

An Evaluation of Two Input Devices for Remote Pointing

I. Scott MacKenzie¹ and Shaidah Jusoh²

¹Department of Mathematics & Statistics
York University, Toronto, Ontario, Canada M3J 1P3
mack@yorku.ca

²Department of Computer Science
University of Victoria, Victoria, British Columbia, Canada V8W 3P6
shaidah@csc.uvic.ca

Abstract. Remote pointing is an interaction style for presentation systems, interactive TV, and other systems where the user is positioned an appreciable distance from the display. A variety of technologies and interaction techniques exist for remote pointing. This paper presents an empirical evaluation and comparison of two remote pointing devices. A standard mouse is used as a baseline condition. Using the ISO metric throughput (calculated from users' speed and accuracy in completing tasks) as the criterion, the two remote pointing devices performed poorly, demonstrating 32% and 65% worse performance than the mouse. Qualitatively, users indicated a strong preference for the mouse over the remote pointing devices. Implications for the design of present and future systems for remote pointing are discussed.

1 Introduction

Pointing operations are fundamental to human interaction with graphical user interfaces (GUI). The most common pointing device is the mouse, but other devices are also used, such as trackballs, joysticks, or touchpads. Pointing is the act of moving an on-screen tracking symbol, such as a cursor, by manipulating the input device. The tracker is placed over text, icons, or menu items, and actions are selected by pressing and releasing a button on the pointing device.

Remote pointing is an emerging variation of mouse pointing whereby the user is positioned an appreciable distance from the display. Common applications are presentation systems and interactive TV.

This paper presents an empirical evaluation of two remote pointing devices. We begin by describing applications and two representative devices chosen for evaluation. Following this, the method and results are presented.

1.1 Presentation Systems

For *presentation systems*, a typical setup includes a notebook computer and a

projection system driven by the computer's video output signal. The user interacts with the system using the built-in pointing device of the notebook computer, but this displaces the presenter's focus of attention and can diminish the impact of the presentation. Alternatively, a remote pointing device is substituted, allowing the presenter to engage the audience more directly. Pointing operations, when necessary, are performed without physically moving to the notebook computer to acquire the system's built-in pointing device. A variety of technologies are available for remote pointing, and we will present some of these later. The general idea of remote pointing for presentation systems is illustrated in Fig. 1.



Fig. 1. Remote pointing for presentation systems

1.2 Interactive TV

We use the term *interactive TV* with reference to a variety of current and anticipated developments for home computing and entertainment. There is considerable debate on technology convergence in telephone, cable, and internet services, and it is not our intent to enter into this forum here. A reasonable assumption, however, is that home entertainment systems will change substantially in the near future. New applications are anticipated, such as web browsing, email, home banking, travel reservations, and so on. All of these require an input mechanism more sophisticated than a typical remote control. Today's remote control units are optimized for *selecting* (via buttons), but are less capable of *choosing*, *pointing*, or *entering*. Because of the rich task space, a new interaction paradigm is likely to emerge — the *TV-GUI* perhaps. Regardless of the form, it is likely that pointing tasks will be part of the interaction. Indeed, remote controls are now available with built-in trackballs or isometric joysticks, although a standard on-screen interface has not emerged. For interactive TV, therefore, we are interested in investigating a variety of interaction issues, and, in particular, the mechanism of pointing. This is illustrated in Fig. 2.




Fig. 2. Remote pointing for interactive TV

Related to the anticipated insurgence of devices for remote pointing, is the need to empirically test the devices and interaction techniques that users will engage. Will they afford facile interaction or will they encumber the user with an awkward, slow, and error prone interface? Before presenting our empirical study, we describe the technology and devices used.

2 Devices for Remote Pointing

Two devices were chosen for this study. They use different technologies and, therefore, are good choices since they represent different points in the design space. Both function as mouse replacement devices and do not require special interface hardware. They interact with the system's installed mouse driver. Although these are intended for presentations, the interaction styles also apply to interactive TV.

2.1 GyroPoint

The *GyroPoint* is a product of Gyration, Inc. (Saratoga, CA). The device has two distinctly different modes of operation. First, it can function as a regular mouse. It includes a ball mechanism and can operate on a mousepad in the traditional manner. Second, it includes a solid-state gyroscope, permitting operation in the air. For the latter, the angular movement of the hand/device side-to-side or up-and-down maps to x - y tracker motion on the system's display.

According to the manufacturer, "a unique electromagnetic transducer design and a single metal stamping utilize the Coriolis effect to sense rotation. Analog voltages proportional to angular rates around the two sensed axes are provided relative to a

voltage reference output” [5]. The technology has several applications, including computer pointers, TV remote controllers, robotics, factory automation, antenna stabilization, and auto navigation.

There are corded and cordless versions of the *GyroPoint*. The corded version sells for less than one hundred dollars. The cordless version sells for several hundred dollars. In this study we used the corded version.



Fig. 3. A user holding the *GyroPoint*



Fig. 4. A user holding the *RemotePoint*

When operated on the desktop, there are two buttons in the usual position. Two additional buttons for operation in the air are located on the sides near the front of the device. The usual grip for air operation is illustrated in Fig. 3. The grey button under the thumb is the primary button for selection. Another button on the opposite side is operated by the index finger and acts as a clutch. When depressed, rotary motion of the device by the wrist, forearm, and arm maps to x - y motion of the tracker. Motion is relative, so the clutching action of the index finger is the same as lifting and repositioning a mouse on a mousepad.

2.2 RemotePoint

The *RemotePoint* is a product of Interlink Electronics (Camarillo, CA) [6]. It is a cordless pointing device with a built-in isometric joystick, similar to the “eraser tip” joystick on a notebook computer. The device includes an infrared transmitter that communicates with a base receiver at distances up to 40 feet. The base receiver plugs into the system’s mouse port. It is priced under one hundred dollars.

The device is held with the thumb positioned over the joystick which has a large, round rubber casing. The primary button is pressed by the index finger in a trigger-like fashion (see Fig. 4).

Users may position the *RemotePoint* and *GyroPoint* to point at the system's display, but this is not a requirement of the devices per se.

3 ISO Testing of Pointing Devices

All pointing devices are not created equal. Nor will they perform equally. The evaluation of a device's performance is tricky at best since it involves human subjects. (This is in contrast to, say, the performance evaluation of system hardware using standardized benchmarks.) Although there is an abundance of published evaluations of pointing devices, the methodologies are ad hoc. Experimental procedures are inconsistent from one study to the next, and this greatly diminishes our ability to understand or generalize results, or to undertake between-study comparisons. As a consequence, we have much to examine, but we are in a quandary on what it means.

Fortunately, there is a recent standard from the International Standards Organization that addresses this particular problem. The full standard is ISO 9241, *Ergonomic design for office work with visual display terminals (VDTs)*. The standard is in seventeen parts. Part 9 of the standard is called *Requirements for non-keyboard input devices* [8].

ISO 9241-9 describes a battery of tests to evaluate computer pointing devices. The procedures are well laid out and, if followed, will result in a strong and valid performance evaluation of one or more pointing devices.

The basic quantitative test is the serial point-select task (see [13] for an example). The user manipulates the on-screen tracker (viz., cursor) using the pointing device and moves it back-and-forth between two targets, and selects the targets by pressing and releasing a button on the device. A serial task is used because it is easy to implement and affords rapid collection a large quantity of empirical data. The selections are blocked within, say, 20 back-and-forth selections per task condition. As the task is carried out, the test software gathers low-level data on the speed and accuracy of user actions. Three dependent measures form the basis of the subsequent quantitative evaluation: movement time, error rate, and throughput.

Movement time (MT), or task completion time, is the mean time in milliseconds for each trial in a block of trials. Since the end of one trial is the beginning of the next, the mean is simply the total time for a block of trials divided by the number of trials. *Error rate (ER)* is the percentage of targets selected while the tracker is outside the target.

3.1 Throughput

Throughput (TP) is a composite measure of both the speed and accuracy of performance. The measure was introduced by Fitts in 1954 [4], and it has been widely used in human factors and experimental psychology ever since.¹ See [9, 18] for extensive reviews.

Throughput is a very powerful measure. Unlike movement time or error rate, throughput is relatively *independent* of the task difficulty or the range of task difficulties in the experiment. (In fact, this is the very thesis upon which Fitts' original work was based.)

Throughput is calculated as follows:

$$\text{Throughput} = \frac{ID_e}{MT} \quad (1)$$

where

$$ID_e = \log_2 \left(\frac{D}{W_e} + 1 \right). \quad (2)$$

The term ID_e is the effective index of difficulty, and carries the unit "bits". It is calculated from D , the distance to the target, and W_e , the effective width of the target. Since MT , or movement time, carries the units "seconds", throughput carries the units "bits per second", or just "bps".

The use of the *effective* width (W_e) is important. W_e is the width of the distribution of selection coordinates computed over a block of trials. Specifically,

$$W_e = 4.133 \times SD_x \quad (3)$$

where SD_x is the standard deviation in the selection coordinates measured along the axis of approach to the target. Thus, W_e captures the spatial variability or accuracy in a block of trials. As a result, throughput is a measure of both the speed and accuracy of the user. In a sense, throughput reflects the overall efficiency with which the user was able to accomplish the task given the constraints of the device or other aspects of the interface.

It is important to test the device on difficult tasks as well as easy tasks; so, multiple blocks of trials are used, each with a different target distance and/or target width.

¹ Fitts used the term *index of performance* instead of throughput. The term *bandwidth* is also used to imply throughput.

4 Method

4.1 Participants

Our study included twelve paid volunteer participants, composed of nine males and three females. Eleven of the participants were students at a local university; one was a member of the staff. All participants used computers with a GUI and mouse on a daily basis. None had prior experience with remote pointing devices.

4.2 Apparatus

The experiment was conducted on a PC-class desktop computer with a 15" VGA display. The experiment used the *Generalized Fitts' Law Model Builder* software [17] to present the tasks and capture the data from four pointing devices. The following devices were used:

- *GyroPoint*-air by Gyration
- *GyroPoint*-desk by Gyration
- *RemotePoint* by Interlink Electronics
- *Mouse 2.0* by Microsoft

The *GyroPoint* was treated as two devices, since it was tested in each of its two modes: in the air and on the desktop. By having the same device operated in two ways, differences between the "air" and "desk" conditions can be more narrowly attributed to the interaction technique (air vs. desk) rather than to the ergonomics of the device. We also included a Microsoft *Mouse 2.0* in the study, for a similar reason. By including a standard desktop pointing device, we have a "baseline" condition. This is important, because we want not only to evaluate and compare the performance of remote pointing devices, but to provide context for our results in the highly tested world of desktop pointing.

So, of the four device conditions, two are examples of remote pointing (*GyroPoint*-air and *RemotePoint*) and two are examples of desktop pointing (*GyroPoint*-desk and *Mouse 2.0*).

4.3 Procedure

A simple point-select task was used, conforming to the pointing test in Section B.6.1.1 of ISO 9241-9 [6]. The task was explained and demonstrated to participants and a warm-up block of trials was given. In each trial, two targets of width W separated by distance D appeared on the display. A cross-hair tracker appeared in the left target and a purple X appeared in the right target. Participants moved the cross-hair tracker by manipulating the pointing device. The goal was to move the tracker back and forth between the targets and alternately select the targets by pressing and releasing the primary device button.

An error occurred if the centre of the tracker was outside the target when the button was pressed. An error was accompanied by an audible beep. As each selection occurred the purple X moved to the opposite target to help guide the participant through a block of trials. For each distance/width condition, 20 back-and-forth selections were performed. Participants were instructed to proceed as quickly as possible while trying to minimize errors. Approximately one error per block of 20 trials (5%) was considered acceptable. Participants could rest between blocks of trials at their discretion.

For the two desktop device conditions (*GyroPoint*-desk and *Mouse 2.0*), participants were seated in front of the system as customary with desktop computers. For the two remote pointing conditions (*GyroPoint*-air and *RemotePoint*), participants stood in front of the system at a distance of 1.5 meters (see Fig. 5).

The operation of the remote pointing devices was explained and demonstrated. Participants held the device in their hand and manipulated it as noted earlier. For the *RemotePoint*, the base station was positioned beside the display; thus the device was pointing roughly at the display during the trials. For the *GyroPoint*-air, the device operates in relative mode, so it need not point at the display for normal operation. However, as participants also had to visually monitor the display, they held the device with it pointing at the display.



Fig. 5. Operating the *RemotePoint* while standing in front of the display

4.4 Design

The experiment was a $4 \times 3 \times 3 \times 4 \times 20$ repeated measures factorial design. The factors and levels were as follows:

Device	<i>GyroPoint</i> -air, <i>GyroPoint</i> -desk, <i>RemotePoint</i> , <i>Mouse 2.0</i>
Distance	40 mm, 80 mm, 160 mm
Width	10 mm, 20 mm, 40 mm
Block	1, 2, 3, 4
Trial	1, 2, 3 ... 20

To balance for learning effects, participants were randomly assigned to one of four

groups. Each group received the device conditions in a different order using a Latin square.

For each device condition, participants performed 4 blocks of trials. Each block consisted of the 9 target distance/width conditions presented in random order. For each target distance/width condition, 20 trials were performed. Thus, a block consisted of $9 \times 20 = 180$ trials. Participants were able to perform all four device conditions in about one hour. The conditions above combined with 12 subjects resulted in 34,560 total trials in the experiment.

The nine target distance/width conditions were chosen to cover a range of task difficulties. The easiest task combined the shortest distance (40 mm) with the widest target (40 mm). The index of difficulty was

$$ID = \log_2 \left(\frac{D}{W} + 1 \right) = \log_2 \left(\frac{40}{40} + 1 \right) = 1.0 \text{ bits} . \quad (4)$$

The hardest task combined the largest distance (160 mm) with the narrowest target (10 mm):

$$ID = \log_2 \left(\frac{D}{W} + 1 \right) = \log_2 \left(\frac{160}{10} + 1 \right) = 4.1 \text{ bits} . \quad (5)$$

The dependent measures were movement time (ms) error rate (%), and throughput (bps).

At the end of the experiment, participants were interviewed and asked to complete a questionnaire. Comments were sought on their subjective impressions of the four device conditions.

5 Results and Discussion

The grand means on the three dependent measures were 957 ms for movement time, 2.6% for error rate, and 3.0 bps for throughput.

It is possible that even with counter balancing asymmetrical skill transfer effects may occur across device conditions [14]. This would surface as a significant main effect for *order of presentation*, that being the group to which participants were assigned. This was tested for and found not to have occurred as the group effect was not statistically significant on each dependent measure.

The main effects and interactions on each dependent measure are presented in the following sections.

5.1 Speed

The *GyroPoint*-desk was the fastest device condition with a mean movement time of 598 ms. The other device conditions were slower: by 11% for the *Mouse 2.0* (666 ms), by 56% for the *GyroPoint*-air (930 ms), and by 173% for the *RemotePoint* (1633 ms). Thus, the two remote pointing conditions were substantially slower than the two

desktop conditions. The differences were statistically significant ($F_{3,24} = 199.4, p < .0001$). The main effect for block was also significant ($F_{3,24} = 20.7, p < .0001$), as was the device \times block interaction ($F_{9,72} = 3.1, p < .05$). The main effects and interaction are illustrated in Fig. 6.

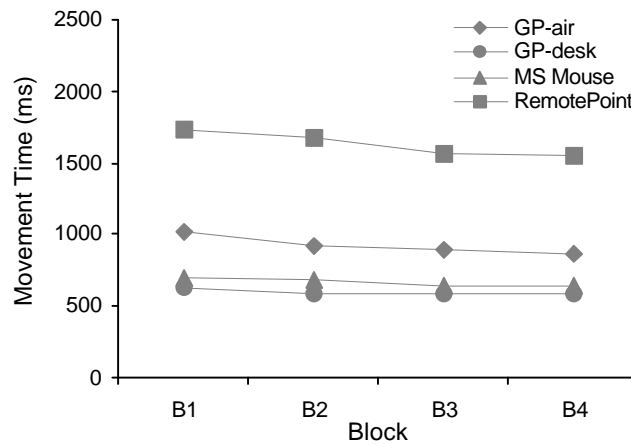


Fig. 6. Movement time by block and device

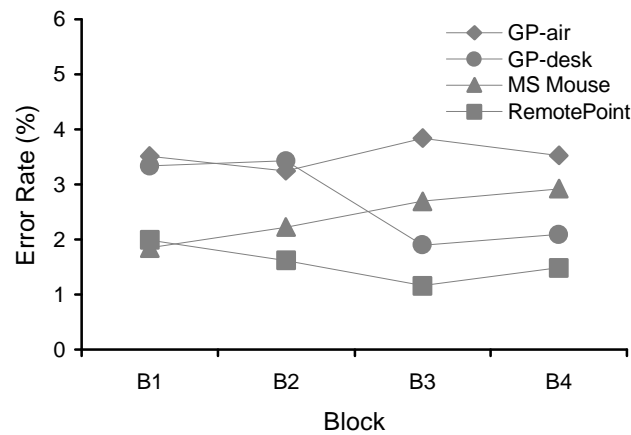


Fig. 7. Error rate by block and device

The very long movement time for the *RemotePoint* is an indication that participants had a difficulty in controlling the movement of the cursor with this device. Problems with finger tremor in controlling isometric joysticks have been noted before [15], so the poor showing here is not surprising.

5.2 Accuracy

The most accurate device was the *RemotePoint* with an error rate of 1.6%. It was followed by the *Mouse 2.0* at 2.4% errors, the *GyroPoint*-desk at 2.7% errors, and the *GyroPoint*-air at 3.5% errors. On the whole, these error rates are quite low. The main effect of device on error rate was significant ($F_{3,24} = 8.6, p < .0005$), but the main effect of block was not ($F_{3,24} < 1, ns$). The device \times block interaction was significant ($F_{9,72} = 3.0, p < .05$). The main effects and interaction are illustrated in Fig. 7.

Standard deviation bars are omitted from Fig. 7 to avoid clutter. However, the lack of statistical significance in the block effect suggests that an interpretation of the trends illustrated over the four blocks in Fig. 7 is not warranted.

The primary observation for accuracy is on the significant main effect of device and the very good showing of the *RemotePoint*. The latter is surprising given the dismal performance on movement time. It is possible that participants proceeded with great caution, given their inability to move expeditiously with this device. Since movement time was much longer, perhaps the very slow positioning tended to make selection — once the cursor was finally positioned inside the target — less error prone.

5.3 Throughput

Throughput is an important dependent measure because it combines speed and accuracy in a single metric, as noted earlier. Thus, tendencies to trade speed for accuracy or vice versa will not necessarily yield differences in throughput. The highest throughput was observed for the *GyroPoint*-desk at 4.1 bps. The other devices exhibited lower throughputs by 10.4% for the *Mouse 2.0* (3.7 bps), by 32% for the *GyroPoint*-air (2.8 bps), and by 65% for the *RemotePoint* (1.4 bps).

These measures are in line with others we have observed. Although there are numerous figures for throughput in the literature, very few were calculated as per the ISO 9241-9 standard, and, therefore, comparisons are difficult. For example, a 1991 study [13] reported throughput for the mouse; however, it was obtained from a regression model of the following form:

$$MT = a + b \times ID_e . \quad (6)$$

The model includes coefficients for the intercept, a in ms, and the slope, b in ms/bit, reported as $a = -107$ ms and $b = 223$ ms/bit. Throughput was reported as the slope reciprocal, namely $1 / 223 = 4.5$ bps. Importantly, the calculation of ID_e included the adjustment noted earlier for the effective target width (see equations 2 and 3). However, the presence of a non-zero intercept tends to weaken the comparison. According to ISO 9241-9, throughput is obtained from the division of means (see equation 1), not from the slope reciprocal in a regression model. All else being equal, the two calculations should yield reasonably similar results provided the intercept is zero, or close to zero.

In our empirical studies, we have recently standardized our calculation of throughput to conform to the ISO standard. As a result, a payoff is now appearing. The payoff is the ability to compare results across studies with confidence that the

comparison is “apples with apples”.

We consistently obtain measures in the range of 3.0 to 5.0 bps for mice [e.g., 10]. Obtaining a similar figure herein is like a “self check” on the experimental apparatus, procedures, data collection, analysis, etc. We also consistently find lower figures for other pointing devices — in the range of 2.6 to 3.1 bps for trackballs [10], 1.0 to 2.0 bps for touchpads [3, 11], and 1.6 to 2.6 bps for isometric joysticks [3, 16].

The main effect of device on throughput was statistically significant ($F_{3,24} = 91.2, p < .0001$), as was the main effect of block ($F_{3,24} = 18.7, p < .0001$) and the device \times block interaction ($F_{9,72} = 2.4, p < .05$). These effects are illustrated in Fig. 8.

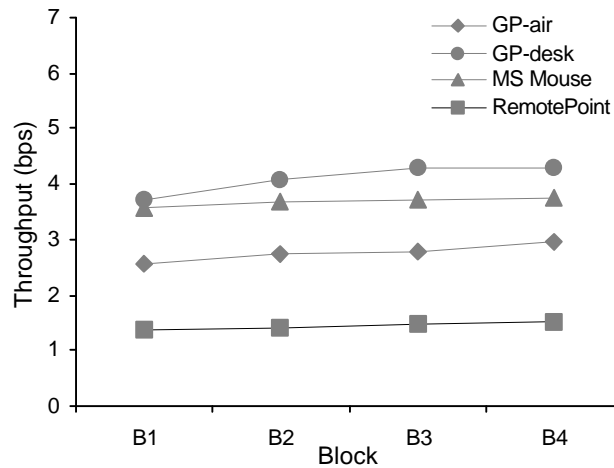


Fig. 8. Throughput (bps) by block and device

The most important observation in Fig. 8 is that the two remote pointing devices fared poorly in comparison to the two desktop pointing devices. These trends are consistent with observations on movement time, but not on error rate. Although, this might suggest that throughput is less sensitive to accuracy than to speed, this is not necessarily the case. In the present experiment the differences in movement time were much more dramatic than those in error rates; so, the similar rank ordering of the movement time results to the throughput results is expected.²

5.4 Qualitative Results

Participants were surveyed for their impressions and their perceived performance. They completed a questionnaire on ease of use, giving each device a rating from 1 (very difficult) to 5 (easy). The results are shown in Table 1. All participants gave the *Mouse 2.0* and the *GyroPoint-desk* the top rating: "easy to use". Nine of twelve participants rated the *RemotePoint* either a 1 "very difficult" or 2 "difficult". The

² Note that throughput is calculated from movement time and “spatial variability”. Error rates per se, although highly correlated with spatial variability, are not used in the calculation of throughput.

GyroPoint-air fared reasonably well with nine participants rating its ease of use as either "neutral" or "slightly easy".

Table 1. Results of Questionnaire on Ease of Use

Participant	Device			
	<i>Mouse 2.0</i>	<i>GyroPoint</i> desk	<i>GyroPoint</i> air	<i>RemotePoint</i>
1	5	5	4	2
2	5	5	3	2
3	5	5	3	4
4	5	5	4	3
5	5	5	3	3
6	5	5	2	1
7	5	5	2	1
8	5	5	4	2
9	5	5	2	1
10	5	5	3	2
11	5	5	4	1
12	5	5	3	1
Mean	5.0	5.0	3.1	1.9

Note: 5 = easy, 4 = slightly easy, 3 = neutral, 2 = difficult, 1 = very difficult

Although the learning trends over the four blocks given in the preceding sections are not dramatic (due to the simplicity of the task), all participants indicated it took effort to get accustomed to the *GyroPoint*-air and the *RemotePoint*. Given that participants had substantial mouse experience but no previous remote pointing experience, it is possible that an extended study would reveal less difference between the devices after prolonged practice.

Participants were also asked to rate their performance with each device, and they all recognized that they took the longest with the *RemotePoint* and were quickest with the *Mouse 2.0* or the *GyroPoint*-desk.

5.5 Experiment and Interaction Considerations

Ideally, in experimental research one tries to hold all variables constant except those under investigation. In this experiment, we had difficulty in achieving this because remote pointing and desktop pointing, by their very nature and by the representative devices chosen, involve a complex shift in interaction styles and psychomotor issues. For example, two devices were operated while seated, two while standing. For three devices, the primary button was activated by the index finger, whereas for one device (*GyroPoint*-air) the button was activated by the thumb. Three devices used position control in which tracker motion was effected by motion of the wrist, forearm, and arm, whereas one device (*RemotePoint*) used velocity control in which motion was effected by force of the thumb.

The use of the corded version of the *GyroPoint* should be considered. Although this may have compromised performance for the air condition, we received no comments on this during the interview.

While we do not feel the issues above impact our overall results or conclusions, it is important to acknowledge the inherent limitations in the experimental design. As remote pointing evolves, further experimentation is warranted in which issues are investigated in isolation (e.g., standing vs. sitting).

Finally, the control-display relationship is altered by the distance of the user from the display. For the seated conditions the eye-to-display distance was about 0.75 meters, for the standing position the distance was 1.5 meters. Thus the "same" targets appeared smaller for the standing position. There is an abundance of inconclusive research on the human performance impact of control-display relationships (e.g., [1, 2, 12]). In one study, viewing distance was actually controlled over a range of 1.5 to 8 meters [7]; however, no difference was found in the operating characteristics of three remote pointing devices. So, we do not feel this is an issue in the present study. For distances greater than a few meters however, display size and the size and form of the tracker are certainly issues to be explored further.

6 Conclusions

This study has shown that two representative devices for remote pointing performed poorly in comparison to a standard mouse. Although the *RemotePoint* had the lowest error rate, it was by far the slowest of the four devices tested and had the lowest throughput, measured as per the ISO standard on pointing devices. Furthermore, subjects gave this device the lowest subjective rating. The *GyroPoint*, while operated in the air, fared slightly better than the *RemotePoint*, but was favoured much less than a standard mouse. These results suggest that remote pointing, while a requirement of some systems, needs further development to support facile interaction.

6.1 Acknowledgements

We thank Shawn Zhang for creating the two concept sketches, and Aleks Oniszczyk for helpful discussions while we were designing the experiment. This research is funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

References

1. Arnault, L. Y., and Greenstein, J. S. Is display/control gain a useful metric for optimizing an interface? *Human Factors* 32 (1990) 651-663
2. Buck, L. Motor performance in relation to control-display gain and target width, *Ergonomics* 23 (1980) 579-589
3. Douglas, S. A., Kirkpatrick, A. E., and MacKenzie, I. S. Testing pointing device performance and user assessment with the ISO 9241, Part 9 standard, *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI '99*. New York: ACM (1999) 215-222
4. Fitts, P. M. The information capacity of the human motor system in controlling the

- amplitude of movement, *Journal of Experimental Psychology* 47 (1954) 381-391
5. Gyration, Inc., Saratoga, California. <http://www.gyration.com/>.
 6. Interlink Electronics, Inc., Camarillo, California. <http://www.interlinkelec.com/>.
 7. Ishiyama, K., and Yano, S. A study of characteristics of pointing devices for television operation, *IEEE Conference on Systems, Man, and Cybernetics - SMC 2000*. New York: IEEE (2000) 1307-1312
 8. ISO. Report Number ISO/TC 159/SC4/WG3 N147: *Ergonomic requirements for office work with visual display terminals (VDTs) - Part 9 - Requirements for non-keyboard input devices (ISO 9241-9)*, International Organisation for Standardisation (May 25, 1998)
 9. MacKenzie, I. S. Fitts' law as a research and design tool in human-computer interaction, *Human-Computer Interaction* 7 (1992) 91-139
 10. MacKenzie, I. S., Kauppinen, T., and Silfverberg, M. Accuracy measures for evaluating computer pointing devices, *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI 2001*. New York: ACM (in press)
 11. MacKenzie, I. S., and Oniszczak, A. A comparison of three selection techniques for touchpads, *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI '98*. New York: ACM (1998) 336-343
 12. MacKenzie, I. S., and Riddersma, S. Effects of output display and control-display gain on human performance in interactive systems, *Behaviour & Information Technology* 13 (1994) 328-337
 13. MacKenzie, I. S., Sellen, A., and Buxton, W. A comparison of input devices in elemental pointing and dragging tasks, *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI '91*. New York: ACM (1991) 161-166
 14. Martin, D. W. *Doing psychology experiments*, 4th ed. Pacific Grove, CA: Brooks/Cole (1996)
 15. Mithal, A. K., and Douglas, S. A. Differences in movement microstructure of the mouse and the finger-controlled isometric joystick, *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI '96*. New York: ACM (1996) 300-307
 16. Silfverberg, M., MacKenzie, I. S., and Korhonen, P. Predicting text entry speed on mobile phones, *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI 2000*. New York: ACM (2000) 9-16
 17. Soukoreff, W., and MacKenzie, I. S. Generalized Fitts' law model builder, Companion *Proceedings of the ACM Conference on Human Factors in Computing Systems - CHI '95*. New York: ACM (1995) 113-114
 18. Welford, A. T. *Fundamentals of skill*, London: Methuen (1968)