Pressure Marks

Gonzalo Ramos, Ravin Balakrishnan
Department of Computer Science
University of Toronto
www.dgp.toronto.edu
bonzo, ravin@dgp.toronto.edu

ABSTRACT
Selections and actions in GUI's are often separated – i.e. an action or command typically *follows* a selection. This sequence imposes a lower bound on the interaction time that is equal to or greater than the sum of its parts. In this paper, we introduce *pressure marks* – pen strokes where the variations in pressure make it possible to indicate both a selection and an action simultaneously. We propose a series of design guidelines from which we develop a set of four basic types of pressure marks. We first assess the viability of this set through an exploratory study that looks at the way users draw straight and lasso pressure marks of different sizes and orientations. We then present the results of a quantitative experiment that shows that users perform faster selection-action interactions with pressure marks than with a combination of lassos and pigtails. Based on these results, we present and discuss a number of interaction designs that incorporate pressure marks.

ACM Classification: H5.2 [User Interfaces]: Interaction styles.

General terms: Design, Human Factors

Keywords: pen input, pressure widgets.

INTRODUCTION
In pen-based interfaces, the choice of a command generally follows the selection of a group of objects. Despite a variety of instantiations, this *selection-action* [8] interaction is typically serial in nature, i.e., the scope and command specifications occur in sequence, one after the other. This sequential nature not only makes the time necessary to perform a selection-action interaction at least equal to the sum of the time it takes to execute its parts, but it can also impose a sequential structure on potentially concurrent or integrated tasks [10]. Moreover, the nature of the delimiter between the selection and action components of the interaction can further increase the overall execution time.

In this paper, we introduce and investigate *pressure marks* (Figure 1), pen strokes where the variations in pressure a user applies while drawing, or *pressure signature*, has meaning. These marks can potentially improve selection-action interactions by allowing the selection and action to be specified *concurrently*. In addition, they have the potential benefit of providing pen gestures that are orientation independent, i.e., a command corresponds to a particular pressure signature (signature for short), rather than to a stroke direction or orientation. Thus, pressure marks can be useful in scenarios where the user's orientation relative to the display varies. Examples of such scenarios include cases where an artist rotates the underlying drawing canvas, or when one or more users interact with the same surface from different directions as commonly occurs in co-located collaborative tabletop environments. Finally, pressure marks can be used to enhance traditional marking techniques, such as pie or marking menus [12], which are most effective when the relative orientation of the screen and the user remains constant. For example, if one is to use a stroke's direction as part of a command, as with marking menus, pressure marks can potentially increase the number of available commands at a user's disposal.

Despite their potential, pressure marks present interesting design challenges. For example, in order for novice users to browse through a given set of available pressure marks, one has to produce visual designs that prompt them to interact through both pressure and x-y spatial movements.

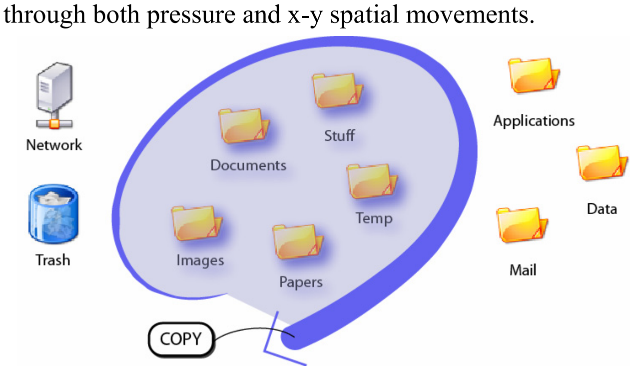

Figure 1: A pressure mark is used to select and copy (e.g., to a clipboard) a group of items in a single stroke. The selection is indicated by the enclosure of the stroke, while the command is specified by the pressure signature (thin-THICK) over the stroke. The selection and action components occur concurrently, with no delimiters between them.

While browsing is straightforward when the interaction only uses x-y space, this is not the case when both x-y and pressure spaces are used concurrently. For example, items in a pie menu are laid out in different x-y positions making it easy for the user to glance at them. However, different pressure marks can have the same x-y spatial movement, making a visual representation that permits browsing non-trivial since overlapping icons may be confusing. Providing an effective browsing experience for pressure marks is key both for assisting novice users in becoming experts and for allowing the exploration and discovery of otherwise hidden commands. However, it is not obvious what the basic set of pressure marks should be, and how well users will perform when using such a set. Complex marks can potentially expand the interaction vocabulary, but they can also impose a steep learning curve and lead to high error rates.

After reflecting on related work, we formally present pressure marks and proceed to investigate their viability through an exploratory study that looks at the way people input different types of pressure marks. Encouraged by our initial results, we conducted a second study that compares pressure marks – a concurrent selection-action technique – with lassoing + pigtailing [8] – a sequential selection-action technique. We then offer a number of designs that leverage pressure marks. Finally, we outline improvements and future work in the area.

PREVIOUS WORK

Pen-based systems generally offer users a *marking* and an *inking* mode. For the purposes of this paper, we are assuming that non-dominant hand mode indication is present as researched by Li et al. [14] and have chosen to concentrate on pen interactions while in marking mode. While selection-action patterns are traditionally sequential, there have been efforts to improve on this experience. Guimbretière et al.'s FlowMenus [7] fluidly connect command selection and direct manipulation. FlowMenus consist of eight octants arranged around a central rest area. By entering and/or leaving the rest area and moving through the octants, a user can, in one fluid motion, navigate through the menu hierarchy, adjust a parameter or manipulate an object on a display. Similarly, Hinckley et al.'s pigtail delimiters [8] allow selection-action patterns to be performed in one continuous fluid stroke. A user explicitly creates a pigtail by intersecting his/her own stroke and then uses the stroke's direction to specify an action or manipulate an object. Pigtails provide a way to integrate an explicit command invocation in a fluid stroke following the selection specification. This is unlike previous selection-action schemes where users signal a command selector using buttons or timeouts. Saund and Lank [19] present a technique that guesses the user's intent by using the stylus' trajectory and context. While in some cases this protocol does not need an explicit command, the system presents a selector widget if the stroke drawn is ambiguous. While fluid, these techniques remain sequential – i.e., a delimiter separates selection and action.

Previous research suggests that interactions where parallelism occurs can outperform sequential tasks. For example, researchers have shown how bi-manual interaction techniques permit parallelism [4] as well as outperforming [13] one-handed ones. Baudisch et al.'s marquee menus [2] are a technique where the selection-action pattern occurs concurrently. The marquee menu's selection is specified by the rectangular area defined by a straight stroke and its action is determined by drawing the stroke in one of four directions. While providing a compact interaction phrase, this type of menu is sensitive to both a mark's point of origin and direction. Although promising, the authors did not elaborate on if and how this technique scales for non-straight strokes with arbitrary orientations, or for larger command sets.

The control of pressure has the potential to be used concurrently with spatial movement. Srinivasan and Chen [21] presented evidence that when provided appropriate visual feedback, users can control variations of pressure over time. In their study, participants applied force to a pressure sensor using the pad of their index finger. Their results however, do not include situations where the applied force changes as the user's finger slides over a rigid surface. With Zliding [16], Ramos and Balakrishnan explore integrated panning and zooming by concurrently controlling input pressure while sliding in x-y space. While this work provides insight on the issue, the authors did not study the possibility of using integrated spatial movement and pressure input for concurrent selection-action operations.

While the use of marks for both selection and action patterns in the GUI is not new, there has not been significant examination of marks capable of concurrent selection-action. Similarly, while there have been efforts supporting the use of an interface from different orientations, they often assume a system that senses a user's position, or has preordained rules of engagement – i.e., regions of the display are meant to be approached from a particular direction [6, 20, 22]. Kara and Stahovich [11] use polar coordinates to recognize gestures independently of their orientation in the context of their Simusketch system. However, some gestures such as the marks in hierarchical marking menus are ambiguous when seen from different orientations. We believe that the exploration of interactions that do not depend on how the user approaches the interface deserves further exploration. This exploration has the potential to benefit the usability of both portable small displays such as those on PDA's and larger form factors such as tabletop displays.

PRESSURE MARKS

Through history, people have used hand-made marks both as a channel of artistic expression and as a way to encode information, i.e. writing. These marks can be made by someone's bare hands and with instruments such as chisels, quills, brushes and pens. They are a testament to the fine pressure control achievable by the human hand. However, people's skills with a stylus can vary dramatically from one person to the next. To be successful it is thus important to design a simple and significant set of signatures that people can perform effectively.

Writing or drawing a predetermined stroke, such as a pressure mark, can be described as an open-loop task, as it was early reported by Woodworth and reviewed by Elliot et al. [5]. In open loop tasks, people develop a motor program, or plan, and later execute it without using sensory feedback. Open-loop tasks can be very fast once people start them, so fast that their sensory system cannot keep up with all the events that occur during it. Because of this perceptual limitation, once the task execution has started it is very difficult to modify the task's goal. However, a person's sensory feedback can affect the task at the planning stage or during execution if the speed at which the task is reduced to below a certain threshold. Building on this information and acknowledging that people's skills vary, we use a set of design guidelines for pressure marks:

Small and Simple: we will aim for a small number of pressure signatures. In addition, signatures should be simple enough so an average user can do them.

Continuous Feedback: appropriate visual feedback should enable users to be aware of what they are doing and the effect pressure has on the interaction.

Browsable: a mechanism for browsing through the available pressure marks should be available.

We initially consider two basic signatures when a mark is drawn: a) pressure remains constant within a certain margin of variance, and b) pressure changes. We can then increase the number of signatures by a factor of two by considering constant pressure at low or high levels, and pressure changing in a monotonically increasing or decreasing way. This process produces a set of four basic signatures (Figure 2): low-low (LL), low-high (LH), high-high (HH) and high-low (HL). While it is possible for users to draw pressure marks with signatures that do not match exactly either of these four classes, we will later describe a way in which we can reduce any pressure mark into one of these four classes.

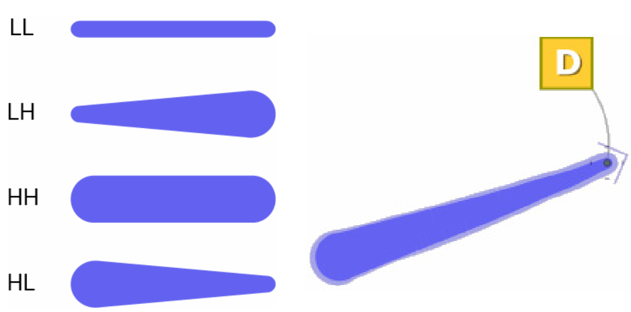

Figure 2: (left) Profiles of the four proposed pressure signatures. (right) User-made stroke with recognized profile overlaid on top. The icon "D" represents some action associated with the HL profile.

In order to give meaningful information to users as they create a mark, we use different types of feedback. First, we use a pressure cursor similar to the one described by Ramos et al. [16] in order to provide continuous information about the pressure being applied through the stylus. Second, we draw a stroke whose width at any point is proportional to the amount of pressure applied at that point. This last feedback is analogous to strokes found in Japanese sumi-e paintings. Finally, when the user lifts the stylus from the display's surface (i.e., the mark is finished) a stylized semi-transparent representation of the pressure mark is overlaid on it, confirming the interaction (Figure 2). We call this stylized visual representation of a pressure mark its *profile*. We use a drawing pause-timeout delimiter to enter a *browsing mode* where novice users can browse through the available set of pressure marks. Many selection-action scenarios use delimiters in this way and in our case, it reverts pressure marks into a standard sequential pattern.

Browsing Through Pressure Marks

During the early stages of our research, we explored different designs for this browsing mode. We first considered a design where available profiles of the pressure marks appear at different locations in x-y space. While this design has the advantage of showing at-a-glance all the available options, pilot studies showed that it also elicits an unwanted response from users in that users often drew marks towards the different profiles as if they were targets in a spatial menu. People exposed to this browsing mode were confused, as they were not sure what it meant and they often treated it as a spatial menu. In the end, this type of design did not facilitate the exploration of available pressure marks. In light of the unwanted response we got from users to these "spatial browsing" designs, we developed an alternate mechanism for browsing pressure marks.

The browsing mechanism we propose exploits the property that each of the four proposed pressure signatures can be connected like domino pieces – e.g., a user can "draw a stroke" that passes through all four signatures: … LL → LH → HH → HL → LL … These "connections" allows us to suggest follow-up marks when the browsing mode is active (Figure 3).

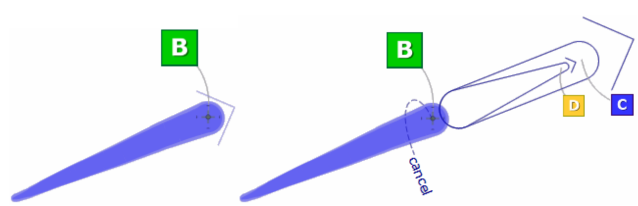

Figure 3: Browsing through pressure marks. A pause-timeout delimiter while marking (left) triggers a browsing mode. We suggest available pressure profiles from that point (right). The boxed letters represent actions associated with the profile they are connected to. The dashed line (right) indicates that pigtailing at that point cancels the mark.

In particular, our implementation draws profiles of signatures that the user can execute next and that differ from the mark currently in progress. These suggested profiles follow the last known direction of the mark in progress. In addition to drawing the profiles of potentially different marks, we draw a dashed line that indicates an available "cancel" crossing gesture. This cancel gesture is simply a self-intersecting pigtail at the end of the mark.

This browsing mechanism follows a quasi-sequential way of exploring the set of available pressure signatures. While spatial browsing can elicit unwanted responses from users, sequential browsing can be tedious and time consuming. We believe that our proposed type of quasi-sequential access is a good compromise between the random access that full spatial browsing provides, and purely sequential exploration. This way of revealing otherwise hidden gestures pushes Scriboli's stroke extension [8] further and addresses the issue of revealing more than one available gesture. In contrast to the original "spatial layout" approaches we considered, informal user feedback on our quasi-sequential browsing mode has been very positive. The question of how useful this browsing mechanism is for first-time users merits exploration, but is outside the scope of this paper.

Pressure Marks' Anatomy: Reduction and Parsing

The four basic types of pressure marks we propose do not capture all possible signatures that a user can make. What do we do if a mark's signature does not fit our prescribed types? There are three options: a) treat the mark as null, as if nothing happened; b) treat the mark as ambiguous and solicit additional feedback to resolve the ambiguity; and c) reduce the mark into one of the four types. After conducting preliminary heuristic evaluations at the early stages of our research, we decided in favor of the third option because it allows us to provide a response that better matches users' expectations. For example, users who made a mark wherein pressure first increases and later decreases often expected the system to recognize a HL signature.

Analysis of a pressure mark consists of both detecting the selection and the action that it encodes. The selection part of a mark, which determines which UI object(s) are selected, depends on the mark's shape – e.g., straight line, lasso, etc. As this is a well-studied topic [8, 9, 15], we will concentrate instead on how signatures are parsed and reduced. We performed parsing over the curve defined by the distance traveled by the stylus and the pressure applied through it. Figure 4 shows an example of such a curve.

We first discard several data points at the ends of the curve to account for noise from the stylus engaging with and disengaging from the digitizing surface. We later analyze the resulting curve and simplify it using a piecewise linear approximation. This scheme tries to fit a straight line on the curve, and then, if the error is above a certain threshold, it divides the curve at the point of maximum error. Later, the algorithm recursively finds a linear fit to those pieces. The analysis stops when it meets a convergence threshold.

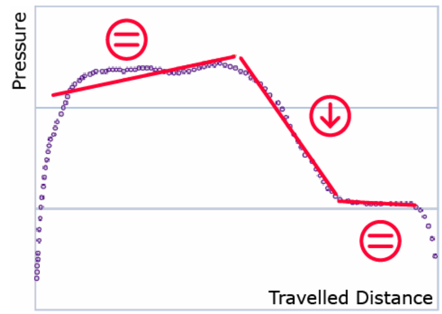

Figure 4: Pressure vs. traveled distance. The red segments show fitted lines on the pressure curve. In this case, the curve is parsed as a HL (↓) mark.

For a given signature, this analysis produces a set of n straight lines; each with characteristic features such as length, % length, slope, and pressure change. We use these features to label each line as being constant (=), ascending (↑) or descending (↓). The labeling takes into consideration how people perceive variations in pressure at different levels, a phenomenon that can be described by the Weber-Fechner Law [3]. Finally, we classify a signature according to its last observed ↑ or ↓ label as LH or HL, respectively. If we observe no ↑ or ↓ labels, we use a simple threshold to classify the signature as LL or HH.

Further details of the reduction-parsing algorithm and labeling process go beyond the scope of this paper. However, our code is available upon request to researchers who desire to inspect or use it.

While straightforward, our reduction-parsing technique does not consider variability among different users. Though we are not considering this issue at this point, one can imagine using other parsing schemes that utilize a training set of gestures, such as Rubine's [18].

It might be interesting in the future to explore using the speed or tilt variations of the pen for command specification, however these properties may be less suitable than pressure since speed is highly variable across users and tilt is orientation dependent.

EXPLORATORY USER STUDY

It is important to determine if users can successfully execute pressure marks using the proposed set of signatures. We also wish to assess how well our reduction-parsing technique performs. Favorable answers to these issues will provide evidence in support of pursuing interaction designs that leverage the use of pressure marks. In order to answer these questions, as well as to gain usability information about pressure marks in general, we performed an exploratory quantitative study, which we describe in this section.

Apparatus

We used a Toshiba Portégé M200 TabletPC running Windows XP Tablet Edition, with a 1400 by 1050 pixel display at ~140 dpi. Participants used the TabletPC in slate mode and interacted with it using a wireless stylus that has a pressure-sensitive tip. The tablet was set flat on a desk, but we allowed participants to adjust the tablet's position on the desk, according to their preference.

Participants

16 people (11 male, 5 female), 18-44 years old, recruited from our university population through e-mail lists, participated in the study. One was left handed. All had some familiarity with the TabletPC. We provided no compensation.

Task and Stimuli

We studied two styles of marks: straight lines and lassos. For each mark's style, we also considered three possible lengths and various different orientations. For the straight line case, we asked participants to draw lines of different *length*: *small, medium* and *large* in four possible *orientations*: *east (E), west (W), south (S)* and *north (N)*. For the lasso case we also asked participants to draw lassos of different *length*: *small, medium* and *large*, but in eight possible *orientations*: *clockwise north (CN)*, *south (CS)*, *east (CE), west* (*CW*) and *counterclockwise north* (*CCN*), *south (CCS)*, *east (CCE)*, and *west (CCW)*.

For each experimental trial, we presented as stimulus a stylized representation of the pressure mark users should make and dashed lines that users should cross in order to complete marks of a particular length. Figure 5 illustrates the stimulus for straight lines and lassos. In the case of lassos, we also showed a gray circle that indicated the particular object that users had to lasso. Start and stop icons, showing where a mark should start and end respectively, reinforced the task's orientation.

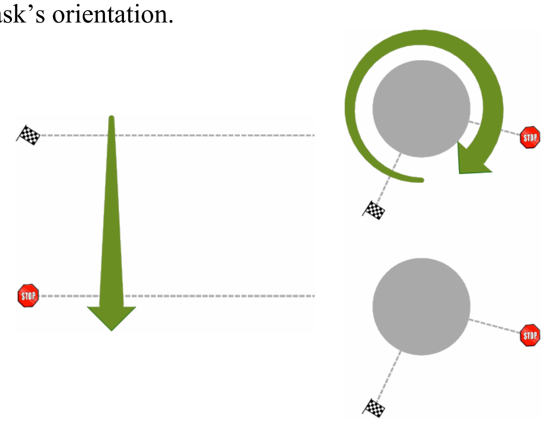

Figure 5: An example of experimental stimuli for the straight line (left) and lasso (right) cases.

Each trial ran the same way for both the straight line and lasso cases. After the stimulus was displayed, users were required to draw a pressure mark with a profile and trajectory similar to the one displayed. After the mark was completed, we tested for two conditions: a) that the mark's trajectory is similar to the stimulus' – i.e., dashed lines are crossed in the right order; and b) that the mark drawn by the user is parsed as one with a similar signature as the stimulus. While failing either test causes an error sound to be played and the trial to be repeated; only failing b) was counted as an error since our primary goal was to determine user ability to generate the given pressure signatures and not their ability to draw straight lines or lassos per se. The mark's browsing mode was disabled for this study.

Procedure and Design

For straight marks, we used a 3 length (*small – 2cm, medium – 4cm, large – 8cm*) × 4 orientation (*N, S, E, W*) × 4 signature (*LL, LH, HH, HL*) within-subjects design. For lasso marks, we used a 3 diameter (*small –2.5cm, medium – 4cm, large – 6.5cm*) × 8 orientation (*CN, CS, CE, CW, CCN, CCS, CCE, CCW*) × 4 signature (*LL, LH, HH, HL*) within-subjects design. For both types of marks, the dependent variables were *trial time* and *errors*. We computed trial time as the time elapsed between the moment a participant touches the tablet's surface with the stylus after a trial's stimulus was presented and the trial's successful completion. A trial was erroneous if the system could not match the user's input with the stimulus. Since one could only advance to the next trial after being successful in the preceding, participants were motivated to perform well. For straight lines, participants completed three blocks of trials. Each block consisted of 48 marking tasks repeated twice. Presentation of trials within a block was randomized. For lasso marks, participants completed two blocks of trials. Each block consisted of 96 marking tasks repeated twice. Again, presentation of trials within a block was randomized. All users did the straight line case first, followed by lassos. In summary, the experiment consisted of:

16 participants × ((3 blocks × 3 lengths × 4 orientations × 4 straight signature marks × 2 repetitions) + (2 blocks × 3 lengths × 8 orientations × 4 lasso signature marks × 2 repetitions)) = 10752 trials.

Prior to the study, the experimenter explained the task to the participants. Before each type of mark was presented, participants practiced with two warm-up blocks of 48 trials. The experimenter then told participants to do the trials as quickly and accurately as possible. Participants completed a questionnaire at the end of the experiment.

Results

This study averaged 1 hour per participant. For the straight lines case, we conducted a 3 (block) × 3 (length) × 4 (signature) repeated measures analysis of variance (RM-ANOVA) on the logarithmically transformed trial time and on the errors. For the lasso case, we conducted a 2 (block) × 3 (length) × 4 (signature) RM-ANOVA, also on the logarithmically transformed trial time and on the errors. The logarithm transform corrects for the skewing present in human response data, and removes the influence of outliers.

Errors

In this study, we are most interested in error rates across blocks of trials because they will give us information as to learning effects, as well as an indication of the performance of our reduction-parsing algorithm. Figure 6 illustrates how, for both straight lines and lassos, error rates decrease as the study progresses until it reaches levels of about 4%.

There was no main effect for blocks in the non-warmup experimental trials, both for straight lines ($F_{2,30}$=2.006, p=0.152) and for lassos ($F_{1,15}$=3.99, p=0.064). However, when we examine this data in the context of the warm-up data, we observe a marked improvement in the users' accuracy as they progress through the entire experiment.

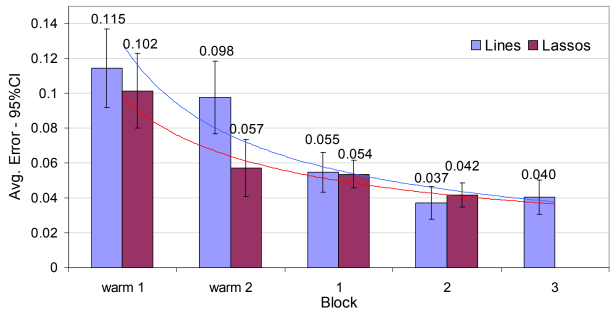

Figure 6: Average errors per block for straight lines and lassos. Power regression lines suggest the presence of learning taking place.

We found a main effect for signature in the case of straight lines ($F_{3,45}$=3.391, p=0.026) and lassos ($F_{3,45}$=8.82, p<0.0001) (Figure 7). Participants were most accurate when performing a LH line or lasso. Unlike the case of straight lines where accuracy was similar for LL, HH and HL, with lassos people made more mistakes when trying to maintain constant pressure, especially the HH mark. This is consistent with observations found in the literature [16, 17].

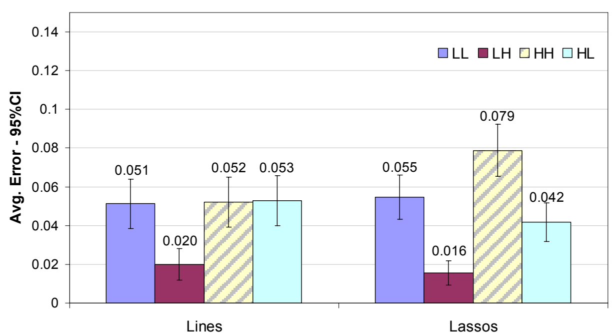

Figure 7: Average errors per pressure signature for lines and lassos.

Trial Times

For straight lines, we found a main effect for blocks ($F_{2,30}$=20.9, p<0.0001), where users performed faster as the trials progressed. For lassos, there were no main effects for blocks ($F_{1,15}$=0.233, p=0.636). This was probably due to users' prior exposure to the straight lines case. However, as Figure 8 illustrates, we still observe a small speed improvement. We also found a main effect for signature in the case of straight lines ($F_{3,45}$=31.023, p<0.0001) and for lassos ($F_{3,45}$=16.784, p<0.0001). Bonferroni-corrected pairwise comparisons reveal significant differences between all signatures (p<0.003), except between LH and HH (p=1.0) for straight lines. A similar trend was seen for lassos where all pairs of signatures were significantly different (p<0.05) except for LH and HH (p=0.746). Figure 9 illustrates this.

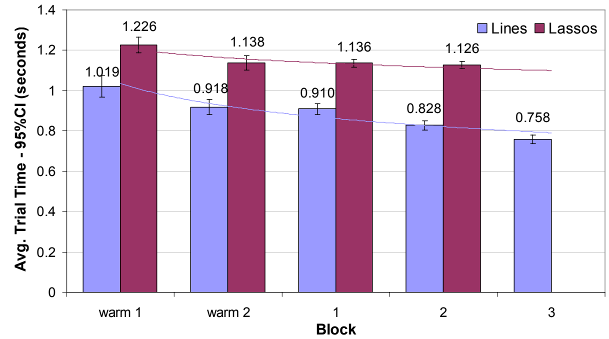

Figure 8: Average Trial Time per block for lines and lassos. Power regression lines are shown.

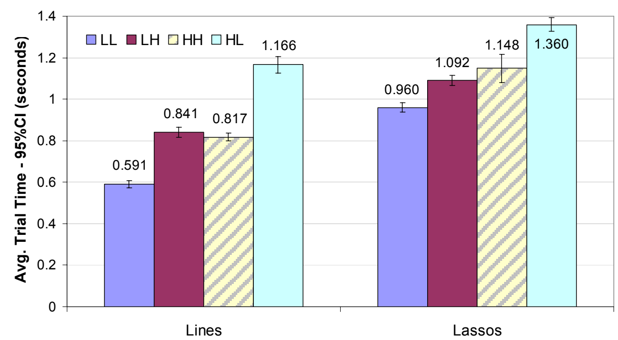

Figure 9: Average Trial Time per pressure signature for lines and lassos.

We can explain these differences by reflecting on the nature of the pressure profiles of the four signatures. LL was the fastest. It requires users to keep pressure relatively constant once the tip of the stylus touches the screen. Users commented on how easy it was to perform LH marks. We observed how they did them in an almost ballistic way. Once the stylus touches the screen, users start increasing pressure as they drag the stylus on the screen's surface. It took users almost the same time to do HH marks as to make LH marks. For the HH case, users reached a high level of pressure in a ballistic way, before dragging the stylus. Our observations showed that when drawing lassos keeping pressure constant was more challenging than it was when drawing straight lines. While users took the longest do HL, they did achieve good levels of accuracy. This is consistent with general user feedback where users describe the HL mark as the most difficult to perform of the four pressure signatures.

Qualitative results are consistent with our experimental observations. Users rated in a scale from 1 (strongly disagree) to 7 (strongly agree) on the ease of use of the different marks. On average, LH (6.3) was rated between agree and strongly agree; HH (5.4) and LL (5.2) between somewhat agree and agree; and HL (4.6) rated least easy, between neither agree/disagree and somewhat agree.

As we might expect from Fitts' law, there was a main effect for length for lines ($F_{2,30}$=191.733, p<0.0001), as well as for lassos ($F_{2,30}$=212.097, p<0.0001).

Summary

The results from this study provide encouraging evidence supporting pressure marks as a viable interaction technique – i.e., people both can learn how to use them and are able to perform them accurately. Moreover, the initial set of pressure signatures we propose seems adequate, albeit with improvements required to our reduction-parsing algorithm. In particular, our heuristics seem to be sensitive to pressure variations while doing a LL or HH mark.

CONTRAST USER STUDY

A concurrent selection-action technique like pressure marks has the potential to produce faster interactions than sequential techniques. However, it is not clear whether the increased complexity inherent in concurrently controlling both pressure and spatial x-y positioning would negate the benefits of concurrency. Therefore, we wanted to gather data as to the performance of pressure marks in comparison to a fluid serial selection-action technique. Accordingly, we ran a study that delves further into the use of pressure marks and contrasts its performance with lassoing + pigtail2 (LP2) [8], one of the latest state-of-the-art fluid serial selection-action technique available to date.

Apparatus and Participants

For this contrast study, we used the same apparatus as in the first study. 14 people (9 male, 5 female), 18-44 years old, recruited from our university population through e-mail lists, participated in the study. None of these 14 people participated in the first study. No compensation was provided.

Task and Stimuli

For this study we used an experimental task similar to the one used by Hinckley et al. [8]. Users were asked to lasso (i.e., select) elements in a selection region and apply the correct action to the selected elements using either a pigtail menu or a concurrent pressure mark. The selection region consisted of 9 squares arranged in a 3 x 3 grid (Figure 10). The squares' size and spacing in our study were chosen to match the experimental setup in Hinckley et al. [8].

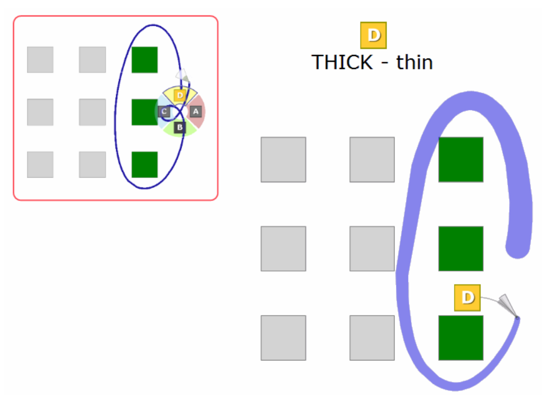

Figure 10: a) Example of an experimental trial for the pressure marks condition. b) Thumbnail of similar trial for the LP2 condition –not shown in the study.

For each trial, we highlighted the squares to be selected in bright green and indicated the action to be taken by displaying the word "North", "South", "East" or "West" for LP2 and "thin-thin", "thin-THICK", "THICK-THICK" or "THICK-thin" for pressure marks. We chose this type of text stimuli instead of a graphic one, because we did not want to impose on users any prescribed way to lasso the targets. Also we showed above the text stimuli the icon that corresponded to the action users had to apply. In our study, the icons were the letters "A", "B", "C" and "D" framed inside a colored box (Figure 10). After the stimulus was displayed, users were required to lasso the green squares and to indicate the requested action. After users completed the selection-action pattern, we tested for two conditions: a) that the lasso included all of the green squares and no distractors; and b) that the action performed matched the one presented as stimulus. While failing either test caused an error sound to be played and the trial to be repeated; only failing condition b) was counted as an error since our primary goal was to study user ability to lasso some number of targets and specify a command, rather than their ability to perfectly lasso a given number of targets per se. A target was inside a lasso if the target's center was inside it. The mark's browsing mode was disabled for this study.

Procedure and Design

We used a 2 technique (*pressure mark, lasso+pigtal2*) × 2 selection type (*single, multiple*) × 6 selection × 4 action (*N*, *S*, *E*, *W* for LP2 and *LL*, *LH, HH, HL* for pressure marks) within-subjects design. For *multiple* selection tasks, the selection was always 3 contiguous squares randomly selected as a row or column [8].

The dependent variables were *trial time* and *error*. Trial time was the time between the moment the stylus touched the tablet's surface after a trial's stimulus was presented and the trial's successful completion. A trial was erroneous if the user lassoed the targets, but performed an incorrect action. Since one could only advance to the next trial after completing the preceding one, participants were motivated to perform well. We divided participants in two groups, according to the order in which techniques were presented to them (pressure marks first or LP2 first). This order was included as a between-subjects factor. For each technique, we asked participants to complete three blocks of trials. Each block consisted of 48 selection-action tasks repeated twice. Presentation of trials within a block was randomized. In summary, the study consisted of:

14 participants × 2 techniques × 3 blocks × 2 selection types × 4 lasso signature marks × 6 tasks × 2 repetitions = 8064 trials.

Prior to the first use of a technique, we explained to participants the nature of the task. Participants practiced with two warm-up blocks of 48 trials. We also instructed participants to be as quick and accurate as possible.

Results

This study averaged 1 hour per participant. We conducted a 2 (technique) × 2 (block) × 2 (selection type) RM-ANOVA on the logarithmically transformed trial times and on the errors. The logarithm transform corrects for the skewing often present in human response data, and removes the influence of outliers. The presentation order of the techniques had no effects on the trial times or the errors.

Trial Times

There was a main effect for technique ($F_{1,11}$=18.22, p<0.001), with pressure marks being an average of 320 msec (27%) faster that LP2. As expected, selection type ($F_{1,11}$=396.38, p<0.0001) had a significant effect – i.e., it takes longer to lasso a larger target. Post-hoc pairwise comparisons show significant differences between technique for both single (p=0.001) and multiple (p=0.002) selections. Pressure marks were consistently faster (Figure 11).

There was a main effect for block ($F_{2,22}$=13.644, p<0.0001), and a marginal technique*block interaction ($F_{2,22}$=3.338, p=0.054). Average trial times improved for both techniques as the study progressed. However, trial times decreased more drastically for the LP2 condition. Bonferroni-corrected post-hoc comparisons show no significant differences between the last two experimental blocks for either LP2 or pressure marks; and also reveal that the difference between techniques at the last block is still significant (p=0.001). Figure 12 illustrates these results.

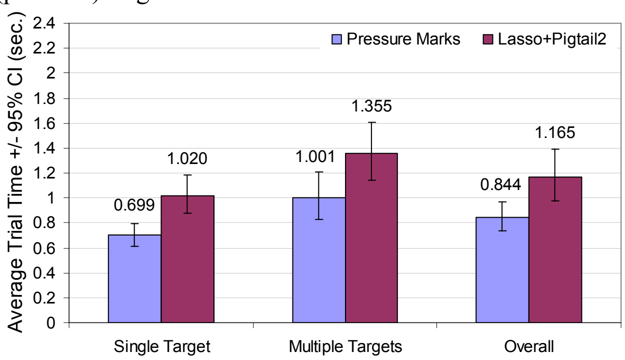

Figure 11: Average trial time per technique and selection type.

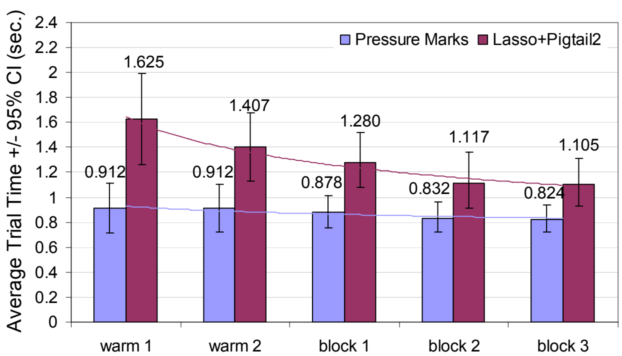

Figure 12: Average trial time per block. Power regression lines are shown.

Errors

We found no significant effects for technique ($F_{1,11}$=0.294, p=0.598) on errors (Figure 13). While participants made slightly fewer errors with LP2 when selecting multiple targets, Bonferroni-corrected post-hoc comparisons show that this difference was not significant (p=0.226).

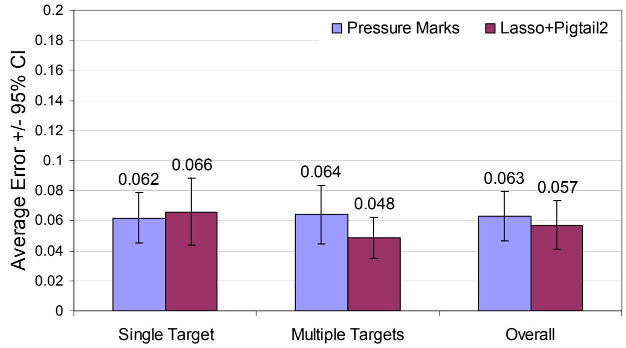

Figure 13: Average error per technique and selection type.

While we saw some improvement in the users' accuracy as the study progressed, we did not found effects for blocks ($F_{2,22}$=0.072, p=0.930) or technique*block ($F_{2,22}$=0.143, p=0.868). Figure 14 illustrates these results.

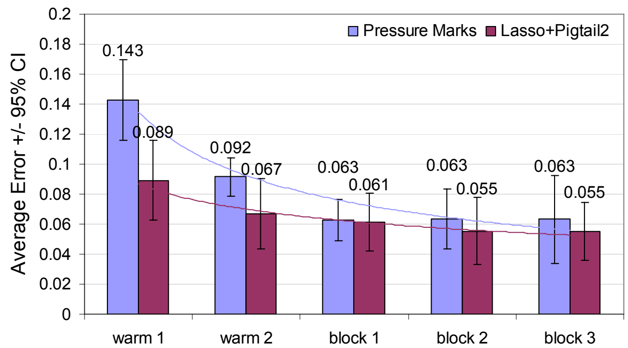

Figure 14: Average error per block per technique. Power regression lines are shown.

Subjective Observations

Many participants reported how they developed different strategies for performing pigtails efficiently, depending on the (action) stimulus. These strategies usually involved starting the lasso from a particular point relative to the target(s) and doing a clockwise or counterclockwise motion. We also observed that participants had a preferred starting position and orientation that stayed almost unchanged when using pressure marks. Participants liked not having to think about the direction and orientation of a lasso.

ON THE USE OF PRESSURE MARKS

Our studies provide us with evidence in support of pressure marks as a viable interaction technique whose ability to specify both selection and action concurrently outperforms existing techniques that require these operations be performed in a sequentially. This evidence encourages us to explore different designs where we can leverage the use and properties of pressure marks.

Pressure Marking Menus

Ramos et al. suggested the idea of pressure marking menus [17], but they did not elaborate upon it. This idea considers both the direction and the signature of a mark to increase the number of items available at any particular level or depth of a marking menu. With our proposed set of signatures, a menu's breadth can increase by a factor of four. While a single-level menu has a straightforward design, the visual design for novice users becomes challenging for menu depths greater than one. Whereas expert users will move through the marking menu's levels by changing a mark's inflexion, novice users need a visual design to help them browse through the menu's options at the current and sometimes at the next levels. Figure 15 shows a design that aims to address this issue.

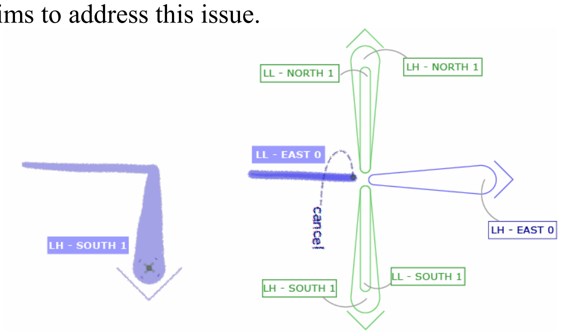

Figure 15: Pressure marking menu design. (left) Expert mark. (right) Feedback shown in the novice / browsing mode. The labels indicate the mark/level.

If we detect a pause-timeout delimiter while a stroke is drawn, we enter a browsing mode that shows available options for the menu's next level that are not collinear with the current mark. We use the direction of the current mark for browsing within the current level. Finally, users can pigtail on the mark to go back to the previous menu level.

Simple Pressure Marks

We can leverage pressure marks to expand other marking schemes such as Zhao and Balakrishnan's simple marks technique [24]. As with compound marks, our proposed signatures can increase simple mark's breadth by a factor of four. However, we think that it is interesting to explore a variation of the simple marking scheme, one that is orientation invariant. We call this design *simple pressure marks*. While traditional marking schemes rely on the presence of a "north" direction, simple pressure marks do not. An orientation-invariant marking scheme can be advantageous in situations when users engage an interactive surface from an arbitrary orientation – e.g., when an artist draws on a sheet of e-paper, or at collaborative tabletop environments.

With simple pressure marks, users specify a command by concatenating pressure signatures made in any direction (Figure 16). This defines a menu structure with four choices per level. We argue that this type of "arbitrary flicking" makes connecting marks easy and independent from a user's handedness, screen layout or orientation. For example, preliminary user observations revealed that users tend to develop a zigzag flicking pattern, which varies in orientation depending on the user.

We consider two simple pressure marks to be connected if a user draws them within a certain time window. Whenever a user draws a mark, if it is not a leaf of the menu tree, we display a "ripple" originating at the mark's end. This animated ripple lasts for as long as the connection time window and aims to make users aware of the opportunity to concatenate marks. If the user starts drawing a mark while the ripple is visible, the new mark connects to the previous one. We provide feedback for this concatenation by displaying the sequence of icons/labels (i.e. menu options) selected up to that moment. Figure 17 illustrates this.

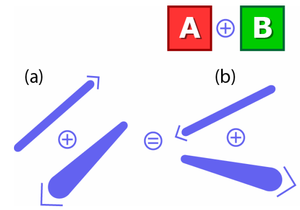

Figure 16: Different simple pressure mark flicks. (a) zig-zag pattern. (b,c) random directions. All correspond to the same command.

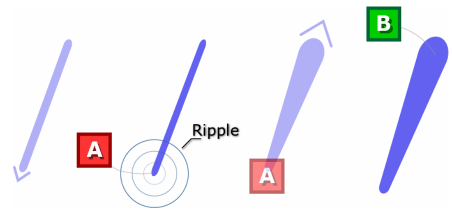

Figure 17: Ripple feedback. From left to right: a LL mark is made that triggers an expanding circular ripple. A LH mark is made while the ripple was active resulting on a LL+LH compound mark.

Whereas we envision expert users performing simple pressure marks straightforwardly, we argue that novice users can take advantage of the visualization and browsing mode discussed previously in the pressure marks section.

Pressure Tails

Pressure marks can also leverage pigtail delimiters to produce a technique called *pressure tails*, which allows for a fluid selection-action-manipulation phrase where the pressure signature comes into play only when the pigtail gesture is performed. The advantage of integrating pigtails and pressure marks is that users do not need to be concerned as to the direction they are pigtailing, since the action is specified by the pressure signature. Instead, a pigtail delimiter marks the beginning of a parameter manipulation, such as the position of an object or its scaling factor. The self-crossing gesture defines a crossing interaction, wherein the pressure before and after the cross can be used to produce simple heuristics for the mark's parsing. Figure 18 illustrates an example of pressure tails.

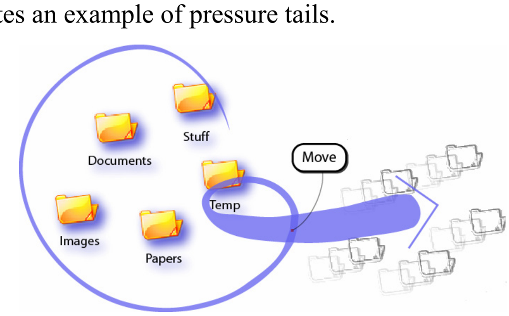

Figure 18: Pressure tails example. A LH crossing signature lets users move a group of objects.

Pressure Fanning

There are situations when users need to inspect information inside a container such as a folder. Similarly, information or GUI elements can be structured in piles. Agarawala and Balakrishnan [1] explore fanning as an interaction technique for revealing the content of piles in the interface. Pressure marks offer the means to provide additional semantics to such a fanning gesture – e.g., depending on a stroke's pressure signature, one can fan out the contents of a pile sorted in ascending or descending order, unsorted or as a means to break the pile (Figure 19). We have implemented this technique and initial feedback is encouraging.

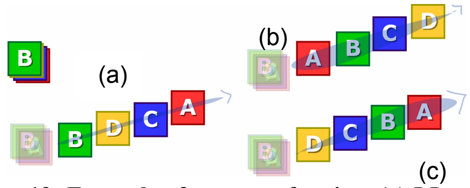

Figure 19: Example of pressure fanning. (a) LL mark fans the elements of the pile in their normal order. (b) HL mark fans the elements in ascending order. (c) LH mark fans the pile's elements in descending order.

CONCLUSION and FUTURE RESEARCH

Pressure marks are a novel way to use the pressure variations within a pen stroke in the user interface. In contrast with most techniques used today, pressure marks can encode selection-action patterns in a concurrent, parallel interaction. The results we present not only show that pressure marks are a viable interaction technique, but also reveal that their use can result in a significant reduction in the time it takes to perform selection-action patterns.

In addition to these positive results, pressure marks have potential as orientation-independent marks, thus enhancing existing marking techniques. We present several designs that explore this possibility. Although novel, these designs remain to be evaluated with a future user study. There are several paths of future work ahead of us. We plan to investigate extensions to the signature set to include compound marks – i.e., L-H-L, or H-L-H. Similarly, we can see if a stroke is drawn clockwise or counterclockwise to expand a given signature set by a factor of two. Finally, there is potential to use the ideas and techniques we presented with input devices that do not support pressure – for example, we could use Zeleznik et al.'s *Pop-Through mouse* [23], count the number of fingers touching an interactive surface, or measure the area of a fingertip in contact with an interactive surface.

REFERENCES

1. Agarawala, A. and Balakrishnan, R. (2006). Keepin' It Real: Pushing the Desktop Metaphor with Physics, Piles and the Pen. *CHI*, 1283-1292.
2. Baudisch, P., Xie, X., Wang, C. and Ma, W.-Y. (2004). Collapse-to-zoom: viewing web pages on small screen devices by interactively removing irrelevant content. *UIST*, 91-94.
3. Boothe, R.G. (2001) Perception of the Visual Environment. Springer
4. Buxton, W. and Myers, B. (1986). A study in two-handed input. *CHI*, 321-326.
5. Elliott, D., Helsen, W.F. and Chua, R. (2001). A century later: Woodworth's, 1899. Two-component model of goal-directed aiming. *Psych. Bulletin* (127). 342-357.
6. Fitzmaurice, G., Balakrishnan, R. and Kurtenbach, G. (1999). An exploration into supporting artwork orientation in the user interface. *CHI*, 167-174.
7. Guimbretière, F., Martin, A. and Winograd, T. (2005). Benefits of merging command selection and direct manipulation. *TOCHI*, *12* (3). 460-476.
8. Hinckley, K., Baudisch, P., Ramos, G. and Guimbretiere, F. (2005). Design and analysis of delimiters for selection-action pen gesture phrases in scriboli. *CHI*, 451-460.
9. Hinckley, K., Guimbretiere, F., Agrawala, M., Apitz, G. and Chen, N. (2006). Phrasing techniques for multi-stroke selection gestures. *GI*, 147-154.
10. Jacob, R., Sibert, L., McFarlane, D. and Mullen, M. (1994). Integrality and separability of input devices. *TOCHI*, *1* (1). 3-26.
11. Kara, L.B. and Stahovich, T.F. (2004). Hierarchical parsing and recognition of hand-sketched diagrams. *UIST*, 13-22.
12. Kurtenbach, G. and Buxton, W. (1993). The limits of expert performance using hierarchical marking menus. *CHI*, 35-42.
13. Latulipe, C., Kaplan, C.S. and Clarke, C.L.A. (2005). Bimanual and unimanual image alignment: an evaluation of mouse-based techniques. *UIST*, 123-131.
14. Li, Y., Hinckley, K., Guan, Z. and Landay, J.A. (2005). Experimental analysis of mode switching techniques in pen-based user interfaces. *CHI*, 461-470.
15. Mizobuchi, S. and Yasumura, M. (2004). Tapping vs. circling selections on pen-based devices: evidence for different performance-shaping factors. *CHI*, 607-614.
16. Ramos, G. and Balakrishnan, R. (2005). Zliding: fluid zooming and sliding for high precision parameter manipulation. *UIST*, 143-152.
17. Ramos, G., Boulos, M. and Balakrishnan, R. (2004). Pressure Widgets. *CHI*, 487-494.
18. Rubine, D. (1991). Specifying gestures by example. *SIGGRAPH*, 329-337.
19. Saund, E. and Lank, E. (2003). Stylus input and editing without prior selection of mode. *UIST*, 213-216.
20. Shen, C., Vernier, F.D., Forlines, C. and Ringel, M. (2004). DiamondSpin: an extensible toolkit for around-the-table interaction. *CHI*, 167-174.
21. Srinivasan, M. and Chen, J. (1993). Human performance in controlling normal forces of contact with rigid objects. *ASME*, *49*. 119-125.
22. Wu, M. and Balakrishnan, R. (2003). Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *UIST*, 193-202.
23. Zeleznik, R., Miller, T. and Forsberg, A. (2001). Pop through mouse button interactions. *UIST*, 195-196.
24. Zhao, S. and Balakrishnan, R. (2004). Simple vs. compound mark hierarchical marking menus. *UIST*, 33-42.