movements. It is our hope that these results will help guide and influence the creation of successful interaction techniques and a new generation of highly interactive mobile devices.

2. Related work

The current work is related to research in multiple areas of human–computer interaction: (1) studies of human appendages using Fitts' law; (2) studies of thumb- and finger-driven techniques for mobile device interaction; and (3) previous prototypes that have used back-of-device interaction techniques. Although there are relevant findings in each of these areas, they have not been brought together to understand human performance with handheld devices or back-of-device interaction techniques.

2.1. Fitts' law studies of human performance

2.1.1. Prior work on throughput calculations

Fitts' law is a robust human performance model that predicts target acquisition times in rapid aimed movements (Fitts, 1954). A rapid aimed movement is a closed-loop movement in which the performer can update and correct his motion. This contrasts with an open-loop movement (Schmidt et al., 1979), in which the performer executes an initial ballistic movement that determines the entire course of uncorrected action, like throwing a dart. Attempting to acquire targets on a screen while looking at the screen usually involves closed-loop movements that can be modeled using Fitts' law (Card et al., 1978). We therefore use Fitts' law as a basis for comparing hand postures in our first of three studies.

Because our related work and the current study rely heavily on it, we present the equation for Fitts' law here, using the popular Shannon formulation (MacKenzie, 1989, 1992):

$$MT = a + b\log_2\left(\frac{A}{W} + 1\right) \qquad (1)$$

In Eq. (1), *MT* is the predicted movement time, *A* is the distance (or amplitude) of movement, and *W* is the target size, a one-dimensional width. The *a* and *b* terms are empirically derived regression coefficients specific to a device and task. The $\log_2$ term is called the *index of difficulty* (*ID*), and is measured in bits. The greater the *ID*, the more difficult, and therefore slower, the predicted acquisition. Importantly, $1/b$ is the *index of performance* (*IP*), which is a task-independent throughput measure in bits/second. *IP* values can be compared across studies as a way of determining performance.

An important consideration in examining prior studies of throughput is whether or not the speed–accuracy tradeoff was normalized using Crossman's correction for *effective target width* (Crossman, 1957; Welford, 1968). This correction adjusts the *W* parameter in the *ID* calculation based on the dispersion of selection endpoints; that is, the accuracy of the performer. Owing to its information-theoretic basis, Fitts' law assumes an approximate 4% error rate. When this error rate is not observed, participants are being overly careful or not careful enough in their aimed movements. They are either under utilizing

or over utilizing the target width W. Crossman's correction adjusts W to enforce a 4% *post hoc* error rate so that throughputs are comparable.[1] It relies on the following calculation of W_e, the effective target width:

$$W_e = 4.133 \times SD_x \tag{2}$$

The SD_x term refers to the standard deviation of endpoint coordinates in the *x*-dimension for a one-dimensional horizontal movement task (SD_y would replace SD_x for a vertical movement task). A theoretical explanation of Crossman's correction is beyond the current scope; interested readers are directed to prior literature (Crossman, 1957; Welford, 1968; MacKenzie, 1992; Soukoreff and MacKenzie, 2004).

Crossman's correction then defines an *effective ID*, called ID_e, using W_e:

$$ID_e = \log_2\left(\frac{A}{W_e} + 1\right) \tag{3}$$

The result of using Eq. (3) in the equation for *MT* prediction (Eq. (1)) is that throughputs (*IP*) can be equitably compared even when participants perform at different points along the speed–accuracy continuum. Another result is that after the removal of clear outliers, error trials can be retained and not discarded (MacKenzie and Soukoreff, 2003; Soukoreff and MacKenzie, 2004). When comparing our throughput results in this paper to those of prior studies, it will be important to note whether ID_e or ID was used; if the latter, throughputs are generally inflated. In this paper, our throughput results for different hand postures are calculated using ID_e, although we include results using ID to facilitate comparisons with prior studies.

2.1.2. Prior studies of appendages using Fitts' law

Some prior studies that used Fitts' law to quantify the performance of human appendages are relevant to the current work. The study by Balakrishnan and MacKenzie (1997) compared throughputs of the index finger, wrist, and forearm, finding that the throughput of the unsupported index finger, at about 3.0 bits/s, was much less than previously reported by Langolf et al. (1976). The wrist and forearm were shown to have throughputs of about 4.1 bits/s.

Zhai et al. (1996) compared a glove, which did not employ fingers, to a *FingerBall*, which did employ fingers, in a six degree-of-freedom docking task. They found that the latter achieved much better throughput, supporting the popular hypothesis that involving fingers in input device manipulation leads to high performance.

Gokturk and Sibert (1999) examined this hypothesis directly by studying the index finger as a pointing device. They found pointing with the index finger to have a

throughput of 5.53 bits/s for horizontal movement and 4.97 bits/s for vertical movement. However, these values are probably based on uncorrected *ID*s and would probably be lower if normalized. Nonetheless, their finding that horizontal movement results in higher throughputs than vertical movement is a finding verified by the current work.

As part of developing a model for text entry speeds on mobile phones, Silfverberg et al. (2000) conducted a Fitts' law study of one-handed thumb tapping and two-handed index finger tapping on a mobile phone keypad. They found that the throughput of the index finger was about 23.1% higher than that of the thumb for this task.

Other work has shown the benefits of having the non-dominant hand set the frame of reference for interaction while the dominant hand manipulates details within that frame (Buxton and Myers, 1986; Guiard, 1987; MacKenzie and Guiard, 2001). These findings are supported by our results, which indicate that throughputs for two-handed postures, where one hand holds a device while another operates it, are generally better than throughputs for one-handed postures, where the same hand holds the device and operates it. While this may not be particularly surprising, a contribution of these findings is in providing the magnitudes of these results, and supporting comparisons across formerly untested postures.

2.2. Thumb- and finger-based mobile device interaction

Owing to the widespread popularity of PDAs and hybrid PDA-phones, much recent attention has been paid to improving interaction with these devices. In particular, users often do not want to remove a stylus to interact with a device, preferring instead to use their index finger or thumb for quick access (Karlson et al., 2006). Guided by this motivation, Parhi et al. (2006) studied thumb-based interaction on PDA screens to determine the minimum target size necessary for accurate input. They found that targets of size 9.6 mm exhibited the same low error rate as larger targets, and that this value serves as a recommended minimum target size. Despite having a stated awareness of Fitts' law, this study did not evaluate enough target sizes and distances with a sufficient number of trials to create throughput measures of thumb performance.

Other efforts have focused on prototyping new designs for thumb- or finger-based interaction with mobile devices. Hirotaka (2003) clarified the importance of designing mobile phones for one-handed use, focusing on changes to the phone's form factor to enhance one-handed thumb performance. Specifically, the thumb's effective range of motion was measured at $68.1°$. However, no formal evaluation was carried out on the new design, although participants were reported as having liked its feeling.

AppLens and *LaunchTile* were designed specifically to support one-handed thumb-based use of mobile phones and PDAs (Karlson et al., 2005). AppLens is a mobile

---

[1] Recent work by Zhai et al. has indicated that Crossman's correction may not fully compensate for different speed-accuracy biases among participants (Zhai et al., 2004). However, the correction is still the best known method for accommodating such differences.

application shell that provides access to programs in a zoomable 9×9 grid via fisheye distortion, an on-screen object cursor, and a minimal thumb-based gesture vocabulary. LaunchTile uses a central "thumb widget" to zoom into a 36×36 application grid revealing 4×4 sub-grids and individual applications. Studies indicate that participants could quickly learn the AppLens gestures, and that overall they preferred AppLens to LaunchTile.

Pirhonen et al. (2002) developed the *TouchPlayer*, a finger-based gesture interface for controlling music playback on a Pocket PC without visual attention. The Pocket PC was mounted on the user's belt, and whole-screen gestures executed media player commands like next/previous track, and volume up/down. Non-speech audio feedback informed the user of the system's response. Lumsden and Brewster (2003) built on this idea by providing *progressive* audio feedback as finger gestures were made on a Pocket PC, finding that such feedback increased the accuracy of users' gestures.

Shift is a technique for enhancing pointing precision on mobile touch screens using the index finger (Vogel and Baudisch, 2007). Prior finger-based pointing techniques suffered from two problems, which Shift remedies: the occlusion caused by the hand, and the lack of precision caused by fingers that are larger than many targets. Prior to Shift, the *Offset Cursor* was a popular design that placed the selection point a standard distance above the finger contact point (Potter et al., 1988; Sears and Shneiderman, 1991). But the Offset Cursor required users to "aim low" in anticipation of the constant offset, ruining the possibility of direct-touch for large targets. Shift allows direct-touch on large targets, and uses a callout in which fine adjustments can be made for small targets. Like Offset Cursor, Shift uses lift-off to trigger selection, which is also used in the current study.

2.3. Previous back-of-device prototypes

A few recent prototypes and design concepts have emerged that attempt to utilize interaction on the back of a device. The idea is obvious enough when one observes the natural placement of the fingers on the back of a device when the device is being held. However, until now, no research has studied the motor performance of such hand postures.

Behind Touch placed a mobile phone keypad with raised tactile bumps on the back of a device, allowing users to first touch their desired key while previewing it on-screen, and then to confirm it by pressing harder (Hiraoka et al., 2003). *HybridTouch* mounted a touchpad on the back of a PDA to support finger-based vertical and horizontal scrolling with one hand while manipulating the stylus with the other (Sugimoto and Hiroki, 2006). A prototype by Silfverberg et al. (2006) used two separate back-of-device touchpads, one for zooming and another for panning (Fig. 2).

Some futuristic designs have also portrayed back-of-device interaction as concept sketches. *Gummi* is a design concept for a bendable handheld computer that used a rear touch-surface that can be manipulated by either set of fingers holding the pliable device (Schwesig et al., 2004). A study of phones throughout history to inspire new design ideas resulted in a screen-controlling touchpad region placed on the back of a futuristic mobile phone (Cheng and Buur, 2004). Both have the nice property of freeing up screen real estate on the front of the device.

A key problem in back-of-device interaction is that of occlusion, where the user's hands are obscured by the device itself. This problem was evident in Wigdor et al.'s (2006) study of "under the table" interaction, which found that targets had to be at least 4.5 cm in diameter for users to acquire them when unable to see their fingers. *LucidTouch* is a recent prototype that attempts to solve this problem by displaying users' hands on the screen (Wigdor et al., 2007). The prototype accomplished this with a camera mounted on a boom off the back of a device. Ideally, future prototypes would exchange the camera for a more subtle sensor.

Text entry systems have also employed limited back-of-device interaction. The *Twiddler* handheld chording keyboard enforces a grip wherein the user's fingers manipulate
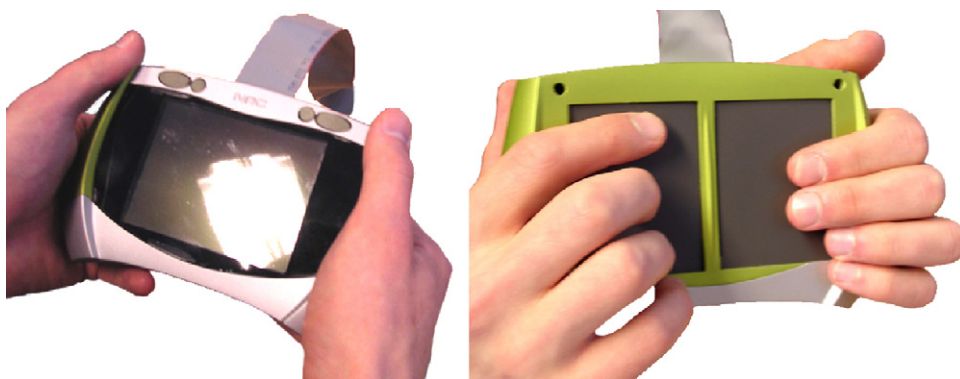

Fig. 2. Silfverberg et al.'s (2003) prototype. In the right image, the left touchpad is used to pan in four directions, and the right touchpad is used to zoom. Used with permission.

a set of keys on the backside of a controller (Lyons et al., 2004). The *ChordTap* system integrated three buttons on the back of a mobile phone that could be used by the fingers to disambiguate which letter the user desired as they pressed a key on the phone's front with the thumb (Wigdor and Balakrishnan, 2004). Recently, Wobbrock et al. (2007) embedded an isometric joystick in the back of a mobile phone to be used with the index finger in entering letter-like gestures from the EdgeWrite alphabet. Their study found that text entry speeds with this back joystick were about 70% of the speeds achieved on the front of the device with an identical joystick operated by the thumb.

3. Throughput of hand postures on a mobile device

In our first of three studies, we examined the performance of various hand postures in Fitts' reciprocal pointing tasks (Fitts, 1954; MacKenzie, 1992). Our goals were (1) to understand which hand postures were most effective for finger-based interaction with a mobile device, and (2) to determine which postural factors affected performance. The factors considered were one or two hands, index finger or thumb, front or back of device, and horizontal or vertical movement. Movement times, error rates, and nominal and corrected Fitts' law models were obtained for these conditions. These findings will help inform designers of mobile interaction techniques and provide modelers with useful parameters.

3.1. Method

3.1.1. Participants

Sixteen participants (9 male, 7 female) participated in our study. These participants were recruited from the local university communities. The only inclusion criterion was that participants were all right-handed. Participants' average age was 23.7 years ($SD = 3.5$). In a posttest questionnaire, all 16 participants indicated that they used a computer for "more than a few minutes each day." Eight participants said they used a touchpad daily as part of a laptop computer. None of the participants indicated that they used a PDA daily, but six said they owned a PDA but no longer used it regularly. Ten participants regularly used mobile phones.

3.1.2. Apparatus

In order to isolate the physical performance of hand postures on a mobile device, we had to create an experiment apparatus that circumvented the two well-known perceptual problems of direct-touch on touch screens, namely the occlusion caused by the user's hand and the unwieldy "fat finger" that creates an uncertain touch location (Potter et al., 1988; Sears and Shneiderman, 1991; Vogel and Baudisch, 2007; Wigdor et al., 2007). We argue that innovations in display and sensing technologies will continue to ameliorate these problems, and therefore they should not confound our results. The simplest and most intuitive way to circumvent these two problems was to separate the output display from the physical mobile device, a simple strategy employed in prior studies (Silfverberg, 2003; Wobbrock et al., 2007). Note that studies of mouse pointing necessarily do the same thing, since any physical mouse is distinct from the desktop monitor. Also, users of touchpads are accustomed to the separation of input and output surfaces. Given the close proximity of our mobile device to the output display, we had no reason to suspect that this decoupling caused any changes in our findings. No participants remarked at discomfort or uncertainty regarding this technique.
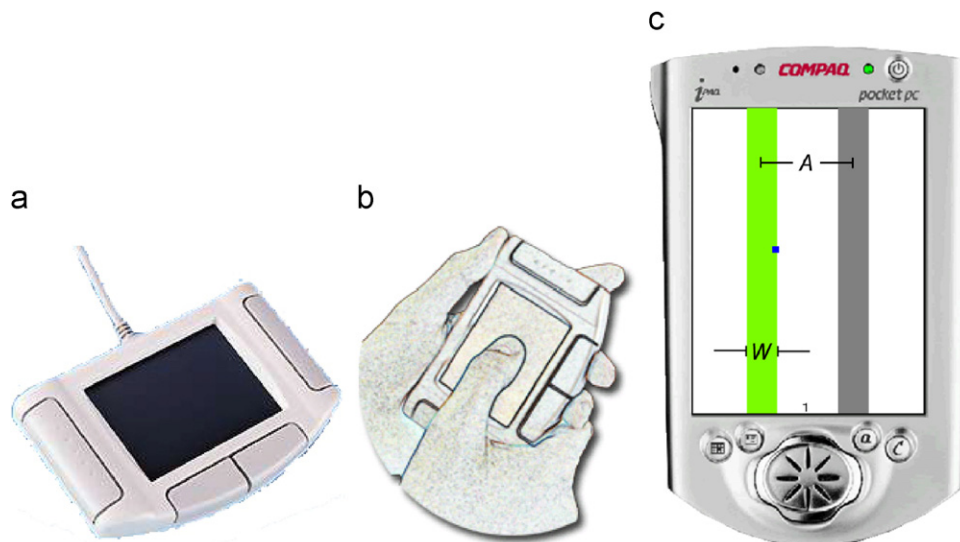

Fig. 3. (a) The synaptics touchpad used as our simulated mobile device. (b) An example posture with two-hands using the thumb on the front of the mobile device. (c) The simulated iPAQ shown on the computer monitor. The current finger position is the blue dot inside the left target (The "A" and "W" labels were not shown during trials).

For our study, we used a stand-alone Synaptics touchpad (Fig. 3a) rotated 90° as a simulated mobile device (Fig. 3b). The touchpad's dimensions were nearly identical to those of the Palm Vx PDA, and its sensing surface had the same aspect ratio as the Palm Vx screen. A simulated iPAQ screen (Fig. 3c) was portrayed on a Dell 1800FP 18″ flat screen monitor set to 32-bit color and 1280×1024 resolution. At this resolution, the iPAQ screen had a physical height and width of about $3.6″ \times 2.7″$. The touchpad's sensing surface measured $2.4″ \times 1.8″$. Thus, the control-display gain was 1.5.

The touchpad was mapped in absolute mode to the simulated iPAQ screen, i.e., the upper-left corner of touchpad directly corresponded to the upper-left corner of the iPAQ screen. We chose absolute mapping for three reasons: (1) it is the customary manner in which mobile device touch screens are used, (2) it has been shown to be faster than relative mapping for difficult selection tasks (Arnaut and Greenstein, 1986; Epps et al., 1986), and (3) absolute mode does not necessitate "clutching" (repeated lifting and swiping) as relative mode often does (MacKenzie and Oniszczak, 1998). Participants sat holding the touchpad in front of the monitor while looking at the simulated iPAQ display. They understood the mapping of the touchpad to the iPAQ screen and gave no indication of confusion in this regard.

We showed participants' finger or thumb location on the iPAQ screen as a bright red dot. When the dot entered the current target, it turned from red to blue (Fig. 3c). The location of the dot for a finger on the back of the device would be as if the finger were visible directly through the device, similar to LucidTouch (Wigdor et al., 2007). The position of the finger on the touchpad was read every 12.5 ms (80 Hz). A C# program running on Windows XP installed on a Dell Optiplex GX270 with 1 GB RAM and a 2.4 GHz processor presented trials and logged movements across the touchpad as (x, y, t) points with timestamps of 10^{-7} s precision and 10^{-3} resolution.

3.1.3. Procedure

The experiment presented participants with a variety of vertical and horizontal bar targets common to 1-D Fitts' law experiments. Participants were instructed to move their thumb or index finger across the touchpad surface to acquire targets quickly and accurately. The current target was highlighted in green, while the inactive target was gray (Fig. 3c). Target selection was implemented using *lift-off* (sometimes called *take-away*) as in many prior systems (Potter et al., 1988; Sears and Shneiderman, 1991; Parhi et al., 2006; Wigdor et al., 2006; Vogel and Baudisch, 2007). Although lift-off has known drawbacks (Buxton et al., 1985; Vogel and Baudisch, 2007), its continued use and refinement makes it the *de facto* standard for these types of interactions.

If the finger was lifted outside the target, an error sound was played. Regardless of whether a given target was hit or missed, the opposite target became the active one. Participants were instructed to return their finger to the touchpad immediately after lifting and proceed to the next target. When the finger was off the touchpad, no location was shown. No targets were placed along the extreme edges of the iPAQ display so that device edges played no role.

Before the first set of trials, participants were told how to use the touchpad. They were shown how to hold the device according to the postures shown in Fig. 4 and given one set of 14 practice trials to perform. During the actual test, participants performed four practice trials and 10 test trials per condition. Participants were instructed not to miss more than one test target every two sets of 10, or about $1/20 = 5\%$ of test trials.
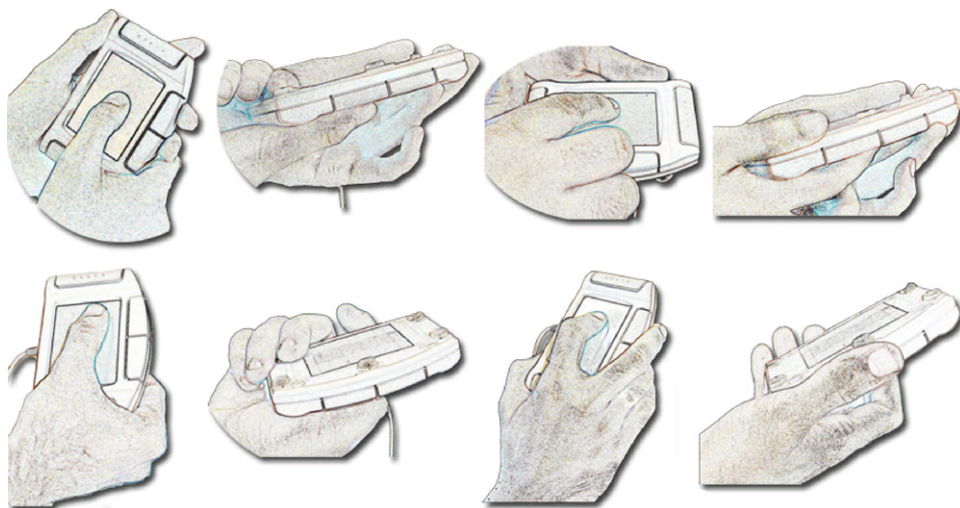

Fig. 4. The eight postures tested in our study, which reflect $2 \times 2 \times 2$ combinations of the *Hands*, *Finger*, and *Side* postural factors. *Top Row*: two-handed postures. *Bottom Row*: One-handed postures. *Left to Right*: thumb-on-front, thumb-on-back, index-on-front, index-on-back.

3.1.4. Design and analysis

The experiment was a within-subjects factorial design with the following factors and levels:

- Hands {1 hand, 2 hands}
- Finger {index finger, thumb}
- Side of device {front, back}
- Direction of movement {horizontal, vertical}
- Index of Difficulty (*ID*) {1.32 to 4.64 bits}
  - Amplitude (*A*) {6, 12, 24 units}
  - Width (*W*) {1, 2, 4 units}
- Trial {1…10}
- Participant {1…16}

For *A* and *W*, one unit equaled eight pixels on the simulated iPAQ screen and about $0.059''$ (1.5 mm) on the touchpad. These levels of *A* and *W* were taken from a prior study (Balakrishnan and MacKenzie, 1997). We use *ID* as a continuous factor ranging from 1.32 to 4.64 bits.

Each combination of *Hands*, *Finger*, *Side*, *Direction*, *A*, and *W* was administered to 16 participants, resulting in $2^4 \times 3^2 \times 10 \times 16 = 23{,}040$ total test trials. The first three factors—*Hands*, *Finger*, and *Side*—are the "postural factors" and determined the eight postures shown in Fig. 4. The postural factors were assigned to participants using a randomized Latin Square. Within each posture, the 18 *Direction*, *A*, and *W* combinations were administered randomly without replacement. Dependent measures were the average movement times and error rates for each set of 10 trials in a given condition.

Our movement time data were analyzed using a mixed-effects model analysis of variance with repeated measures (Littell et al., 1996b; Schuster and von Eye, 2001). *Hands*, *Finger*, *Side*, *Direction*, and *ID* were modeled as fixed effects, and *Participant* was correctly modeled as a random effect (Littell et al., 1996a; Frederick, 1999; Schuster and von Eye, 2001). Mixed-effects models typically retain larger denominator degrees of freedom (*df*s) than traditional multivariate methods. But these larger *df*s do not make detection of significance easier due to the use of wider confidence intervals based on population estimates (Frederick, 1999). After screening high-order effects for non-significance, our analysis model retained interactions to the third degree. In our results, we omit reporting the effects of *ID* since these effects are expected (i.e., trials with higher *ID*s were significantly slower and more error prone than easier ones).

Error rate data often do not conform to the assumptions underlying analyses of variance. Often errors are rare, creating a skew towards 0%. Our error rate data were no exception, violating the normality assumption. However, Poisson regression is well-suited to analyzing these data when they are viewed as error counts (Vermunt, 1997). The overall model was significant ($\chi^2_{(40, N=1280)} = 782.95$, $p < 0.0001$). Upon inspection, it was clear that Poisson regression was more conservative in its determination of

significance than an improperly applied analysis of variance.

3.2. Results

3.2.1. Adjustment of data

A benefit of using Crossman's correction for normalizing the speed–accuracy tradeoff is that all misses do not have to be discarded (MacKenzie and Soukoreff, 2003; Soukoreff and MacKenzie, 2004). However, as in prior studies (MacKenzie et al., 1991; Balakrishnan and MacKenzie, 1997), we regarded outliers as unrepresentative slips and removed them. Outliers were determined by trials whose selection coordinates were more than two[2] standard deviations from the mean selection coordinate on a given set of 10 trials. Means and standard deviations were calculated for each participant, and for each combination of *Hands*, *Finger*, *Side*, *Direction*, *A*, and *W*. There were 865/23,040 outlier trials (3.8%). Also, trials that followed a miss were removed, since the movement amplitude *A* of such trials would not be controlled. There were 1592/23,040 such trials (6.9%), 1389 of which were not outliers. Thus, a total of 2254/23,040 trials were removed from our data (9.8%).

A test for order effects revealed a significant effect of posture order on movement time ($F_{7,1144} = 11.50$, $p < 0.0001$). However, because our counterbalancing approach presents each posture in each test position an equal number of times, improvement over postures is acceptable so long as no posture is affected asymmetrically. Tests for interactions between posture order and all postural factors shows no significant results, indicating presentation order did not unduly favor one posture over another.

3.2.2. Movement times

The grand mean movement time (*MT*) was 1468 ms (*SD* = 166). Means and standard deviations for *Hands*, *Finger*, *Side*, and *Direction* and their combinations are shown in Table 1.

From the table, we can see that two-handed postures were generally faster than one-handed postures at 1360 *vs.* 1575 ms, respectively. This difference was significant for *Hands* on *MT* ($F_{1,1239} = 40.10$, $p < 0.0001$). Index fingers were faster than thumbs at 1408 *vs.* 1527 ms, respectively; this resulted in a significant effect of *Finger* on *MT* ($F_{1,1239} = 12.36$, $p < 0.001$). Coincidentally, the same two means were obtained for horizontal and vertical movement times (1408 *vs.* 1527 ms, respectively), resulting in a significant effect of *Direction* ($F_{1,1239} = 12.44$, $p < 0.001$).

Interestingly, there was no significant main effect of *Side* on *MT* ($F_{1,1239} = 2.20$, n.s.). Table 1 shows that front- and

---

[2]Some prior studies have used three standard deviations as the exclusion criterion. However, with only 10 trials informing each of these calculations, no trials would be excluded under this criterion due to the weight of even one slip.

Table 1
Movement times (ms) and error rates (%) for the eight postures and two movement directions tested in this study

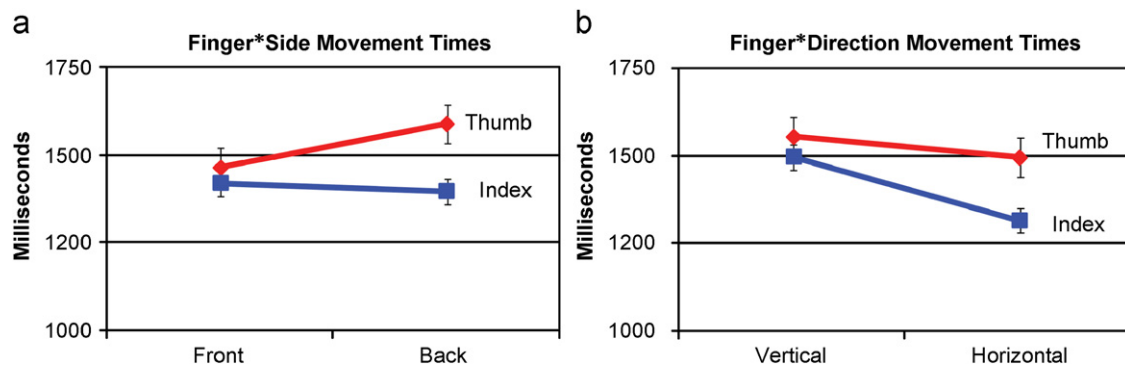|	Hands	Finger	Side	Direction	Movement times (ms)		Error rates (%)	
					Mean	SD	Mean	SD
A	One	Index	Back	Horizontal	1440	476	5.65	11.55
B	One	Index	Back	Vertical	1580	733	6.71	14.15
C	One	Index	Front	Horizontal	1325	541	5.08	11.39
D	One	Index	Front	Vertical	1727	886	10.44	19.90
E	One	Thumb	Back	Horizontal	1689	1471	9.23	17.59
F	One	Thumb	Back	Vertical	1805	1113	18.69	25.36
G	One	Thumb	Front	Horizontal	1537	1067	6.39	14.89
H	One	Thumb	Front	Vertical	1496	970	7.65	15.28
I	Two	Index	Back	Horizontal	1277	659	3.07	5.87
J	Two	Index	Back	Vertical	1288	478	4.36	10.78
K	Two	Index	Front	Horizontal	1228	394	2.72	9.35
L	Two	Index	Front	Vertical	1398	731	4.09	12.11
M	Two	Thumb	Back	Horizontal	1386	522	4.40	9.51
N	Two	Thumb	Back	Vertical	1476	667	6.98	15.44
O	Two	Thumb	Front	Horizontal	1379	789	3.30	9.00
P	Two	Thumb	Front	Vertical	1449	1053	8.03	18.29
Factor means				Horizontal	1408	149	4.98	2.15
				Vertical	1527	171	8.37	4.64
			Back		1493	188	7.39	4.95
			Front		1442	151	5.96	2.65
		Index			1408	170	5.26	2.47
		Thumb			1527	149	8.09	4.70
	One				1575	159	8.73	4.41
	Two				1360	87	4.62	1.90
Grand					1468	166	6.67	3.91

Letters A–P correspond to those appearing in Fig. 8.

Fig. 5. Two-way interactions for (a) *Finger*Side* and (b) *Finger*Direction* on movement time. Error bars are ±1 SE. Lower values represent faster performance.

back-of-device interactions generally took about the same amount of time: 1442 vs. 1493 ms, respectively. Upon closer examination, we see that although index fingers and thumbs performed similarly on the front of the device, thumbs slowed when moving to the back, but index fingers did not (Fig. 5a). This is confirmed by a significant *Finger*Side* interaction ($F_{1,1239} = 4.73$, $p<0.05$). Further, although index fingers and thumbs performed similarly for vertical movements, index fingers sped up more than thumbs did for horizontal movements (Fig. 5b). This is indicated by a marginal *Finger*Direction* interaction

($F_{1,1239} = 3.22$, $p = 0.07$). Finally, these two interactions affected each other disproportionately: when moving from the front to the back of the device, index fingers and thumbs were similarly slowed for horizontal motion but not for vertical motion, as index fingers sped up while thumbs slowed down (Fig. 6). This resulted in a significant *Finger*Side*Direction* interaction ($F_{1,1239} = 4.85$, $p<0.05$).

For one-handed postures, the two fastest conditions were the index finger in horizontal movements on the front and back of the device. Particularly noteworthy is that
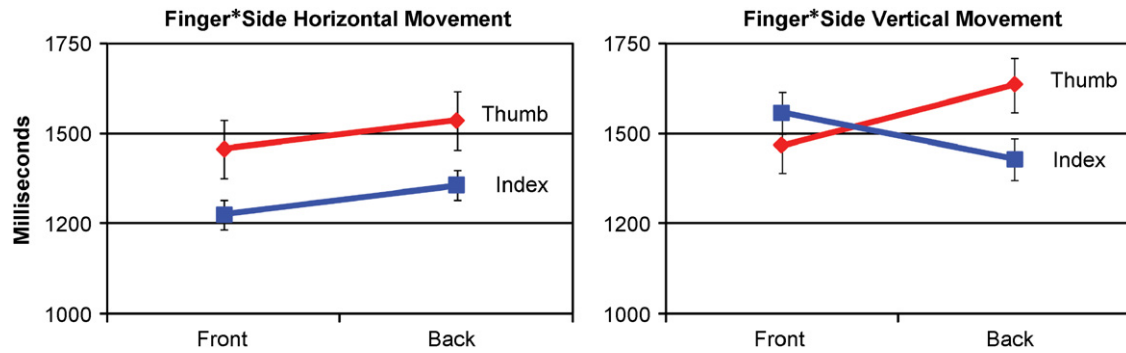
Fig. 6. A significant three-way interaction for *Finger*Side*Direction* on movement time is seen in the transition from one graph to the other. Error bars are ± 1 SE. Lower values represent faster performance.
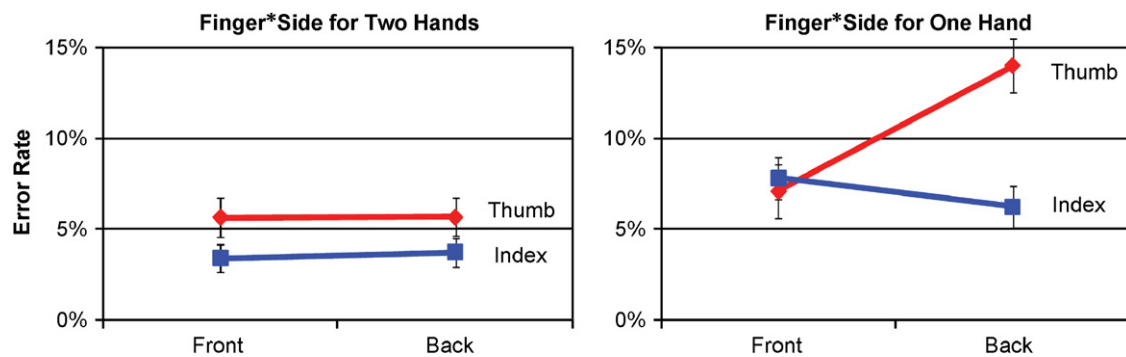

Fig. 7. A significant three-way interaction for *Hands*Finger*Side* on error rates is seen in the transition from one graph to the other. Error bars are ± 1 SE. Lower values represent more accurate performance.

these postures were both faster than the thumb on the front, which is commonly employed in many mobile device interactions. The two slowest postures were the index finger on the front and the thumb on the back, both in vertical movements.

For two-handed postures, the two fastest conditions were again the index finger in horizontal movements on the front and back of the device. As in the results for one hand, these were both faster than the thumb on the front. The two slowest postures were the thumb in vertical movements on the front and back of the device.

3.2.3. Error rates

The grand mean error rate was 6.67% ($SD = 3.91\%$). Along with movement times, error means and standard deviations for *Hands*, *Finger*, *Side*, and *Direction* and their combinations are shown in Table 1.

From the table, we can see that two-handed postures were generally more accurate than one-handed postures at 4.62% *vs.* 8.73% errors, respectively. This difference was significant for *Hands* ($\chi^2_{(1,N=1280)} = 34.51$, $p < 0.0001$). Further, index fingers were more accurate than thumbs at 5.26% *vs.* 8.09% errors, respectively, which was significant for *Finger* ($\chi^2_{(1,N=1280)} = 5.23$, $p < 0.05$). The front of the device was more accurate than the back at 5.96% *vs.* 7.39% errors, respectively. Thus, unlike for *MT*, for errors there was a significant effect of *Side* ($\chi^2_{(1,N=1280)} = 5.34$, $p < 0.05$).

Horizontal movement was also more accurate than vertical movement at 4.98% *vs.* 8.37% errors, respectively, which was significant for *Direction* ($\chi^2_{(1,N=1280)} = 7.80$, $p < 0.01$).

An examination of interactions shows an interesting interaction for *Hands*Finger*Side* ($\chi^2_{(1,N=1280)} = 6.32$, $p < 0.05$). As shown in Fig. 7, when using two hands, error rates increase about equally for thumbs and index fingers when comparing the front to the back of the device. But when using only one hand, the error rates for thumbs on the back dramatically increase, while the index finger error rate actually *decreases*. This again confirms the promise of using index fingers on the backs of devices for interaction. No other effects or interactions were significant for errors.

For one-handed conditions, the two most accurate postures were the index finger in horizontal movements on the front and back of the device. These were also the two fastest postures. They were both more accurate than the thumb on the front of the device. The two least accurate postures were also the slowest: the index finger on the front and the thumb on the back, both in vertical movements. Thus, the one-handed error rates mimic their movement time results.

For two-handed conditions, the two most accurate postures were again the index finger in horizontal movements on the front and back of the device. As in the results for one hand, these were both more accurate than the thumb on the front. Again, the two least accurate postures
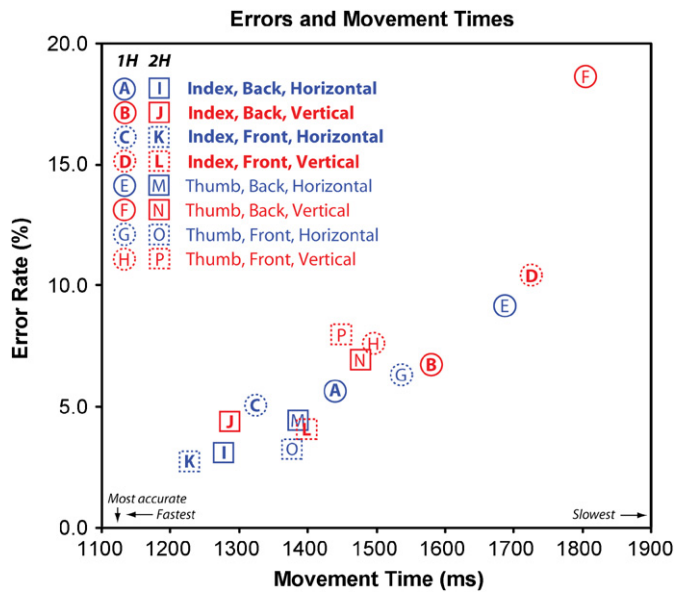
Fig. 8. The eight postures in two movement directions (horizontal, vertical) plotted with error rates (%) and movement times (ms). Points closest to the origin are best. *Key*: One-handed postures are circles, two-handed postures are squares; non-bold letters are thumbs, bold letters are index fingers; dashed lines are front-of-device, solid lines are back-of-device; blue lines are horizontal movements, red lines are vertical movements. See Table 1 for precise numeric values and matching letter codes.

were also the slowest: the thumb in vertical movements on the front and back of the device. Thus, the two-handed error rates mimic their movement time results.

An error rate × movement time plot of all eight postures in both movement directions is shown in Fig. 8. Points closer to the origin are faster and more accurate, while points toward the upper-right are slower and more error prone. Note how bold letters (A–D, I–L) are generally closer to the origin than non-bold letters (E–H, M–P) within each hand, indicating the superior performance of the index finger to the thumb. Note also how blue shapes are generally closer to the origin than red shapes, indicating the superior performance of horizontal to vertical movements. Also, note the general mixing of dashed and solid shapes, indicating the generally comparable performance of front and back of device postures. Finally, note the infiltration of points A and C (circles) into solid two-handed territory (squares). In other words, one-handed index finger horizontal movements appear to be competitive with some two-handed postures on both the front and back of the device, opening new possibilities for design. It is interesting that both of these postures are superior to the one-handed thumb-on-front postures (G, H), which are commonly used to manipulate current mobile devices.

3.2.4. Fitts' law and throughputs

We can formalize the performance of our hand postures while accommodating the speed–accuracy tradeoff by fitting our data to corrected Fitts' law models using ID_e. For comparisons, we also fit our data using the traditional

method of discarding error trials and using nominal IDs (MacKenzie and Soukoreff, 2003; Soukoreff and MacKenzie, 2004). Equations for these models are shown in Table 2. Most have very high correlations ($R^2 \geqslant 0.90$).

For one-handed postures (Fig. 9a), the two conditions with the highest index of performance (*IP*) were index fingers on the front (C) and back (A) in horizontal movement. The lowest *IP*s were for vertical movement with the index finger on the front (D) and the thumb on the back (F). For two hands (Fig. 9b), the highest *IP*s were for index fingers on the front in horizontal motion (K) and on the back in vertical motion (J). The lowest two-hand *IP*s were for vertical motion with the thumb on the front (P) and back (N). These comparisons hold for both the ID_e and ID models, although the entire rank-order of all 16 throughputs is not quite identical.

In general, our corrected throughputs are a bit lower than what has been reported in prior studies (Balakrishnan and MacKenzie, 1997; Gokturk and Sibert, 1999), but our nominal models are similar to prior results. Unlike prior studies of finger pointing, however, our task involved finger dragging, since participants had to be in contact with the device in order to see their finger position. Dragging expectedly lowers throughputs (MacKenzie et al., 1991), but it represents the current state-of-the-art for back-of-device (or under-the-table) interaction, since the finger position must be communicated to the user prior to selection (Wigdor et al., 2006). Although cutting-edge prototypes are attempting to overcome these limitations (Wigdor et al., 2007), much progress is still necessary before back-of-device target acquisition can be accomplished without need of localizing contact with the device. As for front-of-device interactions, although direct touch *is* possible, its drawbacks have prompted the creation of techniques that involve initial contact with the device prior to lift-off selection (Potter et al., 1988; Sears and Shneiderman, 1991; Vogel and Baudisch, 2007). Thus, our experiment and resulting throughput measures are relevant to both current and future designs for front- and back-of-device interaction.

4. Orientations on the front and back of a device

Our first study examined fundamental movement performance in various hand postures on the front and back of a mobile device. Other aspects of front- and back-of-device interaction are also worthy of investigation. One question is how perceptions of *shape orientation* are affected by postures on different sides of a device, and whether or not visual feedback affects those perceptions. Our second study examined these issues.

4.1. Method

The same 16 participants from the first study participated in the second study. The same apparatus was used, except that no Fitts' reciprocal pointing task was visible.

Table 2

Corrected and nominal Fitts' law models for the eight postures and two movement directions tested in this study

	Hands	Finger	Side	Direction	ID_e model				ID model			
					Intercept a (ms)	Slope b (ms/bit)	R^2	IP (bits/s)	Intercept a (ms)	Slope b (ms/bit)	R^2	IP (bits/s)
A	One	Index	Back	Horizontal	681	836	0.992	1.196	617	297	0.961	3.367
B	One	Index	Back	Vertical	319	1688	0.830	0.592	521	367	0.900	2.725
C	One	Index	Front	Horizontal	588	786	0.947	1.272	539	282	0.959	3.546
D	One	Index	Front	Vertical	214	2104	0.957	0.475	184	564	0.907	1.773
E	One	Thumb	Back	Horizontal	359	1588	0.953	0.630	627	305	0.993	3.279
F	One	Thumb	Back	Vertical	48	2680	0.844	0.373	382	521	0.892	1.919
G	One	Thumb	Front	Horizontal	135	1610	0.913	0.621	324	403	0.958	2.481
H	One	Thumb	Front	Vertical	174	1541	0.877	0.649	79	511	0.870	1.957
I	Two	Index	Back	Horizontal	518	783	0.994	1.277	440	297	0.914	3.367
J	Two	Index	Back	Vertical	555	743	0.991	1.346	475	287	0.973	3.484
K	Two	Index	Front	Horizontal	603	599	0.990	1.669	499	265	0.960	3.774
L	Two	Index	Front	Vertical	508	876	0.959	1.142	377	359	0.918	2.786
M	Two	Thumb	Back	Horizontal	537	972	0.953	1.029	552	295	0.963	3.390
N	Two	Thumb	Back	Vertical	299	1310	0.837	0.763	375	386	0.938	2.591
O	Two	Thumb	Front	Horizontal	495	892	0.832	1.121	398	350	0.811	2.857
P	Two	Thumb	Front	Vertical	336	1156	0.910	0.865	196	453	0.904	2.208
Factor means				Horizontal	490	1008	0.947	1.102	500	312	0.940	3.258
				Vertical	344	1512	0.901	0.776	359	431	0.913	2.430
			Back		467	1325	0.924	0.901	499	344	0.942	3.015
			Front		382	1196	0.923	0.977	360	398	0.911	2.673
		Index			498	1052	0.958	1.121	457	340	0.937	3.103
		Thumb			334	1469	0.890	0.756	408	403	0.916	2.585
	One				353	1604	0.914	0.726	456	406	0.930	2.631
	Two				481	916	0.933	1.152	414	337	0.923	3.057
Grand					421	1260	0.924	0.939	434	371	0.926	2.844

Letters A–P correspond to those appearing in Fig. 9.

Instead, participants wrote four capital letters on both sides (front and back) of three surfaces:

- Opaque paper
- Device *without* visual feedback of the stroke
- Device *with* visual feedback of the stroke

In the conditions with opaque paper, participants used a heavy felt pen on thick white paper. In the condition *without* visual feedback, participants could not see the computer monitor as they moved their finger on the touchpad, which was held vertically to mimic a PDA as in the first study. In the condition *with* visual feedback, a persistent stroke trace was displayed on the simulated iPAQ screen. This trace was rendered as if the user's finger were visible directly *through* the device (Fig. 10).

For all 2 sides × 3 surfaces = 6 conditions, participants wrote capital E, O, U, and L letters, which were chosen for their different symmetries with respect to the horizontal and vertical axes (Fig. 11). Whether a participant first wrote on the front or back of a given surface was random, but participants always used the surfaces in the order listed above to prevent bias from the visual feedback conditions affecting the other conditions. The dependent measures were the orientations of the traces relative to the surfaces on which they were drawn.

As shown in Fig. 10, the key issue was whether participants moved their fingers on the back of the device according to the usual *motor pattern* for making such letters, or in a motor-reversed fashion in order to produce a *visually correct* result as if they were looking through the paper or device.

4.2. Results

Participants wrote all four letters E, O, U, and L in a consistent manner regarding orientation. Thus, letter symmetry did not have any effect. In writing these letters in the 2 sides × 3 surfaces = 6 conditions, participants exhibited one of two different orientations:

- *Visual-correct*: If one could see through the surface of the device, the path would appear correct. Relative to the hand that created it, this path shows a motor-reversed pattern. (Fig. 10 is an example of this.)
- *Motor-correct*: If one could see through the surface of the device, the path would appear mirrored. Relative to the hand that created it, this path would exhibit the usual motor pattern.
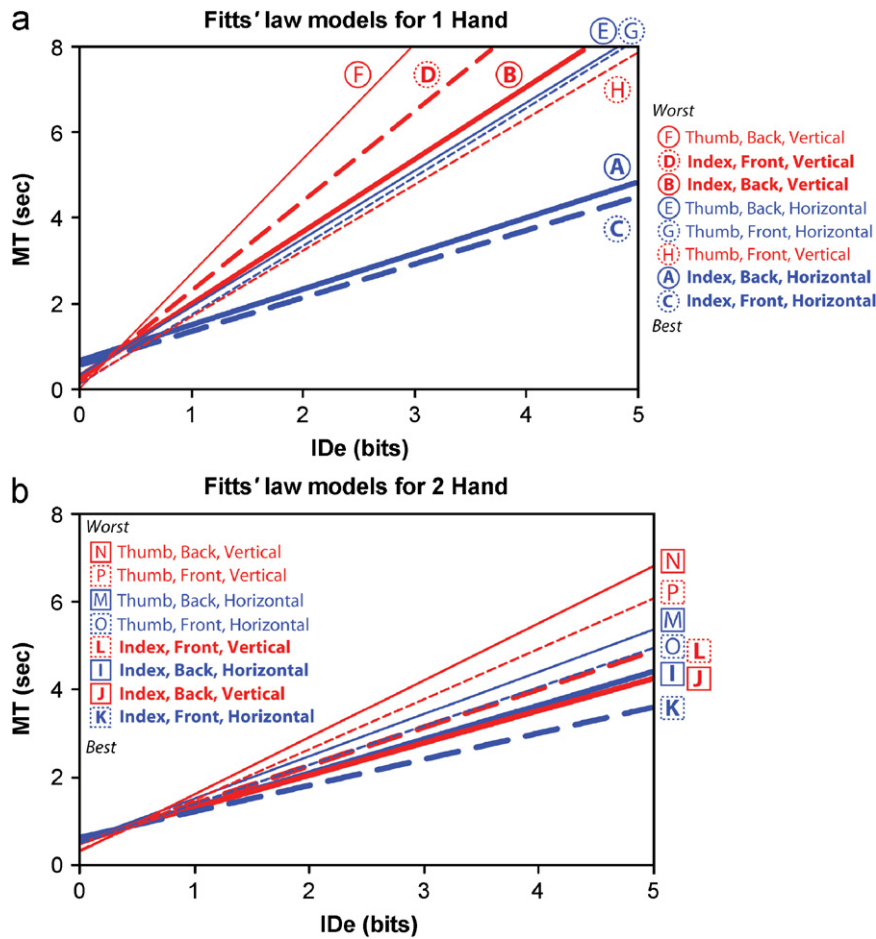
Fig. 9. Fitts' law models using ID_e for (a) one and (b) two hands. *Key*: One-handed postures are circles, two-handed postures are squares; regular lines and letters are thumbs, bold lines and letters are index fingers; dashed lines are front-of-device, solid lines are back-of-device; blue lines are horizontal movements, red lines are vertical movements. Lines with shallower slopes are better. See Table 2 for precise numeric values and matching letter codes.
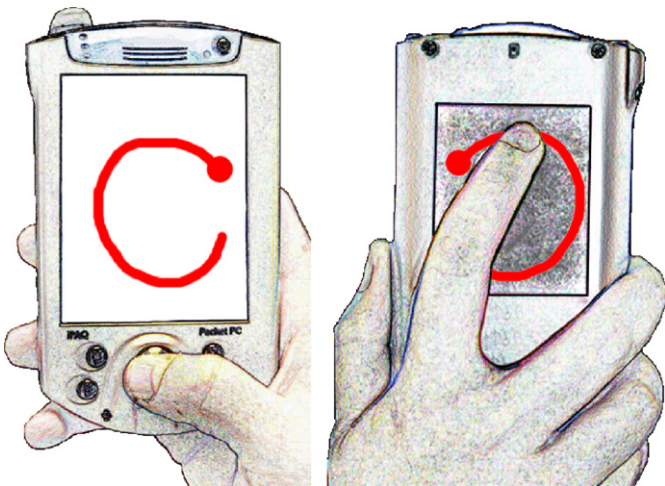

Fig. 10. Our hypothetical device from Fig. 1, now showing a stroke-trace for the letter "C". For "C" to appear correct to a user looking "through" the device, the motor pattern executed by the finger must be reversed (i.e., mirrored). It was unknown if participants would exhibit this behavior, or whether they would prefer to execute familiar motor patterns.

Fig. 11. The capital letters E, O, U, and L entered by participants on 2 sides × 3 surfaces. These letters were chosen because they exhibit different symmetries.

Note that on the front sides of our three surfaces, the *visual-correct* and *motor-correct* schemes describe the same behavior because the motor pattern is articulated in plain sight. On the back of the device, however, these two orientations differ by reflection over the vertical axis.
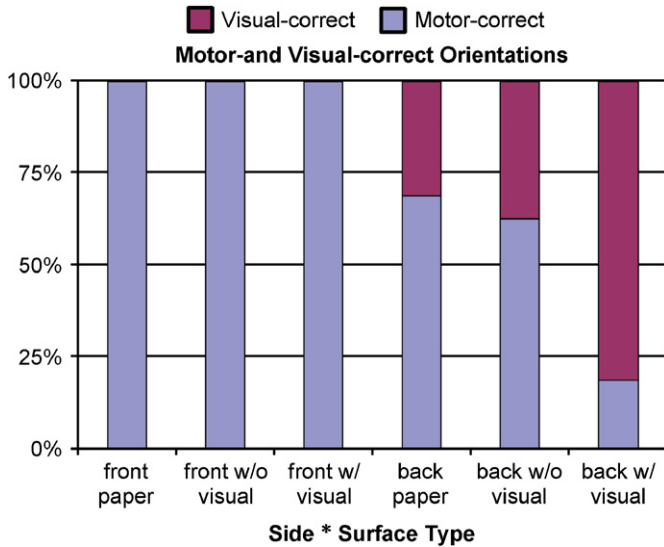
Fig. 12. Proportions of *motor-* and *visual-correct* orientations on 2 sides × 3 surfaces. The interesting result shown here is that the expectation of visual feedback when writing on the back of a device causes most users to switch from a *motor-correct* to *visual-correct* movement scheme.

As shown in Fig. 12, no participants (0/16) wrote in mirror images on the fronts of any surfaces. That is, they all wrote in the *motor-correct* fashion. This is no surprise.

For the opaque paper condition, 11 of 16 participants (68.8%) wrote in a *motor-correct* fashion on the back. When using the device without any visual feedback, 10 of 16 participants (62.5%) wrote in a *motor-correct* fashion on the back. Surprisingly, however, when using the device *with* visual feedback, only 3 of 16 participants (18.8%) wrote in a *motor-correct* fashion on the back. Importantly, the 13 participants who wrote in a *visual-correct* manner did *not* first begin in a motor-correct fashion and then switch directions. In fact, they *began their trace* in a visual-correct fashion despite most of them having just done the opposite in the two previous conditions! It therefore seems that the mere expectation of seeing their output on the simulated iPAQ display caused many of them to reverse their motor pattern. This finding has implications for the design of systems that involve back-of-device motor behavior. If finger movements are to be displayed to the user, a *visual-correct* mapping should be used; otherwise, a *motor-correct* mapping may be used, or at the very least, it should be supplied as an option.

Nonparametric tests of proportions yielded significant differences between the paper and visual feedback conditions ($z = 2.85$, $p < 0.01$) and between the two device conditions ($z = 2.52$, $p < 0.05$). These findings are supported by prior psychological work on identification of letter traces on various surfaces of the body, e.g., on the back and front of the hand (Parsons and Shimojo, 1987). The findings are also supported by recent prototypes using an index finger on the back of a mobile phone for joystick-based text entry (Wobbrock et al., 2007).

5. Unistrokes on the front and back of a device

The third study combined elements from the first two. Like the first study, our third study measured the speed and error rate performance of hand postures. Like the second study, it involved the making of letters on the front and back of a device. But whereas the first study was a "pure performance" study designed to measure and model throughput, the third study was more applied, investigating whether participants could make unistroke characters on the front and back of a device. This would require participants to make composite motions beyond simple horizontal and vertical movements as they did in the first study.

5.1. Method

The same 16 participants entered letters from the EdgeWrite alphabet using various hand postures on the front and back of a device. The EdgeWrite alphabet (Fig. 13) is a unistroke alphabet (Goldberg and Richardson, 1993) designed for accessibility that uses physical edges to guide the finger on a touchpad, resulting in higher physical stability and passive haptic feedback (Wobbrock et al., 2003, 2004). Although another recent prototype put EdgeWrite on the back of a mobile phone using an isometric joystick (Wobbrock et al., 2007), this prototype lacked the physical edges or haptic qualities of the version used for this study.

Participants used the same apparatus as in the two prior studies but with one modification: the touchpad was augmented with a clear plastic template containing a square hole measuring 0.87″ on a side (Fig. 14a). The square hole defined the EdgeWrite input area and corresponded to a square area displayed on the simulated iPAQ screen (Fig. 14b). Strokes made within the physical
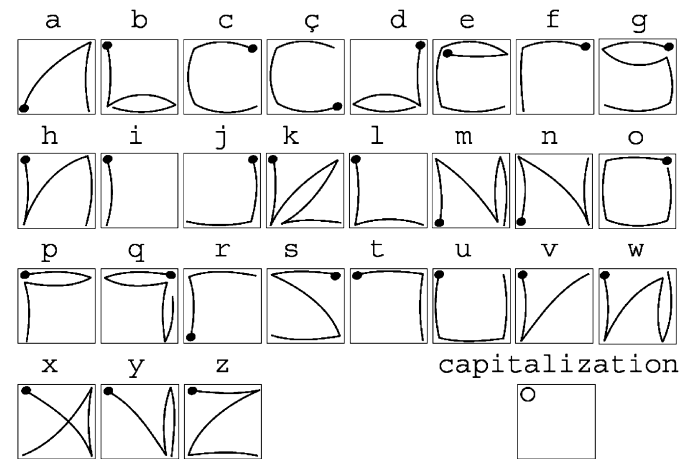

Fig. 13. The EdgeWrite letters. Note that these images are stylized with bowed arcs to indicate the order that the corners should be hit. All movement, however, is in straight lines along edges or diagonals and into corners.
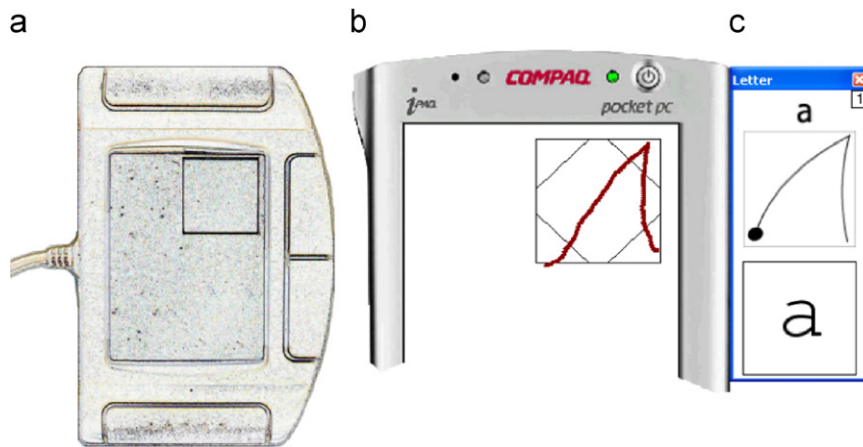
Fig. 14. (a) The device with a plastic overlay and a square hole in the top-right corner. (b) The simulated iPAQ screen with corresponding square. An EdgeWrite "a" has been drawn. (c) The stroke prompt and character output.

square on the touchpad were drawn in the corresponding square on the iPAQ.

For a given posture, participants were prompted with each letter of the EdgeWrite alphabet and asked to enter it three times "quickly and accurately." The EdgeWrite stroke for each letter was also depicted (Fig. 14c) because participants were not familiar with the EdgeWrite alphabet. Our goal was to isolate participants' motor performance, not test their ability to learn the gestures. After a participant performed a stroke, the resulting letter appeared below the EdgeWrite stroke depiction. If the participant's entry was correct, the resulting letter was shown in black; otherwise, it was shown in red and an audible error beep was played. We chose this design over a phrase-based text entry study because of participants' unfamiliarity with EdgeWrite. As stated, our interest was in isolating motor performance.

As in the first study, postures were determined by the combinations of the *Hands*, *Finger*, and *Side* factors. However, due to the poor performance of the thumb on the back, we omitted these two postures. Thus, participants used $2^3 - 2 = 6$ postures in all. The postures were assigned to participants according to a randomized Latin Square. In each posture, they entered 78 letters, for a total of $16 \times 6 \times 78 = 7488$ total letters for the study. Dependent measures were participants' average letter speeds and error rates over the 78 characters for a given posture. Letter speeds were measured from the time the finger contacted the touchpad until the finger was lifted.

5.2. Results

5.2.1. Letter times

A test for order effects revealed no significant effect of posture order on letter time ($F_{5,51} = 1.97$, n.s.). Furthermore, no interactions involving posture order and the three postural factors were significant. These results indicate adequate counterbalancing.
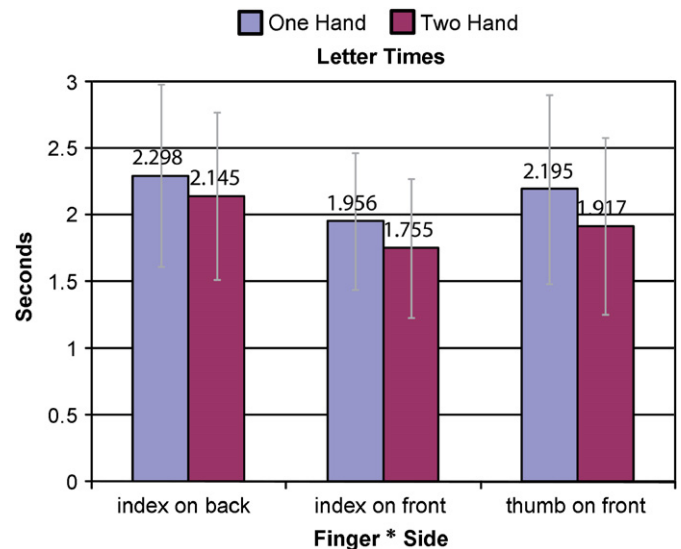

Fig. 15. Average letter times for our 6 postures for one and two hands. Error bars are ± 1 SD. Lower values indicate faster letter-making speeds.

Trials for incorrect letters were omitted from the analysis of letter times. This excluded 1016/7488 entries (13.6%). Average letter times for the six postures are shown in Fig. 15. These data showed significant main effects for *Hands* ($F_{1,75} = 5.14$, $p < 0.05$), *Finger* ($F_{1,75} = 4.46$, $p < 0.05$), and *Side* ($F_{1,75} = 14.88$, $p < 0.001$). However, no interactions among *Hands*, *Finger*, and *Side* were significant.

For one-handed postures, the index finger on the front of the device was fastest (1.96 s), the thumb on the front was next (2.20 s), and the index finger on the back was the slowest (2.30 s). This ordering was the same for two-handed postures: index-on-front (1.76 s), thumb-on-front (1.92 s), and index-on-back (2.15 s). The speed of the index-on-front poses is reminiscent of our first study results in which one- and two-handed postures with the index-on-front moving horizontally were the fastest within their respective number of hands.

5.2.2. Error rate

Unlike our first study, in which errors were relatively rare, errors in the third study were more common and were more appropriately analyzed with an analysis of variance than a Poisson regression model.

As noted in the previous section, the total number of errors was 1016/7488 entries (13.6%). Average letter errors for the six postures are shown in Fig. 16. These data showed significant main effects for *Hands* ($F_{1,75} = 5.58$, $p < 0.05$) and *Side* ($F_{1,75} = 6.54$, $p < 0.05$) but not *Finger* ($F_{1,75} = 0.79$, n.s.). However, as with letter times, no interactions among these factors were significant.

The error rates of two-handed postures followed the rank-ordering of letter times, but the error rates of one-handed postures did not. For one-handed postures, the thumb had the lowest error rate (11.9%), followed by the index finger on front (16.7%), and the index finger on back (20.6%). For two-handed postures, the index finger on the front had the lowest error rate (8.3%), the thumb was second (9.9%), and the index finger on the back was third (13.9%).

In this study, the error rates for two-hands are also reminiscent of the error results in the first study. But for one hand, the thumb, which was the most accurate here, fared relatively worse in the first study. However, the one-handed thumb was about 0.24 s slower than the one-handed index finger on the front, so there is a speed–accuracy tradeoff at work.

One way to combine speed and accuracy into a single measure is to create an adjusted speed by penalizing raw speed based on the number of errors (Wobbrock, 2007). We can do this with the following equation:

$$LT_{adj} = LT_{raw} \times (1 + LE) \tag{4}$$

Eq. (4) defines the adjusted letter time (LT_{adj}) as an accuracy-based increase of the raw letter time (LT_{raw}). In
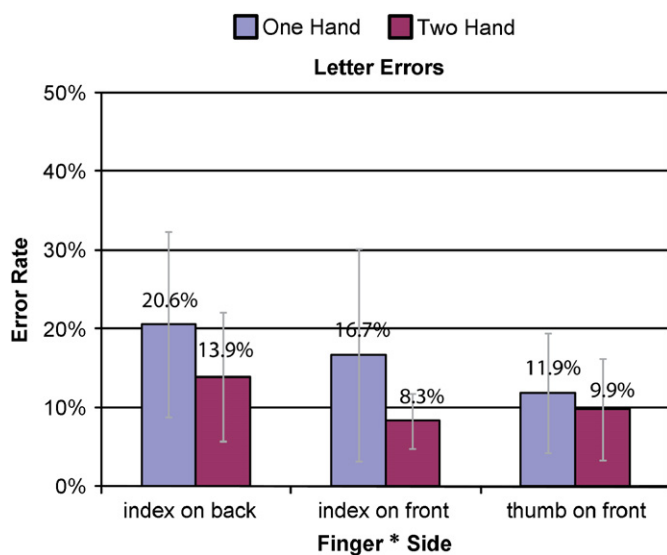

Fig. 16. Average letter errors for our 6 postures for one and two hands. Error bars are ±1 SD. Lower values indicate more accurate letter-making.
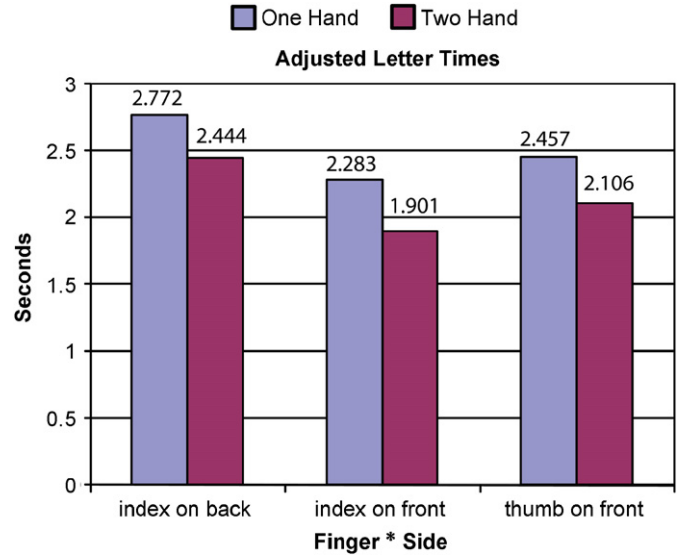

Fig. 17. Adjusted letter times based on raw letter times and letter errors: $LT_{adj} = LT_{raw}(1 + LE)$. Error bars are not applicable because these are single values without a distribution.

this case, the accuracy-based increase is 1 plus the rate of letter errors (*LE*). Thus, perfect accuracy results in exactly the raw letter times. On the other hand, if participants made 10% errors, the raw letter times would be increased by 10%. Fig. 17 shows the adjusted letter times for the six postures. The relative orderings are unchanged from the raw letter times, but there is a greater separation of scores.

6. Discussion and design implications

At the conclusion of all three studies, it is appropriate to ask, "What are the lessons here?" This section answers this question for each of the three studies, highlighting their main findings and implications for design.

The first study showed how various postural factors affect finger speed and accuracy in target selection with a mobile device. For both one and two hands, the index finger performed surprisingly well, especially on the back of the device; in many cases, it had lower movement times and error rates than the thumb on the front. The two best throughputs for both one and two hands involved index fingers on the front and on the back. This is particularly interesting in light of the fact that most modern mobile devices, particularly mobile phones and mini-QWERTY devices, are designed for heavy thumb-based use. These performance results suggest that there may be opportunities for using the index finger on the back of a mobile device—perhaps not as a replacement to thumb-on-front techniques, but as a complement to them.

However, these findings must be accompanied by an awareness of the susceptibility of the postures to movement direction, particularly on the back of the device. Most postures were generally more successful in making horizontal movements than they were in making vertical ones. This was particularly evident for the index finger on the

front and back of the device. For both one and two hands, the index finger on the front of the device moving horizontally was the fastest and most accurate posture. The second best posture in both speed and errors was the index finger on the back of the device, again moving horizontally. But vertical movements were much more challenging, particularly with one hand or with the thumb. The design implications are clear enough: devices and user interfaces should support horizontal motion whenever possible, laying out buttons and other targets in rows rather than columns, and making their devices wide rather than tall and narrow. If only one row of buttons is used, buttons should be tall, allowing for vertical error. Layouts of screen elements such as sliders, scrollbars, or toolbars should also be horizontal so that a finger on the back of a device can operate them in that dimension. Unfortunately, many mobile phones still use traditional telephone keypads that are taller than they are wide.

The index finger was more resilient than the thumb with respect to postural changes. The degradation of the index finger when moving from two hands to one, or when going from the front of the device to the back, was less pronounced than for the thumb, which was affected more negatively by such changes. In particular, the index finger on the front of the device was successful in enough postures that devices could be designed to accommodate an escalation of postures. For example, a device might first be worn on the belt to be operated quickly via the index finger with one hand for common operations (Pirhonen et al., 2002). When greater engagement is necessary, the device could be removed from the belt, manipulated in one hand using the index finger on the back, and then returned to the belt, perhaps while walking. If still greater engagement is required, the user could hold the device in two hands while the thumb and index finger could both become involved. This type of "postural escalation" could provide a model for thinking about how device designs can support increasing levels of interactivity based on posture considerations.

The second study's key finding was that participants' expectation of receiving front-of-device visual feedback of finger movements on the back of a device changed the orientation in which they wrote letters. When the surface on which they articulated their gestures was "inert," offering no expectation of feedback, orientations were generally dominated by motor schemas. But when participants expected the device to *portray* their movement, a visual schema won out, causing participants' motor patterns to reverse relative to the hands that were making them. This provides the first empirical evidence to support the design of feedback systems for interaction on the back of a device (Wigdor et al., 2007).

The third study examined participants' ability to make complex gestures on the front and back of a device. These gestures were letters from the EdgeWrite alphabet (Wobbrock et al., 2003), and were articulated inside a square plastic hole. Letter speed and accuracy results generally favored the index finger on the front, which also had been the most successful in the first study. Unlike the first study, however, the thumb on the front was faster and more accurate in the letter-making study than the index finger the on the back, which was not the case in the first study, where the index finger on the back was more successful, at least for horizontal movement. These differences may be due to three factors. First, the movement amplitudes from the first study to the third study shrunk, which favored the shorter thumb relative to the longer index finger. Second, the first study separated horizontal from vertical movement, which, as we saw, was quite revealing of their performance differences. The third study, in contrast, involved gestures that were composites of horizontal and vertical motions, and indeed, there were more vertical movements in the EdgeWrite alphabet than horizontal ones. An examination of Fig. 13 reveals that in the 26 letters "a"–"z", there are 40.5 vertical movements and 35.5 horizontal ones.[3] Thus, the gestures were less favorable to the index finger on the back, particularly for one hand. Third, EdgeWrite is reliant upon physical edges to provide stability of motion, and edges are known to affect aspects of movement such as variability, accuracy, and speed (Wobbrock, 2003). It may be that the thumb benefited more from these edges than did the index finger on the back. The results from the first and third studies seem to suggest that the thumb is a better path-maker than it is a pointer. This finding is supported by prior work showing the superiority of the index finger over the thumb as a pointer (Silfverberg et al., 2000).

A design implication is that the index finger on the back of the device may be better suited to simple 1-D movements, particularly horizontal ones, and that complex gestures should be reserved for the front of the device. Thus, the index finger could serve for scrolling, panning, zooming, manipulating a slider (e.g., volume control), or choosing among a linear set of possible contexts within which actions on the front of the device could work. Indeed, some recent prototypes have begun to explore these ideas, albeit without the performance data provided here (Silfverberg et al., 2003; Schwesig et al., 2004; Sugimoto and Hiroki, 2006).

7. Future work

The three studies presented in this paper indicate multiple directions for future work. The one-handed posture using the index finger on the back of the device has some nice properties for one-handed mobile device interaction: it does not obscure the screen, it does not alter the natural gripping or holding posture, it allows the user to grasp the bulk of the device, not the device's bottom third (Hirotaka, 2003), and it achieves high throughput for horizontal tasks. A next step would be to design interaction techniques that take advantage of these potential benefits,

---

[3] These numbers treat diagonals as 0.5 vertical + 0.5 horizontal.

and to evaluate working systems in domain-specific tasks. Although certain prototypes have been mentioned throughout this paper, few of them have been evaluated in real-world settings.

Although lift-off is a common technique for achieving selection on touch screens or touchpads (Potter et al., 1988; Sears and Shneiderman, 1991; Wigdor et al., 2006), new techniques should be investigated to provide greater robustness and performance for back-of-device target selection. One possibility is a *lift-and-tap* operation. Unlike the conventional use of *tap*, where the tapped point is the selection location, with lift-and-tap, the tap point would *not* be the selection point; instead, the *lift-point* is the selected location, and tapping *anywhere* merely confirms this. This scheme obviates the need for accuracy in the tapping operation itself, allowing the lift-point to determine the target. But unlike mere *lift-off* selection, the *tap* provides the confirmation, increasing robustness and reducing false positives. The finger could be moved on the back of a device at will, without fear of making accidental selections unless a deliberate tap is made following a lift. If a lift is made without a subsequent tap, no operation occurs.

Another step would be to repeat or extend our experiment using a mobile device while the user is walking or riding in a bus, car, or train. If secondary tasks were involved, the one-handed postures may perform better than those requiring two hands. Although researchers have begun to investigate performance while on the move, few performance studies have been conducted as of this writing. Notable exceptions exist (Lumsden and Brewster, 2003; Lin et al., 2005; Mizobuchi et al., 2005; Oulasvirta et al., 2005; Marentakis and Brewster, 2006), but these generally have not investigated the performance of hand postures or considered the back of the device as a viable interactive surface. Indeed, the challenges of conducting controlled performance studies in the wild are numerous and difficult.

Another avenue for future work is to conduct a study of two-handed *simultaneous* interaction on the back of a two-handed mobile device. Our study assumed a mobile phone or PDA form factor, but a new breed of two-handed form factors called ultra mobile personal computers (UMPCs) is emerging. These devices are larger than mobile phones and PDAs and must be carried and operated using two hands. Thus, up to eight fingers can contact the back of the device, but the performance of such postures remains unknown. New prototypes are already considering such form factors (Silfverberg et al., 2003; Schwesig et al., 2004; Wigdor et al., 2007), but no studies have been conducted. In particular, it may be interesting to revisit certain two-handed models like the *kinematic chain* (Guiard, 1987), and to explore the extent to which back-of-device actions can set a "reference frame" while front-of-device actions operate within that frame. Although some two-handed interaction designs have been explored (Hinckley et al., 1998), they have not linked back- and front-of-device

postures in this way. Future studies exploring this issue might also explore how the size of a device impacts performance, since two-handed postures will often be used with devices that are larger than the PDA- or phone-sized device we have studied here.

Future studies could also include relative positioning instead of just absolute positioning as we have done here. With relative positioning, the mobile device would have an on-screen cursor (Wobbrock et al., 2002; Karlson et al., 2005), and finger movements would cause the cursor to move, much like a mouse. This might reduce the accuracy demands on the fingers, since their actual point of contact with the device would not matter. However, relative positioning introduces the possibility of clutching, which is generally slow and undesirable.

It is important to note that keypad-based interfaces may produce different results than our touch screen results here. Keypad interfaces should be tested, particularly with respect to the use of index fingers on the backs of devices, to see if successful interactions are possible. One difference between the current studies and those using keypads is the need for *pressure* to depress keys. Another difference is the passive haptic feedback that raised tactile keys provide (Silfverberg, 2003). It may be the case that highly tactile keys are necessary on the back of a device to provide the index finger with salient cues for accurate targeting. Or, perhaps elevated keys make it too difficult for the index finger to move over the back of the device comfortably. These and other issues must be examined if back-of-device keypads are to become feasible.

8. Conclusion

We have presented three studies on the performance of hand postures using a mobile device form factor. These studies examined postures composed of one and two hands, index fingers and thumbs, and the front and back of the device. In addition, the differences between horizontal and vertical movements were examined. Our findings indicate that the index finger works well on the front *and* back of a device, and that the thumb on the front of the device is perhaps less dexterous than we might have hoped. We have also investigated letter orientation on the front and back of the device, showing how the expectation of visual feedback reverses the direction of motor articulation. Furthermore, we showed that users could make letter-like gestures on the front and back of a device, and that the index finger on the front of the device performed best. Taken together, these studies indicate that device designs can go further to leverage a richer set of finger and thumb interactions, and that back-of-device performance should be investigated as a viable means for interaction. The data presented here serves as an empirical foundation for these efforts.

References

Arnaut, L.Y., Greenstein, J.S., 1986. Optimizing the touch tablet: The effects of control-display gain and method of cursor control. Human Factors 28 (6), 717–726.

Balakrishnan, R., MacKenzie, I.S., 1997. Performance differences in the fingers, wrist, and forearm in computer input control. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97). Atlanta, Georgia, March 22–27, 1997. ACM Press, New York, pp. 303–310.

Buxton, W., Myers, B.A., 1986. A study in two-handed input. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '86). Boston, Massachusetts, April 13–17, 1986. ACM Press, New York, pp. 321–326.

Buxton, W., Hill, R., Rowley, P., 1985. Issues and techniques in touch-sensitive tablet input. In: Proceedings of the ACM Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '85). San Francisco, CA, July 1985. ACM Press, New York, pp. 215–223.

Card, S.K., English, W.K., Burr, B.J., 1978. Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. Ergonomics 21 (8), 601–613.

Cheng, P., Buur, J., 2004. Qualities of the past—Telephones of the future. In: Companion to the ACM Symposium on User Interface Software and Technology (UIST '04). Santa Fe, New Mexico October, 24–27, 2004. ACM Press, New York, pp. 7–8.

Crossman, E.R.F.W., 1957. The speed and accuracy of simple hand movements. In: Crossman, E.R.F.W., Seymour, W.D. (Eds.), The Nature and Acquisition of Industrial Skills. Report to the M.R.C. and D.S.I.R. Joint Committee on Individual Efficiency in Industry.

Epps, B.W., Snyder, H.L., Muto, W.H., 1986. Comparison of six cursor devices on a target acquisition task. In: Proceedings of the Society for Information Display, San Diego, CA, May 6–8, 1986. Society for Information Display, San Jose, California, pp. 302–305.

Fitts, P.M., 1954. The information capacity of the human motor system in controlling the amplitude of movement. Journal of Experimental Psychology 47 (6), 381–391.

Frederick, B.N., 1999. Fixed-, random-, and mixed-effects ANOVA models: a user-friendly guide for increasing the generalizability of ANOVA results. In: Thompson, B. (Ed.), Advances in Social Science Methodology. JAI Press, Stamford, Connecticut, pp. 111–122.

Gokturk, M., Sibert, J.L., 1999. An analysis of the index finger as a pointing device. In: Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '99). Pittsburgh, Pennsylvania, May 1999. ACM Press, New York, pp. 286–287.

Goldberg, D., Richardson, C., 1993. Touch-typing with a stylus. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (INTERCHI '93). Amsterdam, The Netherlands, April 24–29, 1993. ACM Press, New York, pp. 80–87.

Guiard, Y., 1987. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. Journal of Motor Behavior 19 (4), 486–517.

Hinckley, K., Czerwinski, M., Sinclair, M., 1998. Interaction and modeling techniques for desktop two-handed input. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '98). San Francisco, California, November 1998. ACM Press, New York, pp. 49–58.

Hiraoka, S., Miyamoto, I., Tomimatsu, K., 2003. Behind Touch: a text input method for mobile phone by the back and tactile sense interface. In: Proceedings of Interaction 2003. Tokyo, Japan, February 27–28, 2003. Information Processing Society of Japan, Tokyo, Japan, pp. 131–138.

Hirotaka, N., 2003. Reassessing current cell phone designs: using thumb input effectively. In: Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '03). Ft. Lauderdale, FL, April 5–10, 2003. ACM Press, New York, pp. 938–939.

Karlson, A.K., Bederson, B.B., SanGiovanni, J., 2005. AppLens and LaunchTile: two designs for one-handed thumb use on small devices. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '05). Portland, Oregon, April 2–7, 2005. ACM Press, New York, pp. 201–210.

Karlson, A.K., Bederson, B., Contreras-Vidal, J.L., 2006. Studies in one-handed mobile design: habit, desire and agility. In: HCIL Technical Report 2006-02. University of Maryland, College Park, MD.

Langolf, G.D., Chaffin, D.B., Foulke, J.A., 1976. An investigation of Fitts' Law using a wide range of movement amplitudes. Journal of Motor Behavior 8 (2), 113–128.

Lin, M., Price, K.J., Goldman, R., Sears, A., Jacko, J., 2005. Tapping on the move—Fitts' law under mobile conditions. In: Proceedings of the 16th Annual Information Resources Management Association International Conference (IRMA '05). San Diego, California, May 15–18, 2005. Idea Group, Hershey, Pennsylvania, pp. 132–135.

Littell, R.C., Milliken, G.A., Stroup, W.W., Wolfinger, R.D., 1996a. SAS System for Mixed Models. SAS Institute, Inc., Cary, North Carolina.

Littell, R.C., Milliken, G.A., Stroup, W.W., Wolfinger, R.D., 1996b. Analysis of repeated measures data. In: SAS System for Mixed Models. SAS Institute, Inc., Cary, North Carolina, pp. 87–134.

Lumsden, J., Brewster, S., 2003. A paradigm shift: alternative interaction techniques for use with mobile and wearable devices. In: Proceedings of the Conference of the IBM Centre for Advanced Studies on collaborative Research (CASCON '03). Toronto, Ontario, October 6–9, 2003. IBM Press, pp. 197–210.

Lyons, K., Starner, T., Plaisted, D., Fusia, J., Lyons, A., Drew, A., Looney, E.W., 2004. Twiddler typing: one-handed chording text entry for mobile phones. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '04). Vienna, Austria, April 2004. ACM Press, New York, pp. 671–678.

MacKenzie, I.S., 1989. A note on the information-theoretic basis for Fitts' law. Journal of Motor Behavior 21, 323–330.

MacKenzie, I.S., 1992. Fitts' law as a research and design tool in human–computer interaction. Human Computer Interaction 7 (1), 91–139.

MacKenzie, I.S., Guiard, Y., 2001. The two-handed desktop interface: are we there yet? In: Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '01). Seattle, Washington, March 31–April 5, 2001. ACM Press, New York, pp. 351–352.

MacKenzie, I.S., Oniszczak, A., 1998. A comparison of three selection techniques for touchpads. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '98). Los Angeles, California, April 1998. ACM Press, New York, pp. 336–343.

MacKenzie, I.S., Soukoreff, R.W., 2003. Card, English, and Burr (1978)—25 years later. In: Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '03). Ft. Lauderdale, Florida, April 5–10, 2003. ACM Press, New York, pp. 760–761.

MacKenzie, I.S., Sellen, A., Buxton, W., 1991. A comparison of input devices in elemental pointing and dragging tasks. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '91). New Orleans, Louisiana, March 1991. ACM Press, New York, pp. 161–166.

Marentakis, G.N., Brewster, S.A., 2006. Effects of feedback, mobility and index of difficulty on deictic spatial audio target acquisition in the horizontal plane. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '06). Montréal, Québec, April 22–27, 2006. ACM Press, New York, pp. 359–368.

Mizobuchi, S., Chignell, M., Newton, D., 2005. Mobile text entry: relationship between walking speed and text input task difficulty. In: Proceedings of ACM Conference on Human–Computer Interaction with Mobile Devices and Services (MobileHCI '05). Salzburg, Austria, September 19–22, 2005. ACM Press, New York, pp. 122–128.

Oulasvirta, A., Tamminen, S., Roto, V., Kuorelahti, J., 2005. Interaction in 4-s bursts: the fragmented nature of attentional resources in mobile HCI. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '05). Portland, Oregon, April 2–7, 2005. ACM Press, New York, pp. 919–928.

Parhi, P., Karlson, A.K., Bederson, B.B., 2006. Target size study for one-handed thumb use on small touchscreen devices. In: Proceedings of the ACM Conference on Human–Computer Interaction with Mobile Devices and Services (MobileHCI '06). Helsinki, Finland, September 12–15, 2006. ACM Press, New York, pp. 203–210.

Parsons, L.M., Shimojo, S., 1987. Perceived spatial organization of cutaneous patterns on surfaces of the human body in various positions. Journal of Experimental Psychology 13 (3), 488–504.

Pirhonen, A., Brewster, S., Holguin, C., 2002. Gestural and audio metaphors as a means of control for mobile devices. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '02). Minneapolis, Minnesota, April 20–25, 2002. ACM Press, New York, pp. 291–298.

Potter, R.L., Weldon, L.J., Shneiderman, B., 1988. Improving the accuracy of touch screens: an experimental evaluation of three strategies. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '88). Washington, DC, May 15–19, 1988. ACM Press, New York, pp. 27–32.

Rekimoto, J., Ishizawa, T., Schwesig, C., Oba, H., 2003. PreSense: interaction techniques for finger sensing input devices. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '03). Vancouver, British Columbia, November 2–5, 2003. ACM Press, New York, pp. 203–212.

Schmidt, R.A., Zelaznik, H., Hawkins, B., Frank, J.S., Quinn, J.T., 1979. Motor-output variability: A theory for the accuracy of rapid motor acts. Psychological Review 86 (5), 415–451.

Schuster, C., von Eye, A., 2001. The relationship of ANOVA models with random effects and repeated measurement designs. Journal of Adolescent Research 16 (2), 205–220.

Schwesig, C., Poupyrev, I., Mori, E., 2004. Gummi: a bendable computer. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '04). Vienna, Austria, April 24–29, 2004. ACM Press, New York, pp. 263–270.

Sears, A., Shneiderman, B., 1991. High precision touchscreens: Design strategies and comparisons with a mouse. International Journal of Man-Machine Studies 34 (4), 593–613.

Silfverberg, M., 2003. Using mobile keypads with limited visual feedback: implications to handheld and wearable devices. In: Proceedings of the 5th Int'l Symposium on Human–Computer Interaction with Mobile Devices and Services (Mobile HCI '03) Udine, Italy, September 8–11, 2003. Springer, Berlin, Germany, pp. 76–90.

Silfverberg, M., MacKenzie, I.S., Korhonen, P., 2000. Predicting text entry speed on mobile phones. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '00). The Hague, The Netherlands, April 1–6, 2000. ACM Press, New York, pp. 9–16.

Silfverberg, M., Korhonen, P., MacKenzie, I.S., 2003. Zooming and panning content on a display screen. International patent WO 03/021568 A1, March 13, 2003.

Silfverberg, M., Korhonen, P., MacKenzie, I.S., 2006. Zooming and panning content on a display screen. United States patent 7075513, July 11, 2006.

Soukoreff, R.W., MacKenzie, I.S., 2004. Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. International Journal of Human–Computer Studies 61 (6), 751–789.

Sugimoto, M., Hiroki, K., 2006. HybridTouch: an intuitive manipulation technique for PDAs using their front and rear surfaces. In: Proceedings of the ACM Conference on Human–Computer Interaction with Mobile Devices and Services (MobileHCI '06). Helsinki, Finland, September 12–15, 2006. ACM Press, New York, pp. 137–140.

Vermunt, J.K., 1997. Log-linear Models for Event Histories. Sage Publications, Thousand Oaks, California.

Vogel, D., Baudisch, P., 2007. Shift: a technique for operating pen-based interfaces using touch. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '07). San Jose, California, April 28–May 3, 2007. ACM Press, New York, pp. 657–666.

Welford, A.T., 1968. Fundamentals of Skill. Methuen, London.

Wigdor, D., Balakrishnan, R., 2004. A comparison of consecutive and concurrent input text entry techniques for mobile phones. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '04). Vienna, Austria, April 2004. ACM Press, New York, pp. 81–88.

Wigdor, D., Leigh, D., Forlines, C., Shipman, S., Barnwell, J., Balakrishnan, R., Shen, C., 2006. Under the table interaction. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '06). Montreux, Switzerland, October 15–18, 2006. ACM Press, New York, pp. 259–268.

Wigdor, D., Forlines, C., Baudisch, P., Barnwell, J., Shen, C., 2007. LucidTouch: a see-through mobile device. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '07). Newport, Rhode Island, October 7–10, 2007. ACM Press, New York, pp. 269–278.

Wobbrock, J.O., 2003. The benefits of physical edges in gesture-making: empirical support for an edge-based unistroke alphabet. In: Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '03). Ft. Lauderdale, Florida, April 5–10, 2003. ACM Press, New York, pp. 942–943.

Wobbrock, J.O., 2007. Measures of text entry performance. In: MacKenzie, I.S., Tanaka-Ishii, K. (Eds.), Text Entry Systems: Mobility, Accessibility, Universality. Morgan Kaufmann, San Francisco, pp. 47–74.

Wobbrock, J.O., Forlizzi, J., Hudson, S.E., Myers, B.A., 2002. WebThumb: interaction techniques for small-screen browsers. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '02). Paris, France, October 27–30, 2002. ACM Press, New York, pp. 205–208.

Wobbrock, J.O., Myers, B.A., Kembel, J.A., 2003. EdgeWrite: a stylus-based text entry method designed for high accuracy and stability of motion. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '03). Vancouver, British Columbia, November 2–5, 2003. ACM Press, New York, pp. 61–70.

Wobbrock, J.O., Myers, B.A., Aung, H.H., LoPresti, E.F., 2004. Text entry from power wheelchairs: EdgeWrite for joysticks and touchpads. In: Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '04). Atlanta, Georgia, October 18–20, 2004. ACM Press, New York, pp. 110–117.

Wobbrock, J.O., Chau, D.H., Myers, B.A., 2007. An alternative to push, press, and tap-tap-tap: gesturing on an isometric joystick for mobile phone text entry. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '07). San Jose, California, April 28–May 3, 2007. ACM Press, New York, pp. 667–676.

Zhai, S., Milgram, P., Buxton, W., 1996. The influence of muscle groups on performance of multiple degree-of-freedom input. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '96). Vancouver, British Columbia, April 1996. ACM Press, New York, pp. 308–315.

Zhai, S., Kong, J., Ren, X., 2004. Speed–accuracy tradeoff in Fitts' law tasks—on the equivalence of actual and nominal pointing precision. International Journal of Human–Computer Studies 61 (6), 823–856.