Two-Handed Input in a Compound Task

Paul Kabbash[1], William Buxton[1,2], and Abigail Sellen[3,4]

[1]Computer Systems Research Institute
University of Toronto
Toronto, Ont., M5S 1A4, Canada
<surname>@dgp.toronto.edu

[3]Rank Xerox EuroPARC
61 Regent St.
Cambridge, CB2 1AB, U.K.
sellen@euroPARC.xerox.com

[2]Xerox Palo Alto Research Center
3333 Coyote Hill Rd.
Palo Alto, CA. 94304, U.S.A.

[4]MRC Applied Psychology Unit
15 Chaucer Rd.
Cambridge, CB2 2EF, U.K.

ABSTRACT

Four techniques for performing a compound drawing/color selection task were studied: a unimanual technique, a bimanual technique where different hands controlled independent subtasks, and two other bimanual techniques in which the action of the right hand depended on that of the left hand. We call this latter class of two-handed technique "asymmetric dependent," and predict that because tasks of this sort most closely conform to bimanual tasks in the everyday world, they would give rise to the best performance. Results showed that one of the asymmetric bimanual techniques, called the *Toolglass* technique, did indeed give rise to the best overall performance. Reasons for the superiority of this technique are discussed in terms of their implications for design. These are contrasted with other kinds of two-handed techniques, and it is shown how, if designed inappropriately, two hands can be worse than one.

KEYWORDS: Two-handed input, GUIs, Toolglass, palette menus, compound tasks.

INTRODUCTION

Since they were introduced by the Xerox Star in 1982 [7], graphical user interfaces (GUIs) have become the most common means of interacting with computers. However, in the intervening 12 years, the basic style of interaction with GUIs, called "direct manipulation" [8], has changed relatively little.

Underlying the work described in this paper is our long-standing belief that there is significant improvement to be gained in both the *directness* and degree of *manipulation* that GUIs afford by engaging the use of both hands in interaction. An earlier study [2] showed that not only was this kind of interaction well within the bounds of novices' ability, but that it improved performance of both novices and experts in the tasks studied.

Recently, a more general two-handed direct manipulation technique called *Toolglass* [1] has emerged. A large part of our motivation for the work reported in this paper is to make an initial study of some of the underlying assumptions for Toolglass, and to compare it to other kinds of two-handed techniques.

Proponents of two-handed techniques have claimed two potential advantages. First, the division of labor across two hands means that each hand can remain almost in "home position." Consider a compound task where one subtask requires selection from a menu, and the other involves drawing in a different part of the screen. The use of two hands means that each hand can stay in the approximate vicinity of the area in which the work is carried out (e.g., [3]). Second, assigning one subtask to each hand allows for the possibility of temporal overlap in the performance of the two subtasks, reducing the time to complete the task due to actions being carried out simultaneously (e.g., [2]). Simply on the basis of time and motion, therefore, we might predict that such techniques would lead to improved performance. However, the danger is that engaging both hands might increase the cognitive load of the task in terms of carrying out cognitive processes such as monitoring, decision-making, and task coordination.

We contend that a certain class of two-handed techniques can result in significantly improved performance without imposing additional cognitive load. However, our main motivation is to better understand both the motor and cognitive load imposed by different kinds of two-handed techniques. There are few models with which to inform design in two-handed interaction. Models such as Fitts' Law serve well in predicting aspects of one-handed interaction: the time required to articulate the necessary actions in simple, serial motor tasks. Modeling two-handed interaction is less straightforward: the various ways in which two hands are integrated have implications for both sensorimotor and cognitive aspects of the task.

In this respect, work by Guiard [4] has been helpful. Guiard pointed out that a vast majority of everyday tasks may be described as "asymmetric and dependent" in that the two hands have very different roles to play which depend on each other in three characteristic ways (we assume right-handed subjects in the description that follows) :

1. The left hand sets the frame of reference for the action of the right. For example, in hammering a nail, the left hand holds the nail while the right does the hammering.

2. The sequence of motion is left then right. For example, the left hand grips the paper, then the right starts to write with the pen.

3. The granularity of action of the left hand is coarser than that of the right. For example the left hand brings the painter's palette in and out of range, while the right hand holds the brush and does the fine strokes onto the canvas.

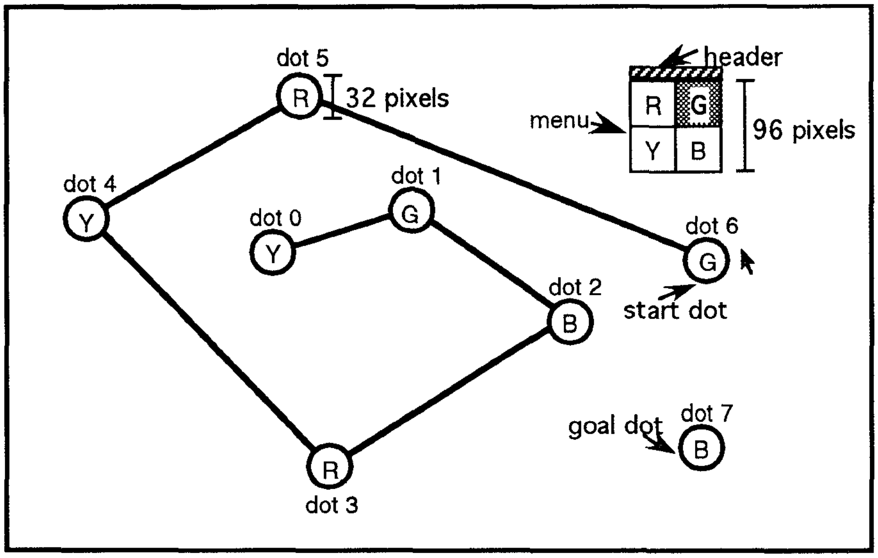

Figure 1. Experimental Task. *Having completed six dots, the subject must draw a line segment from dot 6 to dot 7, first selecting the color "blue" from the menu. The menu could be repositioned by clicking and dragging its header. (All text in the Figure is explanatory and did not appear on the screen. The dot distances and color sequence shown are arbitrary and differed for each set of dots. R = red, G = green, B = blue, Y = yellow.)*

A working hypothesis in the experiment which follows is that consistency with Guiard's characteristics is a good measure of the "naturalness" of interaction. More specifically, we predict that techniques that conform to these principles will impose a lower cognitive load on subjects leading to faster and more effortless task performance.

In our study, we tested four different techniques for performing a compound drawing/color selection task which we believed would tease out important differences in behavior across the techniques. These, hopefully, would shed some light on our understanding of the underlying processes.

METHOD

Task

Following Dillon [3], the experimental task was a colorized version of connect-the-dots in which subjects drew colored line segments between a set of twelve dots displayed on the monitor. The dots were revealed one at a time, in one of four colors, such that no two consecutive dots had the same color. To draw a connecting line, subjects needed to match their drawing color with the color of the next ("goal") dot to appear on the screen. They did this by selecting the appropriate color from a menu palette (see Figure 1)

Interaction Techniques

Subjects executed the task using four interaction techniques. The first, R-tearoff, is a one-handed technique representing the *status quo.* The other three are two-handed techniques. Each is described in detail below, assuming right-handed subjects.

Right-tearoff Menu (R-tearoff): This is modeled on tearoff menus as found in traditional GUIs. The menu is always visible, and can be repositioned by clicking over and dragging from its header. All actions are executed using a single input device (the mouse) manipulated by the right hand.

Left-tearoff Menu (L-tearoff): This is a two-hand/two-cursor version of the R-tearoff technique. Each hand controls a separate cursor on the screen. While each cursor is functionally equivalent, we have observed that in practice they tend not to be used interchangeably. The left-hand cursor is typically used exclusively for color selection, and can therefore remain floating over the menu. The right-hand cursor is usually used both for line drawing and for positioning the menu.

Except when moving the menu, therefore, each cursor/hand can always remain in "home" or "nearly home" position for its designated task. In principle, this should afford minimal movement during task performance.

Palette Menu: This is a new technique resulting from the work of Bier et al. [1] and modeled on the interaction between a painter's brush and palette. As in the previous cases, drawing is done using the right hand. However, in contrast to the L-tearoff technique, the left hand controls the *position* of the menu (or "palette"), rather than another cursor. Like the R-tearoff technique, colors are selected using the cursor controlled by the right hand. However, in theory, hand motion can be minimized by having the left hand move the menu close to the cursor when selections are to be made.

Toolglass Menu: This technique, introduced by Bier et al. [1], is much like the Palette technique in that the right hand is used for drawing and menu selection, and the left hand controls the menu position. However, there are two important distinctions. First, the menu is transparent—the user can see both the items (colors) on the menu and also what is underneath the menu. Second, color selection and the initiation of drawing are integrated into a single action. With this technique, subjects initiate each line drawing by positioning the desired color from the menu over the start dot and then clicking *through* that color. Hence, the same action that initiates drawing also performs the color selection. This integration of tasks means that fewer motor operations are required than in the Palette method, for example.

To normalize across conditions, menu size was the same in all techniques.

Technique Properties & Predictions

Table 1 summarize the properties of each technique. The first two categories (*control* and *limb-to-subtask mapping*) describe relevant physical characteristics of the techniques. The last three categories list the properties that we hypothesize contribute to both cognitive and sensorimotor load.

Control

Since the same physical device is used for each of the two subtasks (color selection and drawing) in the R-tearoff technique, these two subtasks have to be distributed across time. Hence we refer to this technique as *time-multiplexed.*

In the other three techniques, a specific device is dedicated to each subtask, resulting in a spatial distribution of the subtasks. Hence, we refer to them as having a *space-multiplexed* control structure.

Limb-to-Subtask Mapping

This categorizes whether the left (L) or right (R) hand performs the two primary tasks (drawing and color selection) and the secondary task (menu positioning).

Assemblage

In Guiard's terminology [4], the L-tearoff technique may be called an *orthogonal* assemblage because each hand executes an independent subtask (color selection in the LH and line drawing in the RH) in a separate area of the screen.

In contrast, the Toolglass and Palette techniques are *asymmetric dependent* assemblages because: the motion performed by the left hand (menu positioning) serves directly as input for right-hand action, precedes the action of the right hand, and is of coarser granularity than the right hand action. Because they mimic everyday bimanual tasks, we predicted that the Palette and Toolglass techniques would not impose any additional cognitive load over and above the R-tearoff technique. In this respect we also predicted that these two-handed techniques would be superior to the L-tearoff technique which we hypothesized would impose a considerable cognitive load on subjects.

Visual Diversion

This category attempts to quantify a cost due to the degree of "schizovisia" resulting from each technique. The assumption is that if one has to redirect one's visual attention between one part of the screen and another, the cost/load increases with the distance and complexity of the display. Cost is assumed to be due to the need to "reassimilate" the context. This is similar to the cost of assimilation discussed in Robertson et al. [6], which describes problems of assimilation when changes in data displays are made discretely, rather than by smooth animation. We assume the cost is cognitive, adding to the total cognitive load imposed by any technique.

The rankings in Table 1 were arrived at by reasoning that the effect of visual diversion worsens as the distance between the menu and the start dot gets larger. Visual diversion is assumed to be least problematic in Toolglass because color selection is performed in the same area of the screen as the start of line drawing, i.e., over the start dot. In the Palette technique, we assume that the cost of bringing the menu nearer to the work area is smaller than in L-tearoff or R-tearoff techniques. Therefore, we assume this will be done more often in the Palette condition, and thus visual diversion for this condition will be less.

Motor Operations

This category captures the number of basic motor operations required by each technique. In each, the line drawing subtask was executed with the right hand using the same two motor operations; i.e., clicking in the start dot followed by a dragging motion to the goal dot.

Each technique also allowed the menu to be repositioned on the screen. This was always necessary in the Toolglass technique. With the other three techniques, the goal dot occasionally appeared under the menu (about 3% of the time). When this occurred, subjects were forced to reposition the menu because the menu was opaque. Even when the menu was not obscuring the target dot, subjects might choose to move the menu in order to bring it closer to the start dot. This reduces the motion of the right hand when moving between the menu and the drawing task (in the R-tearoff and Palette techniques) and reduces visual diversion for the task (in the R-tearoff, L-tearoff, and Palette techniques). Since such movement was not necessarily performed on every trial, the cost for menu positioning in these three techniques is shown in parentheses in Table 1.

As can be seen from the Table, simply in terms of number of motor operations required, Toolglass has an advantage which is accentuated when menu positioning takes place in the other techniques.

Subjects

Twelve subjects participated (4 females and 8 males) with normal color vision. Subjects were strongly right handed based on the Edinburgh Handedness Inventory [5]. All had prior experience using the mouse, but widely varied in their day-to-day computer use.

	R-tearoff	L-tearoff	Palette	Toolglass
control	time-mpx	space-mpx	space-mpx	space-mpx
limb-to-subtask mapping:				
menu positioning	(R)	(R)	(L)	L
color selection	R	L	R	R
line drawing	R	R	R	R
assemblage	sequential	orthogonal	asymmetric dependent	asymmetric dependent
visual diversion (ranking)	3	3	2	1
number of motor operations	5 (+4)	4 (+4)	5 (+1)	3

Table 1. Properties of the four techniques. *Entries in parentheses indicate subtasks that were not required on every trial. R = right hand, L = left hand.*

Apparatus

Subjects performed the task on an Apple *Macintosh IIfx* with 13-inch RGB monitor, using a standard mouse in the right hand and a Kensington trackball (model *Turbo Mouse ADB Version 3.0*) in the left hand. The devices were placed on either side of a keyboard and were adjusted to control/display ratios equivalent to the second fastest setting on the *Macintosh* Control Panel for the mouse and the second slowest for the trackball. These settings were found to be optimal during pilot testing. Independent movement of the two devices was supported in software.

Design and Procedure

Subjects performed multiple trials on all four techniques in a within-subjects design. Order of techniques was counterbalanced using a Latin Square. The trials, which consisted of connecting two consecutive dots on the screen, were grouped into sets of dots, each set containing twelve dots.

For each technique, subjects were trained by verbal instruction and were given four warm-up sets of dots prior to data collection. They then connected twenty-four sets of dots on each technique. The same twenty-four sets were repeated in random order for each of the four techniques. Since data were not collected for the first trial of each set, each subject completed a total of 240 recorded trials on each technique.

At the end of each set of dots, subjects were given their total execution time over that set as well as their own best time and were instructed to try to beat their best time. The total time taken by subjects to complete the experiment was approximately one hour.

After completing the four tasks, subjects were asked two questions to elicit qualitative ranking. In order, these were "Which of the four techniques did you dislike the most?" and "Which technique did you prefer?".

RESULTS & DISCUSSION

Trial Completion Time

Total trial completion time (TCT) differed significantly amongst techniques ($F_{3,30} = 16.33$, $p < .001$) as shown in Table 2. Peritz post hoc tests for the differences between means [9] confirmed that Toolglass was faster than all other techniques with no other differences between pairs.

TCT got shorter with practice, regardless of the technique used, as shown by a main effect of position on TCT ($F_{3,30} = 3.37$, $p < .05$). Peritz tests revealed significant differences between techniques performed first as opposed to last, but no other differences amongst position. (This was the only dependent variable to show an effect of position.)

Since, as summarized in Table 1, Toolglass required the fewest motor operations (only 3 per trial), it is tempting to point to this as the reason for this technique being the fastest to perform. However, this explanation is far too simplistic. First, the mere number of operations does not take into account differences between the techniques in the level of difficulty or time to complete individual motor operations. For example, in the L-tearoff technique the menu was repositioned on only a very small proportion of the trials (2.9%). When it did not require menu positioning, this condition had only 4 motor operations, including a button click and an easy pointing movement between menu color items (Fitts' index of difficulty = 1). Thus, Toolglass and L-tearoff may have been roughly equivalent in terms of motor costs alone, not accounting for why Toolglass was faster.

Second, it is important to remember that TCT represents the aggregate of the time to complete the elemental actions necessary to carry out the task, as well as the effects of cognitive load imposed by the task demands. One aspect of this which is of particular interest is the extent to which the use of the left hand in parallel with the right hand might impose additional cognitive load.

	total (sec)	LH use	draw (sec)	LH during draw
R-tearoff	2.89	--	.818	--
L-tearoff	2.96	49%	1.01	21%
Palette	2.90	47%	.878	29%
Toolglass	2.43	83%	.991	63%

Table 2. Mean trial completion times. *The second and fourth columns give the sum of the time periods during which the left hand was in motion, as a proportion of total trial time and drawing time, respectively.*

Left-hand Use

Table 2 shows the proportion of time the left hand (LH) was in use as a function of total trial duration for the three techniques involving the LH. The main effect of technique on LH use was significant ($F_{2,16} = 19.56$, $p < .001$) with Toolglass exhibiting much greater use (83% in Toolglass vs. 49% and 47% in the L-tearoff and Palette techniques, respectively). There was no difference between the L-tearoff and Palette conditions. Since the Toolglass resulted in the highest use of the LH, but lowest overall TCT, this suggests that extensive use of the LH is not a major hindrance to overall performance speed.

A more systematic way of examining the effect of LH use on right hand (RH) performance is to look at the time taken to execute the drawing task, as this was, from a purely mechanistic view, the same across techniques. We might predict that an increase in LH use during the drawing task would correspond to an increase in drawing time for the RH.

Draw movement time (draw MT) was measured from the moment the right hand clicked in the start dot until the mouse button was released at the completion of line drawing (see Table 2). There was a small but significant difference between techniques for draw MT ($F_{3,30} = 12.82$, $p < .001$). Peritz tests showed that the techniques fell into two groups; the means for draw MT were greater in L-tearoff and Toolglass than in R-tearoff and Palette, and there were no differences within groups. LH use during the drawing task was found to be higher in the Toolglass condition than both L-tearoff and Palette ($F_{2,20} = 22.46$, $p < .001$) with no difference between the L-tearoff and Palette conditions (Table 2).

Therefore, there is no straightforward relationship between LH use and the time to complete the drawing task. Although the longer time for Toolglass could be interpreted as due to interference from the LH, this explanation does not hold for the longer draw MT for the L-tearoff condition.

Both LH use analyses (in terms of total trial duration and draw MT) fail to provide evidence that an increase in LH use slows down task completion. An alternative interpretation of the data takes into account the fact that the degree of parallelism between hands is under subjects' control. Thus, subjects will only act in parallel to the extent allowed by the demands of the task. If the task is overly demanding, there will be fewer cognitive resources available to devote to LH use. The implication is that degree of parallelism can be taken to be a measure of cognitive load—the higher the degree of parallelism, the less the cognitive load imposed by the task. According to this interpretation, we would conclude that there is a lower cognitive load imposed by the Toolglass technique as compared to the other two-handed techniques. On both measures of LH use, the Toolglass gave rise to the most parallelism.

	between trials	after menu selection	total
R-tearoff	.113	.118	.231
L-tearoff	.258	.231	.489
Palette	.077	.130	.207
Toolglass	.037	--	.037

Table 3. Wait Times (sec.).

Wait Times

An alternative measure of cognitive load is to consider the time that subjects take between actions, or "wait times." Presumably these measures represent time in which subjects must reflect, plan, make decisions, or prepare for the next action.

Two measures of wait time were calculated: the time that a subject paused without moving either cursor, both after the start of a trial and after making a menu selection (Table 3). (Because the Toolglass technique collapsed the menu selection task and the start of the drawing task into a single RH action, wait time after menu selection was not applicable.)

There was a main effect of technique on wait time between trials ($F_{3,30} = 29.47$, $p < .001$) and wait time after menu selection ($F_{2,16} = 5.20$, $p < .05$). In each case, mean wait times were longest in the L-tearoff condition. The post hoc comparisons also showed that Toolglass differed from R-tearoff for wait time between trials, but revealed no differences between R-tearoff and Palette on either dependent variable. The total of these wait times also differed among techniques ($F_{3,30} = 62.39$, $p < .001$), again showing a marked disadvantage for L-tearoff. Peritz tests confirmed that Toolglass produced the shortest total wait time compared to the other three techniques, with no differences between the Palette and R-tearoff conditions.

Most striking is the disparity in total wait time between the Toolglass and L-tearoff techniques. Interestingly, this difference of .45 sec can almost entirely account for the TCT difference of .53 sec. between these techniques. This suggests that the determination of overall TCT may be more a function of cognitive processing between component subtasks, than the time-motion requirements of the subtasks themselves.

	% errors
R-tearoff	.798
L-tearoff	4.31
Palette	1.29
Toolglass	1.04

Table 4. Sequencing errors (%)

Sequencing Errors

In addition to wait times, cost of an interaction can also be evaluated by examining the kinds of errors that occur. One kind of error in particular was of interest which we call *sequencing errors* and which we define as attempting to draw the connecting line before selecting a menu color.

As can be seen from Table 4, sequencing errors were much more frequent on average in the L-tearoff condition than in the other three conditions. This is confirmed by analysis of variance revealing a main effect of technique ($F_{3,30} = 7.28$, $p < .01$), with Peritz tests confirming that this was due to L-tearoff with no other differences between pairs.

We might expect that sequencing errors would be higher when there is no dependency between the tasks performed by the right and left hands. Presumably, when the action of the right hand depends on the action first of the left hand, the subtasks form part of an intergrated sequence of events, and there is less confusion about which hand to use next.

Menu Distance

Part of the rationale behind the design of the Palette menu was that assigning the function of positioning the menu to the left hand would allow subjects to keep the menu closer to the work area, thereby making it easy to reduce the distance between the drawing cursor and the menu. Indeed, the fact that the left hand could move the menu was the only difference between the Palette and the R-tearoff techniques. As seen earlier, however, space-multiplexing for menu position did not give Palette any advantage over R-tearoff in terms of total TCT. This suggested it was important to check whether or not subjects did move the menu toward the start dot in the drawing task for the three techniques in which this was optional.

Distance between the menu and the start dot at the moment of menu selection was measured. Mean distances were 279 pixels, 245 pixels, and 166 pixels for L-tearoff, R-tearoff, and Palette, respectively. The analysis showed a main effect of technique ($F_{2,16} = 38.65$, $p < .001$) with post hoc comparisons revealing significant differences between all three means. It appears, then, that space-multiplexing of menu position in the Palette condition was effective in reducing menu distance.

Qualitative Ranking

The qualitative ranking of the four techniques is shown in Table 5. Overwhelmingly, subjects preferred using the Toolglass technique (9 subjects) with about an equal number (10 subjects) expressing a dislike for either L-tearoff or Palette. A consistent complaint against the Palette technique was that subjects were not sure what strategy to use when moving the menu, in contrast to Toolglass where "you know where to move the left hand." Similarly, subjects felt that in L-tearoff the left hand was performing too complicated a task, requiring them "to think which hand is doing what." The lack of response with respect to the R-tearoff technique is perhaps not surprising, in view of its familiarity.

	worst	best
R-tearoff	1	1
L-tearoff	5	2
Palette	5	0
Toolglass	1	9

Table 5. Qualitative ranking of the four techniques.

Palette Menu

We were surprised by the relatively poor performance of the Palette menu technique. Our assumption was that since the Toolglass and Palette techniques were both examples of asymmetric dependent assemblage, their performance would be much closer. The difference cannot be explained by lack of motor skills, since the skills requisite for Palette menus were demonstrated in using the other techniques. The reason is more likely cognitive. Evidence for this is in the qualitative evaluations, where subjects expressed confusion about the appropriate strategy with the Palette technique. Clearly the real-world analogy of the painter's palette did not transfer in the time available. This is supported by the quantitative data too. While the menu distance measures reported showed that subjects did keep the palette closer to the drawing area than either tearoff technique, it was still far from optimal, where—like the Toolglass technique—the menu would be brought right to the drawing area. That such near optimal use of the technique is possible is supported by the fact that the motor action involved is essentially the same as that required by the Toolglass technique and that we have observed it with experts, outside of the experiment. While the everyday skills did not transfer within the experiment, over a longer interval, they likely would have. This remains for a further study to investigate.

SUMMARY & CONCLUSIONS

The most general point that emerges from this work is a reinforcement of the findings of Buxton and Myers [2] that subjects can use two hands effectively in performing direct-manipulation tasks. They showed improved performance when compared to the *status quo,* and they did so with little or no additional training. Remember, while subjects were experienced in conventional direct-manipulation techniques, none had seen the *Toolglass* technique before. Clearly, to achieve the performance that was observed in the short interval of the experiment, a significant degree of skill transfer must have taken place.

A second general observation is that using two hands, *per se,* does not necessarily result in improvement over a conventional one-handed technique. When we discuss two-handed input, one of the first reactions of many people is something like, "Oh yes. Let's add a second mouse and cursor and split the load between the two hands."

As was seen in the L-tearoff case, in our experiment, this approach may not result in any improvement, and may in fact degrade performance *despite the fact that far less hand motion is required than in the one-handed case.* Our results showed that there was no improvement in task completion time for the two-handed L-tearoff technique over the one-handed R-tearoff technique. Added to this, the L-tearoff technique was worse than the one-handed technique in terms of measures of cognitive load such as wait time. Further, this kind of configuration results in other detrimental side effects such as sequencing errors, which simply do not occur in one-handed techniques.

One lesson from this is that a motor-sensory Fitts-type analysis alone is not adequate in modeling performance. Otherwise, the L-tearoff would have fared much better than it did (both hands being in "nearly home" position.) Cognitive load must also be considered.

The second lesson is that two hands can be worse than one, if the technique is designed inappropriately. Our analysis, with the help of Guiard [4], suggests that techniques which assign independent subtasks to each hand are of these sort. This is akin to a "tapping the head while rubbing the stomach" approach.

However, the case of the Toolglass technique shows that if two-handed techniques can be designed such that they take into account skills that are already in place, two hands for interaction can be very much superior to one. Subjects in our experiment were able to achieve fast performance, with no indication that using two hands imposed any additional cognitive load over and above the familiar one-handed technique. The subjects both in our experiment and [2] had clearly built up the requisite skills through a lifetime of living in the everyday world, and they were able to apply them in the new situation.

The challenge for the designer is to understand the nature of these skills and recognize how they can be applied in interacting with complex systems. It is in this regard that the work of Guiard [4] is so useful, in that it provides a language for classifying and evaluating bimanual actions. Certainly, the experiment suggests that consistency with his principles is a good starting point for identifying two-handed usage that seems "natural."

In conclusion, many people express the goal of HCI in terms of "making systems easy to use." We use an alternative formulation—one which puts a different spin on the problem: "to accelerate the process whereby novices perform like experts." Based on our study and experience, we believe that two-handed techniques such as Toolglass are examples of how this acceleration can be accomplished. Our hope is that this work will similarly accelerate the process whereby some of these techniques come into common practice.

ACKNOWLEDGMENTS

We would like to thank Francisco Garcia and Joseph Wright of Xerox Research Centre Canada (XRCC) for their help in the actual running of this experiment. We would also like to acknowledge the support of the researchers at Xerox PARC, especially Eric Bier, who developed the Toolglass and Palette Menu techniques. Likewise, we thank Michael Chen and Dan Venolia at Apple Computer, who provided invaluable help in designing the driver that enabled us to use two Apple Desktop Bus devices on the same machine.

This research was undertaken under the auspices of the Input Research Group of the University of Toronto. We are indebted to the members of the group for their input. This work has been supported by the Natural Sciences and Engineering Research Council of Canada, the Information Technology Research Centre of Ontario, and Xerox PARC. This support is gratefully acknowledged.

REFERENCES

1. Bier, E. A., Stone, M. C., Pier, K, Buxton, W., & DeRose, T. D. (1993). Toolglass and magic lenses: The see-through interface. *Proceedings of SIGGRAPH '93.* (pp. 73-80). New York: ACM.
2. Buxton, W., & Myers, B. A. (1986). A study in two-handed input. *Proceedings of the CHI '86 Conference on Human Factors in Computing Systems* (pp. 321-326). New York: ACM.
3. Dillon, R. F., Edey, J. D., & Tombaugh, J. W. (1990). Measuring the true cost of command selection: Techniques and results. *Proceedings of the CHI '90 Conference on Human Factors in Computing Systems* (pp. 19-25). New York: ACM.
4. 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.
5. Oldfield, R. C. (1971). The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia, 9,* 97-113.
6. Robertson, G., Mackinlay, J. & Card, S. (1991). Cone trees: Animated 3D visualizations of hierarchical information. *Proceedings of the CHI '91 Conference on Human Factors in Computing Systems* (pp. 189-194). New York: ACM.
7. Smith, D. C., Irby, C., Kimball, R., Verplank, W. & Harslem, E. (1982). Designing the Star user interface. *Byte, 7(4),* 242-282.
8. Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages. *IEEE Computer, 16(8),* 57-69.
9. Toothaker, L. (1991). *Multiple comparisons for researchers.* Newbury Park, CA.: Sage Publications.