Guidelines for the Design of Haptic Widgets

Ian Oakley, Alison Adams, Stephen Brewster and Philip Gray

Glasgow Interactive Systems Group, Dept of Computing Science
University of Glasgow, Glasgow, G12 8QQ, UK
+44 (0) 141 330 4966
io, adamsa, stephen, pdg@dcs.gla.ac.uk
http://www.dcs.gla.ac.uk/~stephen

Haptic feedback has been shown to improve user performance in Graphical User Interface (GUI) targeting tasks in a number of studies. These studies have typically focused on interactions with individual targets, and it is unclear whether the performance increases reported will generalise to the more realistic situation where multiple targets are presented simultaneously. This paper addresses this issue in two ways. Firstly two empirical studies dealing with groups of haptically augmented widgets are presented. These reveal that haptic augmentations of complex widgets can reduce performance, although carefully designed feedback can result in performance improvements. The results of these studies are then used in conjunction with the previous literature to generate general design guidelines for the creation of haptic widgets.

Keywords: Haptic, Desktop, GUI, Multi-target, Design guidelines

1 Introduction

There is a growing body of literature indicating that haptic feedback, or feedback that allows a user to *feel* an interface, can yield performance improvements in target acquisition tasks in GUIs. As early as 1994 Akamatsu & Sate (1994) demonstrated lower task completion times in a target acquisition task using a simple haptic mouse with adjustable friction and vibro-tactile display. In the same year, Engel *et al.* (1994) showed reduced error rates and task completion times using a haptic trackball with 2 degrees of freedom force feedback. Research on this topic has continued and more recently a number of researchers (e.g. Dennerlein *et al.*, 2000, Miller & Zelenik, 1998, Oakley *et al.*, 2000) have reported performance improvements attributable to haptic feedback as presented on a number of devices. In this more recent research haptic targets are typically presented as walled areas, or as wells of attractive force that actively draw the cursor towards their centre. Targets augmented in either of these ways exhibit a "snap-to" behaviour, actively capturing the cursor as it strays over them, and requiring a user to exert effort in order to move off the target.

This research provides compelling evidence for incorporating haptic feedback into GUIs, but in fact, tells only half of the story. While there has been extensive research on the presentation of single haptic targets, there have been few investigations of more realistic scenarios incorporating multiple targets. In such situations, the influence exerted by haptic targets incidentally traversed by users as they move towards their desired destinations must be considered. The extraneous forces these widgets apply have the potential to alter the paths users wish to take, and consequently may reduce their performance and satisfaction. Indeed, this assertion is upheld in a study conducted by Oakley *et al.* (2001) incorporating a condition investigating a standard haptically augmented menu system.

One possible solution to this problem is to try to remove the unwanted haptic feedback by attempting to predict a user's desired destination, and applying the feedback only on this target. Such a manipulation, if successful, would serve to reduce the complexity of the multi-target case to the simplicity of the single target case, and transfer the performance benefits gained there. However, as Dennerlein & Yang. (2001) point out:

"...only enabling one force field is an unrealistic simulation for the implementation of force-feedback algorithms. If one confidently knew the desired target, why not then select that target automatically without using the mouse?"

According to this rationale, Dennerlein & Yang are considering the implications of partially successfully target prediction systems. They describe a study presenting multiple targets to users, and manually control the number of haptic distracter targets between a user and the destination target. They reason that adjusting the number of distracter targets simulates different accuracies of target prediction. Their conclusions are mixed. They suggest that while objective measures of performance may be maintained by using partially successful target prediction algorithms, a user's subjective experience can be negatively affected.

The practicalities underpinning target prediction, however, seem more in doubt than the validity of the idea. Oirschot & Houtsma (2000, 2001) describe several studies investigating the accuracy of prediction of the final destination of a movement given its initial trajectory. They conclude that although the creation of an algorithm to perform such a task may be possible, the parameters that control it would vary substantially from device to device and from user to user. Munch & Dillmann (1997) describe a complete system that provides not only haptic feedback in a GUI, but also a target prediction system that attempts to mediate the application of this feedback. Their target prediction system relies on both trajectory analysis and a model of application behaviour to determine user destination. They suggest that it would only be successful after a learning period for each combination of user and application.

To summarise, the literature relating to target prediction suggests that although it may be an objectively effective solution to the problems of multi-target haptic interaction, it is also a costly and underdeveloped one. Differences between individuals, devices, and even applications may be enough to render such systems useless without substantial training times. More worryingly, the evidence suggests that partially successful systems may exert a damaging influence on subjective satisfaction.

Oakley *et al.* (2001) suggest an alternative solution to the multi-target problem. They describe a study investigating a haptically augmented hierarchical menu system that led to performance improvements similar to those observed in haptically augmented single target interactions. They achieved this by dynamically tailoring the forces in the menu to support, and not obstruct, the motions undertaken by users. This was done through the modification of the magnitude of the force applied (which was in the form of a walled area) through the application of two simple rules. Firstly, when a user was moving slowly the maximum force exerted by a target was reduced. This enabled users to move from one menu item to an immediately adjacent one without being hindered by strong forces. Secondly, when a user was moving rapidly, which in a menu system tends to occur either horizontally or vertically, the forces that opposed that motion were reduced, while the ones that supported it were maintained. Effectively, as a user moved horizontally in the menu, only the vertical forces that aided that motion were presented, and *vice versa*. This allowed users to move across, or along, menu items at speed, gaining the benefits of forces supporting, without the cost of those obstructing, these actions. The results of the study showed that the condition incorporating these dynamic forces bettered a purely visual condition with no haptic feedback by reducing errors, and a condition incorporating the same haptic feedback, without these adjustments, by lowering task completion times, and reducing subjective workload.

Here we extend these ideas by describing two studies applying this kind of dynamic haptic feedback to different multi-target situations. We then build on these results, in conjunction with the previous literature, to create preliminary design guidelines for the creation of complex haptic widgets.

2   Experimental Overview

2.1 Equipment

Both experiments were conducted under Windows NT and force feedback was provided by a PHANToM (from SensAble Technologies) equipped with a pen stylus featuring a button. The PHANToM (pictured in Figure 1) is a sophisticated 3 Degree of Freedom (DOF) output and 6 DOF input force feedback device. The workspace

Figure 1. The PHANToM from SensAble Technologies.

available to participants was restricted to a narrow vertical plane, 110 mm wide by 82.5 high mm (matching the available graphical range of 800 by 600 pixels) and 20 mm deep. Motion along the x and y axes controlled cursor position. No action was mapped to motion on the z-axis.

2.2 Haptic Feedback

There were two types of haptic feedback used in these studies: Standard Haptic and Adjusted Haptic. The first of these comes in the form of simple two-dimensional haptic barriers. To enable these barriers to reside next to one another they had a simple force profile ensuring that either side of the barrier returned to zero force. Four of these barriers arranged to enclose a rectangular area served to produce a haptified target. The force profile of a single barrier is shown in Figure 2. One consequence of this implementation is that the corners of targets are subject to more substantial forces, as barriers in both x and y dimensions independently contribute force. This problem was partially resolved by capping the maximum exerted force to the maximum for a single barrier, but the corners of a target still consisted of larger areas of the maximum force. This makes diagonal motion more difficult than either horizontal or vertical motion. The maximum magnitude of the haptic barriers differed between the two studies: in the first it was 0.25 Newtons (N), in the second 0.65 N. This algorithm has similar properties to those used in other studies of haptified targets; moving over a target causes a user to be pulled into its centre, and leaving a target requires overcoming the barrier forces surrounding it.

The Adjusted Haptic feedback was based on the Standard Haptic feedback, possessing the same basic structure and magnitude (0.25 in the first study, 0.65 in the second). It was created by dynamically applying the following three rules to modify the maximum applied strength of the haptic barriers. One: reduce the maximum force applied if a user is moving slowly (beneath 2 cm per second) to a minimum of one third of its normal value. Two: if a user is moving rapidly (above 2 cm per second) and has only been on a target for a short time (less than 100 ms) reduce the maximum
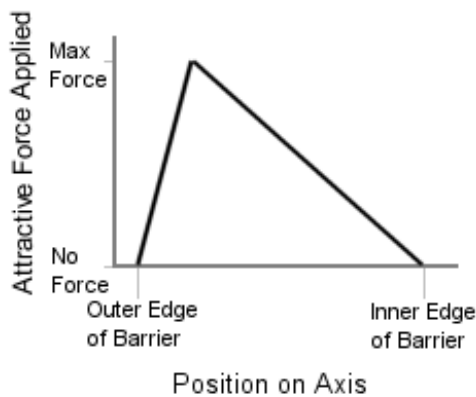

Figure 2. Force profile for a single haptic barrier.

applied force by a factor of two. Three: increase the maximum force applied to three times its original amount if a user has begun to perform a click (by depressing the PHANToM's button) and reduce the force back to normal levels when the click is completed (by releasing the button). The rationale for these choices is that the first will enable users to easily move to adjacent items, the second will facilitate rapid unobstructed movements and the third will increase the likelihood that a clicking action, once begun, will be successfully completed. All transitions between different force magnitudes were gradual, so as not to disrupt users, but took place extremely rapidly. This was made possible due to the PHANToM's native 1000 Hz update rate.

2.3 Experimental Design and Participants

Both studies had the same basic design. Both incorporated three conditions: Visual, Haptic and Adjusted. No haptic feedback was present in the Visual condition. The Haptic condition included the Standard Haptic feedback on each target, while the Adjusted condition featured the Adjusted Haptic feedback.

The first study involved eighteen participants, the second twelve. No participants performed in both studies. The majority of the participants were computing science students, the rest were experienced computer users. No participant had more than trivial previous experience with haptic interfaces. Both studies featured fully balanced repeated measures experimental designs; each included six order conditions. Each order condition was completed by three participants in the first study, and two participants in the second study. Practice in all three conditions in both studies took place immediately before the experiment began and always occurred in the same order as the presentation of conditions in the experimental session.

2.4 Measures

Both studies were subject to the same basic range of subjective and objective measures. Subjective assessment was achieved through the application, after the completion of each condition, of a modified version of NASA TLX (Hart & Staveland, 1988), an established measure of workload. Standard TLX questionnaires consist of the following six scales: Mental Demand, Physical Demand, Time Pressure, Effort Expended, Frustration Experienced and Performance Level Achieved. We included one extra factor: Fatigue Experienced. We feel that this is an important additional factor to consider with regard to haptic interfaces.

Objective measures in both studies included task completion time and a detailed taxonomy of errors. Both studies were essentially target acquisition and selection tasks, and in this situation we consider an error to have occurred when a user moves over the desired target and then off it again without completing the selection process. These errors fall into two categories. Firstly a *slide over*, which occurs when the user simply moves over the correct button, and then off it, without making any attempt to select it. This is arguably part of the normal targeting process. The second and more serious error is a s*lip off* (Brewster, 1998), which occurs when a user initiates the selection process by (at least) depressing the PHANToM's button, but then moves off the target before releasing the PHANToM's button. In this case the target is not activated, although the visual feedback received is typically the same as for a

successful operation. The final classification of errors is simply a *miss* – a selection of the wrong target. In this category we include failed attempts to select the wrong target; situations in which a user begins to select an inappropriate button, but then fails to complete this action by performing a slip off.

2.5 *Hypotheses*

Both studies shared similar hypotheses. We predicted that the Haptic condition would show fewer errors than the Visual condition (as the haptic walls make staying on a target easier), at the cost of an increase in time, workload and possibly slide over errors (as the walls make movement more difficult). We hypothesised that the Adjusted condition would combine the positive aspects of both the other conditions, yielding low task completion times, low workload and a low occurrence of errors.

3   Experiment 1

3.1 *Task*

This study involved the evaluation of a haptically augmented toolbar. Each button in the toolbar was 22 pixels square visually and 3.025 mm square haptically. The toolbar consisted of twenty-five buttons arranged in a square configuration, and is pictured in the centre of Figure 3. The visual representation and behaviour of the toolbar was based on the toolbars that appear in existing GUIs. Moving over a button led to it visually highlighting; depressing the PHANToM's button led to the display of a different highlighted state; releasing the PHANToM's button completed the interaction. To ensure a wide variety of approach angles to this toolbar, it was placed in the centre of the eight large start points, each identified by a number and positioned at 45-degree intervals around it. Each trial involved moving over a specific start point identified in an instruction panel on the far right of the screen. When this occurred, a picture of the target button was displayed in the instruction panel. Selecting this button in the toolbar completed the trial, and caused a new start point to be displayed.
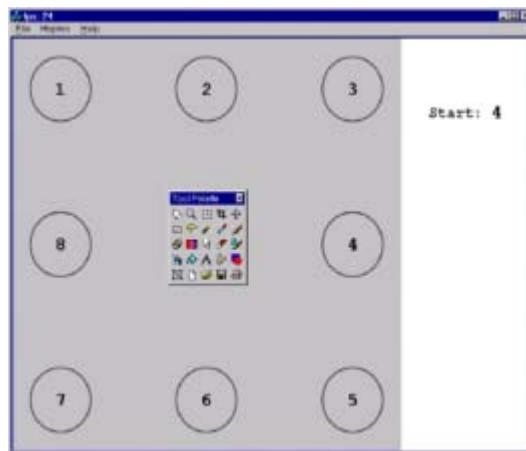

Figure 3. Screenshot of haptic toolbar study.

Given eight start points and twenty-five targets, there were two hundred trials in the experiment: one instance of every possible combination. These were displayed in a random order. Task completion time was measured from the moment a participant left the start point until the successful selection of the appropriate button.

3.2 Results

All analyses of subjective measures, time and errors were conducted using repeated measures single factor ANOVA and *post-hoc* t-tests, using Bonferroni confidence interval adjustments. Results from the TLX questionnaire are presented in Figure 4, adjusted so that higher ratings consistently indicate higher workload. Overall workload was significantly higher in the Haptic condition than in the Visual and Adjusted conditions (both $p<0.001$). The Haptic condition was rated significantly more taxing than the Adjusted condition in all individual scales (all at $p<0.01$) except Time Pressure, and more demanding than the Visual condition in all factors (all at $p<0.01$) bar Time Pressure, Performance Achieved and Mental Demand. There were no significant differences between the Visual and Adjusted conditions in any aspect of the subjective measures.

The timing data are presented in Figure 5. No significant difference in task completion time was found between the Visual and Adjusted conditions, while both yielded significantly faster times than the Haptic condition (both $p<0.01$). Error data are presented in Figure 6. The Adjusted and Haptic conditions produced significantly fewer slip offs than the Visual condition (both $p<0.01$). The Adjusted condition also
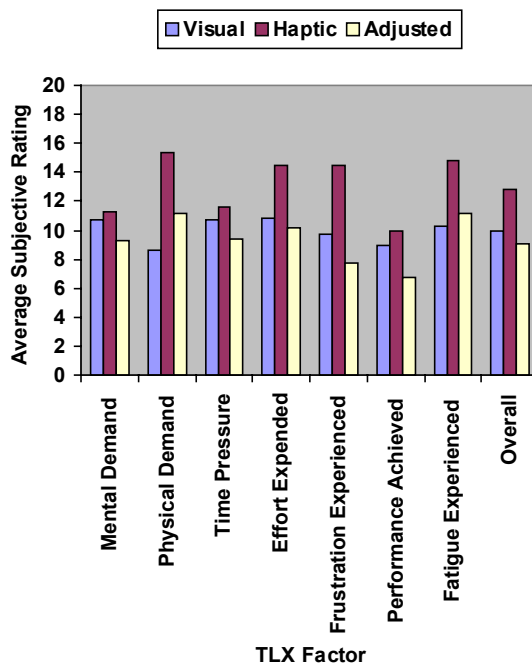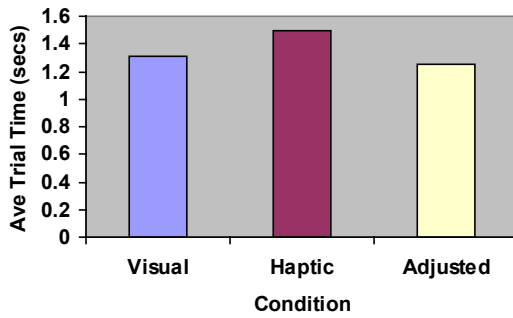

Figure 4. TLX ratings for study 1.

Figure 5. Average task completion times in study 1.

yielded significantly fewer slide overs than the Visual and Haptic conditions (both p<0.01), and the Visual condition fewer than the Haptic (p<0.05). Finally, there were no significant differences in the number of misses.

3.3 Discussion

In this study, the experimental hypotheses were upheld: the Adjusted condition combined the favourable aspects of both other conditions. It attained the low error count apparent in the Haptic condition and the rapid task completion time present in the Visual condition. The Haptic condition was more subjectively taxing than the other two conditions. These results support those reported in Oakley *et al.*'s (2001) study of a haptically enhanced menu. The standard haptic feedback that is effective in a single target situation results in a performance hit when applied in a situation incorporating multiple targets. Appropriately adjusted haptic feedback, however, can lead to performance benefits in these complex environments.
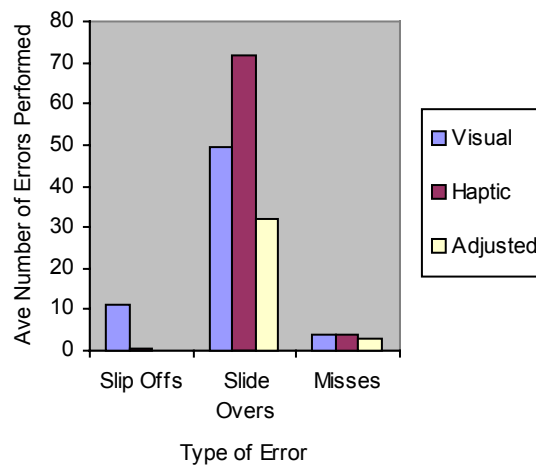

Figure 6. Errors recorded in study 1.

4. Experiment 2

4.1 Introduction

This second study was related to the first, and involved the haptic augmentation of icons spread across a canvas; a typical cluttered desktop. From the perspective of creating haptic augmentations there are several key differences between a group of icons and the buttons on a toolbar. Most importantly, groups of icons, unlike toolbars, do not possess a highly structured and rigid spatial arrangement. There is no guarantee that a target will be adjacent to another, and even adjacent targets tend to be separated by bands of empty space. A second difference is simply one of size – icons are much larger that buttons on a toolbar, and are spread over a greater area. Consequently, it may be reasonable to expect that the speed at which users move may also alter from that used when interacting with a toolbar. These factors seem likely to exert some influence on the effectiveness of haptic augmentations, and we sought clarification as to what this might be.

4.2 Task

The graphical representation of the icons was 32 pixels square, while the text underneath spread further than this (up to 52 pixels). The haptic representation encompassed both these areas and was slightly larger than the graphical representation at 7.7 mm, or 56 pixels, square. This discrepancy is due to the fact that the icons were sensitive to selection events occurring slightly beyond the range of the graphical representation of the widget, and the area of the haptic target was made to match this active range. This behaviour is typical of icons in windowing systems. Clicking once on an icon caused it to highlight (if it was not already) and double clicking activated the icon. Both highlighting and activation were triggered in response to the depression of the controller button, rather than the release. The user was not able to move the widgets.
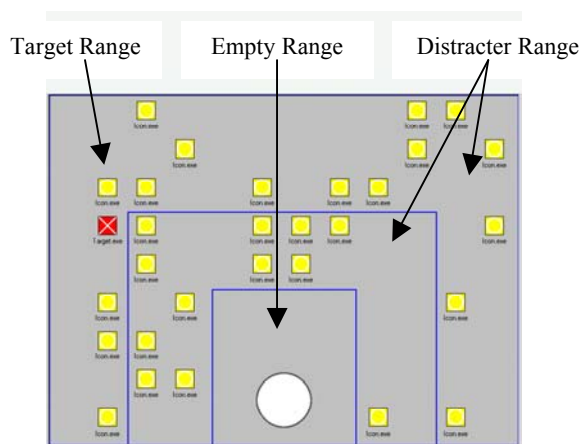

Figure 7 Screenshot of haptic icons study.

    Thirty icons (including the target) were present on the screen at all times, and each trial in the study involved moving over a single, stationary start point (positioned in the centre bottom of the screen) and then moving to and double clicking on a specific icon. The position of each icon was randomised on the completion of each trial within the bounds of the following three restrictions. Firstly, a 12 by 9 grid of valid icons positions was used to ensure that the targets appeared neatly arranged in rows and columns. This grid also ensured that there was always a 1.375 mm (or 10 pixel) space between adjacent targets. Secondly, no icons were ever positioned near the start point (it resided at the centre bottom of a four icon by four icon gap in the grid). Finally, the active target, the one that participants had to select, was constrained to appear outside of a larger 8 by 6 gap around the start point. These manipulations ensured that targets could never intrude on the start point – users always began a trial in a haptically empty space – and increased the likelihood that users would have to traverse distracter targets in order to reach their desired destination. Figure 7 is a screenshot of the experiment, labelled to indicate these positioning constraints. As we were not interested in the cognitive search time involved in locating the target icon, we used only two, very distinct, graphical representations for the icons; the target was a red cross, the distracters were yellow circles. Each condition in the study involved two hundred trials. Task completion time was measured from the moment a participant left the start point until the successful selection (double click) of the target icon.
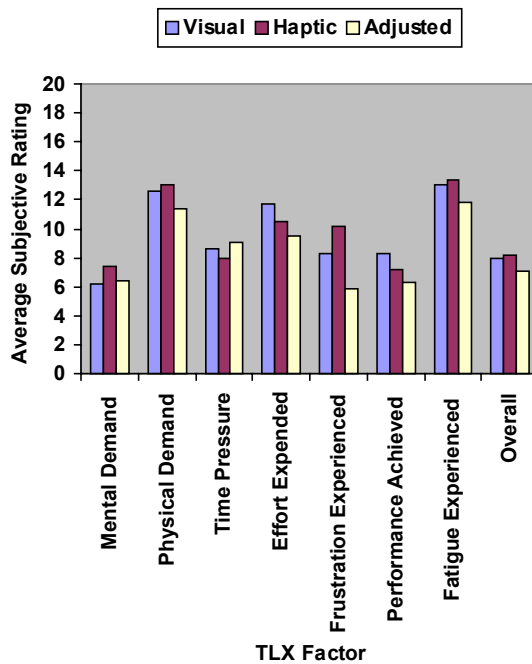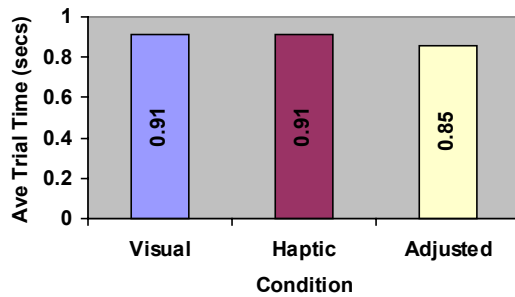

Figure 8 TLX ratings for study 2.

Figure 9. Task completion times in study 2.

4.3 Results

All analyses of subjective measures, time and errors were conducted using repeated measures single factor ANOVA and *post-hoc* t-tests, using Bonferroni confidence interval adjustments. The TLX results are pictured in Figure 8, adjusted so that higher ratings always indicate a higher workload. Few significant differences were revealed between the three conditions. The Adjusted condition yielded significantly improved ratings of Effort Expended and Performance Level Achieved when compared to the Visual condition (respectively $p < 0.05$ and $p < 0.01$), and significantly reduced Frustration Experienced when compared to the Haptic condition ($p < 0.005$).

The timing data are presented in Figure 9. The adjusted condition resulted in significantly faster times than either the Visual or Haptic conditions (both $p < 0.05$). Error data are presented in Figure 10. Fewer slip off and slide over errors are present in the Adjusted and Haptic conditions when compared to the Visual condition (all $p < 0.005$). The Visual condition, however, resulted in fewer misses than the Haptic condition ($p < 0.01$).
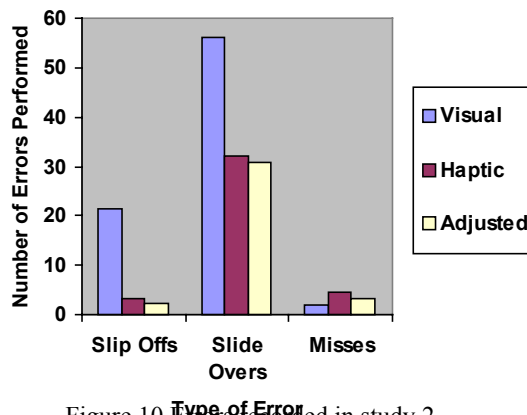

Figure 10 Errors recorded in study 2.

4.4 Discussion

The experimental hypotheses concerning the Adjusted condition were upheld. It combined the fastest task completion times with the lowest error rates, and showed a modest gain over the Visual and Haptic conditions in subjective measures. However, the Haptic condition, while exhibiting the predicted decrease in error rate compared to the Visual condition, did not produce the expected performance hit in terms of task completion time and subjective measures. As the same feedback was responsible for this performance hit in the toolbar study, this may indicate that target acquisition tasks relying on large, spatially separated targets are not sensitive indicators of the effectiveness of a haptic augmentation. Techniques that improve performance in these situations may fail in more challenging scenarios and the ability to generalise from them should be questioned. Interestingly, the Visual condition resulted in fewer miss errors than the Haptic condition. We suggest that this may be a consequence of participants getting "snagged" on nearby distracter targets, when, without the haptic feedback, their velocity would have normally carried them over their desired target.

5 Guidelines for the design of haptic widgets

The two studies described here, in conjunction with some of the previous literature (Oakley *et al.*, 2001), indicate that the haptic augmentation of multiple targets with unadjusted attractive forces is at best not optimal, and at worst can reduce performance and subjective satisfaction. However, they also suggest that appropriate haptic feedback – haptic feedback that provides performance improvements at no cost - can be created through various manipulations of these haptic augmentations. Here we attempt to make explicit, to explain, the process by which these manipulations were designed. We present an initial set of guidelines for the design of effective haptic widgets to function in both single and multi target situations.

Miller & Zeleznik (1998) have previously presented three guiding principles to aid the creation of haptic widgets. Firstly, they suggest that haptic feedback should be used to reduce errors through guidance; to provide forces to support the motions that a user is undertaking. Secondly, they indicate that the forces applied should function as feedback; they should be based upon, but never control, a user's input. Finally, they state that any force feedback applied to a user should be overridable; a user should be able to pop through, or escape, from any haptically augmented area. The guidelines we present here share similar tenets to these principals. However, we try to go further, to more precisely define the problems and solutions involved in adding haptic feedback to desktop widgets.

These guidelines are based around the idea that the force presented should support, and not oppose, a user's intent. This entails drawing a balance between allowing users to move where they want as freely as possibly, and providing forces to improve targeting and reduce errors. An advantage of these guidelines is that they do not require target prediction, a currently immature technology. A disadvantage is that they do assume that an extremely dynamic simulation controls the haptic feedback. The guidelines rely on the rapid, smooth adjustment of force magnitude according to the current state of interaction, and this flexibility may be challenging to implement

on current consumer devices (such as the haptic mouse used by Dennerlein *et al.* (2000)).

5.1 Guiding Strategy

Haptic feedback has the potential to improve objective user performance in two ways: reducing the number of errors made, or decreasing task completion times. We suggest that it is more profitable to design haptic augmentations to achieve this first aim, to reduce errors. There are several reasons for this. Firstly, a reduction in errors can be linked to improvements in other metrics: it has been associated with decreases in task completion time (Dennerlein & Yang, 2001), and some studies have linked it to subjective satisfaction (Oakley *et al.*, 2000). Secondly, to gain an increase in task completion time, users must adopt a movement strategy supported by the haptic augmentation, and there is no guarantee that this will occur. There are more assurances that forces to prevent errors will be successful. For instance, an attractive basin supports faster movement times, because targeting is simpler, and a reduction in errors, because a conscious effort is required to leave the target. However, although users may move towards the target more rapidly, this is a choice they make. The decrease in errors, on the other hand, is simply a property of the attractive forces applied over the target.

5.2 Choice of haptic augmentation

Widgets augmented with attractive basins or haptically walled areas have typically provided the best performance improvements (Oakley *et al.*, 2000). However when designing the haptic feedback for a widget, it is also important to consider its shape, and the likely path a user will take over it. For instance, when using these standard targeting augmentations in conjunction with square or rectangular widgets, diagonal motion is more difficult than horizontal or vertical motion. This may have an impact on performance and subjective satisfaction.

5.3 Interaction between force strength and widget size

The maximum strength used for any widget, or set of widgets, should be dependent on both the size of the widgets and density of the arrangement that they are presented in. As the toolbar study described in this paper indicates, a dense arrangement of small widgets requires small forces, as large forces will severely hamper motion from one widget to an adjacent one. Also motions over small, well packed widgets are likely to be slower, as only a short distance must be traveled. Consequently small forces are sufficient to aid targeting. Correspondingly, large, spatially separated widgets suit much stronger forces, as illustrated in the second study presented here. With the absence of nearby widgets, the presence of these stronger forces is less likely to be disruptive. Also users often approach large spatially distributed widgets at considerable speed. Thus, targeting benefits are likely to be maximised by increasing the strength of the forces applied, to match the increase in approach speeds.

5.4   Range of useful force magnitudes

The literature suggests that the maximum strengths of haptic targets should be in the range of 0.25 N (used in the toolbar study here) to 0.8 N (Dennerlein *et al.*, 2000). There is little research indicating whether these figures would be device dependant, and it is worth noting that for use in multi-target situations, feedback of these magnitudes will typically need to be adjusted as described later in these guidelines. Maximum applied strength is also likely to be highly dependent on individual differences. In any real system, it would be essential that maximum strength be user configurable. We suggest however, that the general strength ratios between different sizes and densities of widgets would stay more or less the same across users; irrespective of the maximum strength a user chooses, the proportions between the magnitude of the forces applied over a large target, and of that applied over a small target seem likely to remain the same.

5.5   Exploit patterns of user behaviour

The haptic feedback present on a widget should capitalise on patterns of motion afforded by that widget. This is often related to the shape of the widget. In Oakley *et al.*'s (2001) study of haptically augmented menus, they exploited the fact that motion in a menu typically occurs either vertically or horizontally to provide only supportive forces. Given the similar shape of the widgets, this same manipulation has the potential to apply to scrollbars. The scrollbar could exert strong targeting forces as a user moves along its length, but fade out these forces when a user attempts to move off it, in a direction perpendicular to its length. In a scrollbar it may also be appropriate to increase the strength of the targeting forces with increased speed along the scrollbars length.

5.6   Exploit widget behaviour

Widget behaviour can also often be exploited to increase the effectiveness of a haptic augmentation at no cost. In both studies reported in this paper the Adjusted conditions incorporate haptic feedback designed to aid the completion of an action that has been begun. The strength of the haptic walls increases when a user begins to select a target, and reduces to normal levels on the completion of that selection. A similar strategy could be applied to any widget interaction that incorporates more than a single explicit stage. Beginning the interaction could trigger a change in the haptics, such as an increase in magnitude, designed to support the interactions successful completion.

5.7 Dynamic response to slow movement

Force should vary according to speed: slow motions require low forces. In situations incorporating densely packed widgets this is especially important. It is hard to traverse from one widget to an adjacent widget when opposed by even a relatively low force. Users often end up moving further than they intended, "popping through" the target widget onto one beyond it. They are then faced with the same task again – moving to an adjacent widget. This can lead to very frustrating interactions, and is arguably the biggest problem with multi-target haptic widgets. Varying the applied

force such that slower motions are opposed by lower forces can overcome this problem, allowing users to move freely, while still providing a sufficiently strong force that accidental movements off a target are prevented. The strength of force required to support targeting clearly reduces in tandem with a reduction in the speed at which a user is moving, and to produce effective multi-target haptic augmentations it is essential to capitalise on this fact.

5.8 *Dynamic response to rapid movement*

Equally, an extremely rapid motion over a target typically indicates that it is not a user's final destination, and thus requires the application of low forces. Users do not want to be impeded by widgets that are nowhere near their final destination. Again this is especially important in situations where there is a high density of widgets. In these situations it is likely that a user will traverse over numerous irrelevant widgets before reaching his or her desired target. One mechanism to achieve this is that used in the studies described here, where weaker forces are applied during the first 100 milliseconds that a user is over a widget. A disadvantage of this manipulation is that it may decrease the effectiveness of the behaviour observed by Dennerlein & Yang (2001) in which users throw themselves at speed towards a haptic target, relying on the forces it exerts to halt them. Reducing these forces for the first few moments that a user is over a target may make it less effective at capturing a rapidly moving user.

6 Conclusions

We have presented two studies investigating the use of haptic feedback to support targeting tasks. Both studies indicate that the best performance can be gained through the application of carefully designed haptic feedback, which is dynamically responsive to current states of interaction and directly measurable aspects of user behaviour (such as velocity). Following on from these studies, we have produced preliminary guidelines to make explicit the reasoning behind the design of the successful feedback, and to allow others to apply these techniques in different situations. While we focus on augmenting standard GUIs, we suggest that the findings, and guidelines, presented here may have further applicability. Many fish-tank VR systems incorporate both haptic feedback and interface widgets, and the research described here will translate easily to these systems. Other uses may include haptics systems for scientific visualisation, or for visually impaired people. In both these scenarios users are often required to explore complex arrangements of haptic targets. Applying the techniques outlined in these guidelines could make these tasks simpler and less demanding.

Acknowledgements

References

Akamatsu, M. & Sate, S. (1994), "A multi-modal mouse with tactile and force feedback", *Int. Journal of Human-Computer Studies*, 40(3), 443-453.

Brewster, S.A. (1998). "The design of sonically enhanced widgets", *Interacting with Computers*, 11(2), 211-235.

Dennerlein, J.T., Martin, D.B., & Hasser, C. (2000), "Force-feedback improves performance for steering and combined steering-targeting tasks", in *Proceedings of CHI'2000,* ACM Press, 423-429.

Dennerlein, J.T. & Yang, M.C. (2001), "Haptic force-feedback devices for the office computer: Performance and musculoskeletal loading issues", *Human Factors*, 43(2), 278-86.

Engel, F.L., Goossens, P., & Haakma, R. (1994). "Improved Efficiency through I- and E-Feedback: A Trackball with Contextual Force Feedback", *International Journal of Human-Computer Studies*, 41(6), 949–974.

Hart, S. & Staveland, L. (1988), "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research", in *Human Mental Workload*, eds. Hancock and Meshkati., Amsterdam, 139-183.

Oirschot, H.K.V., & Houtsma, A.J.M., (2000), "Cursor Trajectory Analysis", in *LNCS Haptic Human Computer Interaction*, LNCS Press 127-134.

Oirschot, H.K.V., & Houtsma, A.J.M. (2001), "Cursor Displacement and Velocity Profiles for Targets in Various Locations", in *Proceedings of Eurohaptics 2001*, Birmingham, UK, 108-112.

Miller, T. & Zeleznik, R. (1998). "An Insidious Haptic Invasion: Adding Force Feedback to the X Desktop", in *Proceedings of UIST'98*, ACM Press, 59-64.

Munch, S. & Dillmann, R. (1997). "Haptic Output in Multimodal User Interfaces", in *Proceedings of IUI'97*, ACM Press, 105-112.

Oakley, I., McGee, M.R., Brewster, S.A., & Gray, P.D. (2000), "Putting the 'feel' into 'look and feel'", in *Proceedings of CHI'2000*, ACM Press, 415-422.

Oakley, I., Brewster, S.A., & Gray, P.D. (2001), "Solving Multi-Target Haptic Problems in Menu Interaction", in *Proceedings of CHI'2001*, ACM Press, 357-358.