Interacting at a Distance: Measuring the Performance of Laser Pointers and Other Devices

Brad A. Myers, Rishi Bhatnagar, Jeffrey Nichols, Choon Hong Peck, Dave Kong, Robert Miller, and A. Chris Long
Human Computer Interaction Institute, School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
bam@cs.cmu.edu http://www.cs.cmu.edu/~pebbles

ABSTRACT
It is difficult to interact with computer displays that are across the room. A popular approach is to use laser pointers tracked by a camera, but interaction techniques using laser pointers tend to be imprecise, error-prone, and slow. Although many previous papers discuss laser pointer interaction techniques, none seem to have performed user studies to help inform the design. This paper reports on two studies of laser pointer interactions that answer some of the questions related to interacting with objects using a laser pointer. The first experiment evaluates various parameters of laser pointers. For example, the time to acquire a target is about 1 second, and the jitter due to hand unsteadiness is about ±8 pixels, which can be reduced to about ±2 to ±4 pixels by filtering. We compared 7 different ways to hold various kinds of laser pointers, and found that a laser pointer built into a PalmOS device was the most stable. The second experiment compared 4 different ways to select objects on a large projected display. We found that tapping directly on a wall-size SmartBoard was the fastest and most accurate method, followed by a new interaction technique that copies the area of interest from the big screen to a handheld. Third in speed was the conventional mouse, and the laser pointer came in last, with a time almost twice as long as tapping on the SmartBoard.

Keywords: Laser Pointers, Remote Interaction, Input Devices, Interaction Techniques, User Studies, Pebbles, Handhelds, Palm Pilots.

INTRODUCTION
In meetings and presentations, the speaker and other attendees often want to point to and interact with large displays they may not be next to. As ubiquitous computing [11] becomes more common, rooms will contain many more devices, appliances, and displays that are computer-controlled, which the user might want to control from across the room. Another motivation for interacting at a distance is interactive TV, such as WebTV and even on-screen TV listings, where users typically sit at a distance from the screen and today use slow step-keys or joysticks. Many research and commercial systems have investigated using laser pointers to interact with screens across the room [2, 4, 5, 9, 12], but these interactions seem awkward and slow.

Our measurements, reported here, show that this is due to inherent human limitations. People do not know exactly where the laser beam will point when they turn it on, and it takes about 1 second to move it into position. People's hands are unsteady, so the beam wiggles, and even filtering techniques cannot improve the accuracy to better than about 2 to 4 pixels. This can be compared to a mouse, where there is basically no wiggle, or to a stylus on a touch-screen, where more than a ±1-pixel wiggle is unacceptable. When the button on the laser pointer is released, the beam often flies away from the target before the beam goes off. Although these problems are well known, we were unable to find any prior studies that measured their magnitudes and effects. We were also interested to know if the design of the laser pointer or the user's stance would affect the amount of wiggle in the beam while trying to hold the pointer steady.

Next, we wanted to compare laser pointing to other techniques for interacting with large screens. We tested a laser pointer, a regular mouse, a wall-sized touch-sensitive SmartBoard with a projected display, and a new interaction technique called "semantic snarfing" [8], where the picture from the big screen is copied to a small handheld and the interaction is performed on the handheld's touchscreen using a stylus.

RELATED WORK
A number of researchers have looked at laser pointer interaction at a distance, but none have reported performance numbers or user studies as described in this paper. Proxima sold commercial data projectors that incorporated a camera for tracking the laser dot [4] and thereby controlled the mouse cursor. Eckert and Moore present algorithms and interaction techniques for a system with a camera looking for a laser dot [2]. They describe interaction techniques, including waiting for the laser pointer to be off for at least one second to signify a mouse action. Kirstein and Muller describe a simple system with minimal interaction techniques for laser point tracking [5]. The Stanford iRoom project is investigating the use of a laser pointer with special gestures and pie menus for the interaction [12]. The XWeb system provided a

(a) (b) (c) (d)

Figure 1. Devices from the first experiment: (a) Conventional, inexpensive laser pointer, (b) pen-like laser pointer, (c) Symbol SPT 1700 Palm with built-in laser scanner / pointer, and (d) laser pointer mounted on a toy gun. *(simulated laser beams)*

variety of laser pointer interaction techniques, including ways to use a laser pointer for menu selection, scrolling, and Graffiti-based text entry [9]. The times to detect *laser on*, *off,* and *dwell* in XWeb were at least one second each. In order to draw Graffiti accurately, the strokes had to be almost full screen, and text was entered at about 1/3 the speed for text entry on the Palm.

MacKenzie and Jusoh investigated various techniques for interacting at a distance with a computer display [6]. They compared two forms of “air mice”: a GyroPoint with a gyroscope that measured angular movement in the air, and a RemotePoint mouse with a joystick for moving the cursor. They found the GyroPoint was about 56% slower than a regular mouse, and the RemotePoint was 173% slower. Error rates were also higher for the air mice. We were interested in comparing our interaction devices with the GyroPoint and the RemotePoint, so we used MacKenzie and Jusoh’s setup in our second experiment.

EXPERIMENT 1: LASER POINTER PERFORMANCE

We were frustrated that prior papers about laser pointer interaction techniques did not differentiate between problems due to limitations of today’s technology (camera resolution, processor speed, etc.) and problems inherent in the human use of laser pointers. We felt studying these issues would improve the design of future interaction techniques. Therefore, we performed two studies to discover some of the fundamental parameters of laser pointing and, as a result, aid in the design of more usable laser interaction techniques. Specifically, we hoped to discover answers to some of the questions related to selecting objects on a projection screen using a laser pointer.

We measured how long people take to acquire a target, how steadily people can keep their beam on the target, and whether people can turn off the beam within their intended target. We also measured accuracy at different distances. We conjectured that the *angle* of wiggle should be the same at all distances, so the *diameter* of wiggle would be proportional to the distance. We further conjectured that different laser pointer devices would perform differently, and might even affect whether the laser beam wiggle would be different in the horizontal and vertical directions.

Laser Pointers for the First Experiment

The first experiment used 4 laser pointer devices held in various ways. The first device was a conventional inexpensive laser pointer (Figure 1-a) held both near the body and with the arm fully outstretched. All participants held this laser pointer between their thumb and forefinger (as shown in Figure 1-a) because of the pointer’s small size. The second device was a pen-like laser pointer (Figure 1-b), which we asked participants to hold like they would hold a pen (as shown in 1-b). We speculated that the pen grip might be more stable for precise movements and for holding still. The third device was a Symbol SPT 1700 handheld device [10], which runs the PalmOS software and has a built-in laser and radio transmitter (Figure 1-c). The Symbol laser is designed for bar code scanning, but on the SPT 1700 model, the rotating mirror can be turned off so the laser will be stationary and can serve as a laser pointer. We were particularly interested in the SPT since it provides a single platform that serves as both a pointing and interacting device.

Finally, we thought about various situations where people need to hold their hands steady. First, we investigated mounting a laser pointer to a fencing foil grip, but this proved technically difficult and the grip was essentially the same as for the other laser pointers. Next, we tried mounting the laser pointer to a piece of a glove worn on the index finger, so the users could just point with their finger. Unfortunately, this did not work either, because the laser pointer was too heavy to support with only the index finger, and users did not like wearing the glove.

For the fourth device, we settled on a laser pointer mounted on a toy handgun (Figure 1-d) since guns are designed to point at a target and remain stable. We wanted the laser beam to turn on when the user pulled the gun trigger, but this proved to be too challenging, so instead we made the beam stay on continuously. For the other devices, the user turned the beam on and off in each trial using the button on the laser pointer.

Detecting the Laser Beam

To measure the accuracy and path of the laser beam, we wanted very high resolution so that we could investigate fundamental human limitations instead of camera and processor limitations. Therefore, we placed a camera 3 inches behind a piece of tracing paper with a small dot drawn as a target in the center (Figure 2). The camera view was about 3¼ inches by 3¼ inches. To get this resolution on a typical projected screen at a distance of about 5 feet would require the camera to have a resolution of about 7680×7680, which may be possible in the future but is not feasible today.

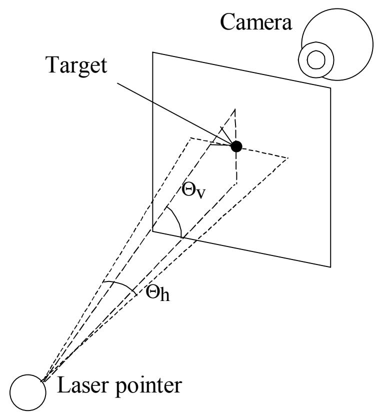

Figure 2. Experimental setup.

We used a Sony Digital Video camera with a Winnov image capture board in our PC. Our laser tracker software is written in MFC Visual C++ and uses the Video For Windows API for grabbing frames from the device. We grab each frame and subtract it from the previous frame to get a difference frame. In this difference frame, we find the laser dot by looking for the pixel with the highest red value that is above the noise level. On a 450MHz Pentium III, we can process images with a 320×240 resolution in RGB format at a rate of 18–20 frames per second.

For each trial, the user was asked to aim at the target indicated on the tracing paper using one of the laser pointers. Figure 3 shows a sample plot that was captured from our program on one trial. The plot is superimposed over the actual view captured by the camera. The horizontal lines were on the original tracing paper, and are ¼ inch apart. The system was calibrated by indicating the position of a one-inch line that was drawn on the paper. The experimenter placed the innermost square in the center over the picture of the target dot at which the user was aiming, and then hit the "Place Target" button to configure the software. The computer program recorded each trial in three phases. The acquisition phase recorded the entry trail of the laser as the user moved towards the target. In the trace, the user's starting position is marked with an "S", and the portion of the trace before target acquisition is shown in yellow on the screen (from the "S" to the center).

The second phase recorded the dwell of the laser point on the target. When the target was acquired, the computer beeped and the user tried to hold the laser pointer as steadily as possible on the target for 3 seconds. The part of the trace near the center in Figure 3 is this phase, and is shown in red on the screen. A second beep signaled the user to turn off the laser. The third phase recorded the laser's exit trail. This goes to the "E" in Figure 3, which is where the beam disappeared, and is shown in purple on the screen. The polygon shows the convex hull of the points during the dwell phase, and the text display at the bottom shows the area of this polygon in square inches. The larger square in the middle of Figure 3 is centered on the average of all the points recorded during the second phase, and shows that for this trace, the average position of the beam during the dwell phase was very close to the target position. The size of the larger square in the figure does not represent any measurement taken by the system.

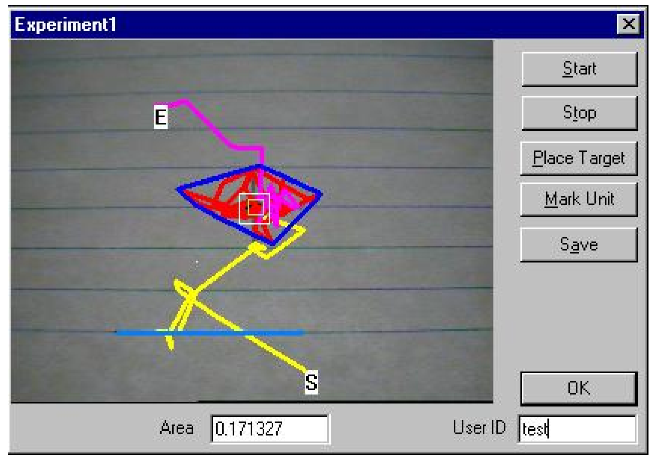

Figure 3. Sample plot as captured by the camera.

Experimental Setup

Seventeen participants between the ages of 21 and 31 took part in this experiment, 5 women and 12 men. After the experiment, the users rated their experience using a computer and a laser pointer on a questionnaire. The mean scores of the users, on a scale of 1 to 7, were 5.7 for computer experience (high) and 2.3 for laser pointer experience (low).

Each participant used seven laser pointers at three different distances from the target: 5, 10 and 15 feet. The seven different devices were: (1) the laser pointer of Figure 1-a held close to the body, (2) the same laser pointer held at arms-length (extended) as if the user were pointing at something, (3) the Symbol device of Figure 1-c held in one hand, (4) the Symbol device held with two hands, (5) the pen laser held in a pen grip, as if the user was writing on the wall, as shown in Figure 1-b, (6) the same pen laser of Figure 1-b held like the small laser pointer as shown in Figure 1-a (the "normal" way to hold a laser pointer), and (7) the laser pointer mounted on the toy gun shown in Figure 1-d.

This was a within-subjects study, and each participant performed the task once at each of the 3 distances with 7 devices, to make 21 trials per participant. The order of the devices and distances were randomized to avoid ordering effects. The study took about 10–15 minutes for each participant to complete.

Results of the First Experiment

The average time across devices that participants took to acquire the target in phase one was about 1 sec. The minimum was 0.1 sec. and the maximum was about 3.5 sec. We could not measure the maximum *distance* that the beam started away from the target, because it often started out of the camera's view. We observed that it was not unusual for the beam to start several inches away from the target.

We found no significant effects on the deviation of the beam (hand wiggle) correlated with the ordering, which means that there was little effect of learning or fatigue in this experiment. There was a significant effect by user ($p < 0.001$), with the largest wiggle being about four-times the smallest: the deviation ranged from 0.14 to 0.6 inches between the most and least accurate users.

Distance and device yielded the most interesting results. Table 1 and Figure 4 show the average wiggle across all devices by user's distance from the target. This distance is the average Euclidean distance from the target measured in inches. As predicted, the amount of wiggle gets worse as the user gets further away ($p < 0.001$). Converting the wiggle into an angle shows the average angle of deviation for each distance. However, the angle is not constant across the three distances. We conjecture that people had less angular wiggle (although still higher magnitude of wiggle) because they were concentrating harder and the size of the laser dot was larger, which affected the camera's view.

There is also a significant difference in wiggle by device ($p < 0.001$). The most stable device was the SPT 1700 held with 2 hands, followed by the SPT 1700 held with only one hand, the pen held the normal way, and the small laser pointer held the normal way (Figure 5). Users performed poorly with the gun pointer as well as with the pen pointer when held as a pen, contradicting our predictions. Least stable was the laser pointer held with the hand outstretched.

Across all devices, we find a statistically significant difference for the horizontal vs. vertical angles ($p < 0.001$). In particular, the vertical deviation is more, sometimes by a factor of 10, than the horizontal deviation, so people seem to be steadier left-to-right than up-and-down. We did not find that the ratio between vertical and horizontal deviations was affected by device.

We were also interested in the accuracy that could be achieved by filtering the stream of points recorded from the laser. Figure 6 shows a graph of the horizontal deviation for a typical trial for one user. The DX TIME plot is the distance from the target. The CUMX TIME plot is the cumulative average of all the points that have been recorded so far. The AVGX TIME plot is filtered using a ½ second moving window. The graph shows that using the ½ second window, the moving average still varies about ±0.20 inches from the target. The cumulative average shows that using a longer filter would still only get to within ±0.10 inches of the target. For some users, the cumulative average did not get much closer than ±0.20 inches from the target.

Table 1. Average deviation for each device in radial distance from target (inches) and angle (degrees).

(dist in inches, angle in degrees)	5 feet		10 feet		15 feet	
	dist	angle	dist	angle	dist	angle
1. laser close	0.17"	0.16°	0.29"	0.14°	0.36"	0.11°
2. laser extend	0.18"	0.17°	0.35"	0.17°	0.46"	0.15°
3. Symbol 1 hand	0.18"	0.17°	0.28"	0.14°	0.41"	0.13°
4. Symbol 2 hands	0.15"	0.15°	0.24"	0.11°	0.31"	0.10°
5. pen w/pen grip	0.19"	0.18°	0.28"	0.13°	0.39"	0.12°
6. pen normal grip	0.17"	0.16°	0.27"	0.13°	0.33"	0.11°
7. gun	0.18"	0.17°	0.26"	0.13°	0.42"	0.13°

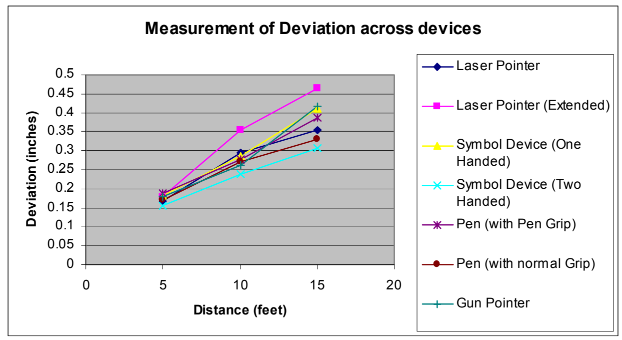

Figure 4. Graph of deviation by distance.

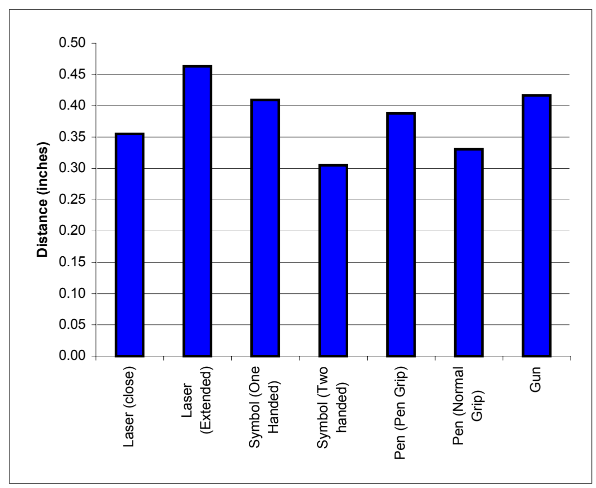

Figure 5. Graph of deviation by device at 15 feet.

We also found that when we told the users to turn the beam off, it took an average of about ½ second for the beam to actually disappear. In this time, the beam would often drift from the target, sometimes farther than we could measure (about 1½ inches).

On the questionnaire at the end of the experiment, participants were asked to rate the devices for ease of use on a scale from 1 to 7, with 1 being best. Users rated the SPT 1700 with two hands as the easiest to use with an average score of 2.6. In comments they mentioned that they liked that the pen pointer and the laser pointer were short and nice looking. These devices received average scores of 2.98 and 2.94 respectively. Most of the users liked the idea of using a laser pointer mounted on a toy gun, but mentioned that they would be uncomfortable using it in a conference. They also mentioned that the SPT 1700 was too heavy and cumbersome. One user mentioned that he would like to take it to a technical conference because “it would look really cool.”

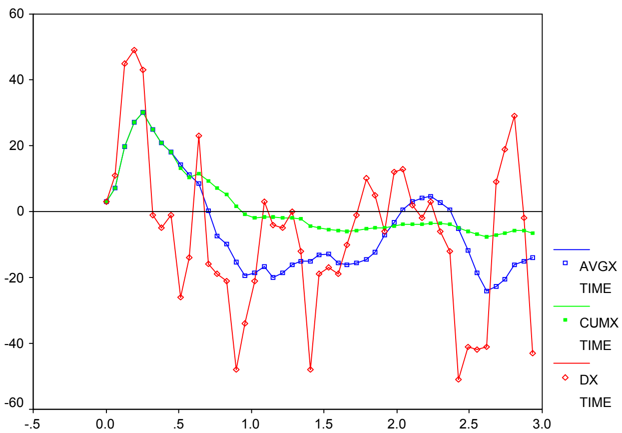

Figure 6. Example trace of data recorded for one trial. The x-axis is seconds. The y-axis is distance, measured in 1/100th of an inch.

Discussion of the First Experiment

In the first experiment, we found that users could not turn the beam on where they wanted, could not hold it still where they wanted, and could not turn it off where they wanted. We found that our participants needed at least 1 second before the beam was where they intended. For a typical projection screen in a small room where the user is 5 feet from a 6-foot-wide screen at 1024x768, the size of the wiggle will be about ±3 pixels using the measured angles. Standing 10 feet from a large 8-foot-wide screen, the wiggle will be about ±4 pixels. This implies that widgets designed for laser interaction must be fairly big. Even using a moving average filter will not improve the accuracy by much, since we still are only getting within ±2 to 4 pixels of the target. In reality, the amount of wiggle will be larger since our numbers assume an ultra-high resolution camera. The low resolution of today’s cameras increase the inaccuracy by about a factor of 3.

These results mean that to correctly track a user’s dwell on a location with any accuracy, we have to track the laser for about 1 second to wait for users to get the beam where they want it, and then wait another ½ second to get a moving average. This means that about 1½ second are required to execute a selection. If a start and end-point pair is desired (e.g., when dragging an object), the location where the beam disappears cannot be used, so the system must detect the first and last dwell-points, which doubles the total delay.

The moving average may also create delays during movement. Suppose a laser tracking system uses a moving average with a window of ½ second to improve accuracy. When the user moves the laser beam from one side of the screen to the other, the position of the cursor will lag behind with slightly more than a ½ second delay. To solve this problem, we suggest the use of a distance threshold that governs when the moving average window is used. When the distance between two successive points is sufficiently large, the system guesses that the user is moving a long distance and is not worried about accuracy. Since we measured the wiggle at about ±4 pixels, it is reasonable to assume that once successive points are closer than 8 pixels, the system should begin applying the moving average, since it is likely that the user is beginning to home in on a particular target. We have found that this noticeably improves the usability of a laser tracking system with a moving average.

Different types of laser pointers can have an impact on the accuracy of the tracking, but our attempts at creating a more stable laser pointer were not successful. It seems that the wiggle is an inherent property of human pointing and cannot be fixed by device design. The “gun” grip performed poorly. Holding a laser pointer like a pen was also not helpful. It is no surprise that holding a solid, heavy device (the Symbol SPT 1700) with two hands would be the most stable or that holding a laser pointer outstretched is the least stable.

We conjecture that people wiggle more vertically than horizontally across all devices due to properties of people’s muscles. However, we did not find that shape of the wiggle correlated with the device even though different devices were held differently, and presumably used different muscles.

EXPERIMENT 2: TRACKING ACROSS THE ROOM

In order to experiment with using a laser pointer for interacting at a distance, we needed a laser tracking system that would work across the area of an entire projected display, rather than the tiny area used in our first experiment. We initially tried pointing an inexpensive camera at a front-projected SmartBoard, but had little success. The problem was that the automatic gain-control on cheap cameras made the bright laser dot saturate to white and become indistinguishable from other white areas of the screen. Therefore, we needed to use a more expensive camera which has a manual brightness control. Turning the brightness nearly off allowed the camera to still see the laser dot, but little of the projected image. This made the tracking quite easy and accurate. Also, the more expensive camera did not have the curved-image problems that were reported by Olsen [9]. We used the same Winnov image capture board and software as for the first experiment. We were again able to track the laser dot with a 320x240 resolution at a rate of 18–20 frames per second.

Configuring the laser tracker was a two-step process. First, the user pushed the “Calibrate” button, which caused the software to calculate the normal noise level of the video system by looking for the maximum difference in the value of corresponding pixels in successive frames without any laser or other change to the picture. Next, the user indicated four points on the screen. The four points enabled the software to perform a linear mapping of the camera's view to the screen coordinates, since the image seen by the camera was often trapezoidal.

To control the PC's cursor, we connected the laser tracker's results to our RemoteCommander software [7], which maps operations on handhelds' screens to the PC's cursor and keyboard actions. We used a special PC-side version of the RemoteCommander client that accepts input from the laser tracker program as if it had come from a handheld. Since we were exploring multi-machine user interfaces, we assumed that the function usually performed by the mouse buttons would be performed by another device. In the future, we might add the ability for a dwell of the laser dot to signify a button press, as in other laser tracking systems.

Currently, our laser tracker software is only set up to track a single laser dot. Techniques described elsewhere (e.g., [12]) could be used to track multiple dots.

Experimental Setup

In this experiment, we wanted to use our laser tracking setup to see how the speed and accuracy of interaction using the laser pointer compared to other techniques such as a mouse. Therefore, we repeated the classic pointing experiment from Card, Moran and Newell (Experiment 7B, pp. 250) [1], which is based on Fitt's "dotting task" [3]. In particular, we used the setup described by MacKenzie in studying air mice [6]. A picture of the test application is shown in Figure 7.

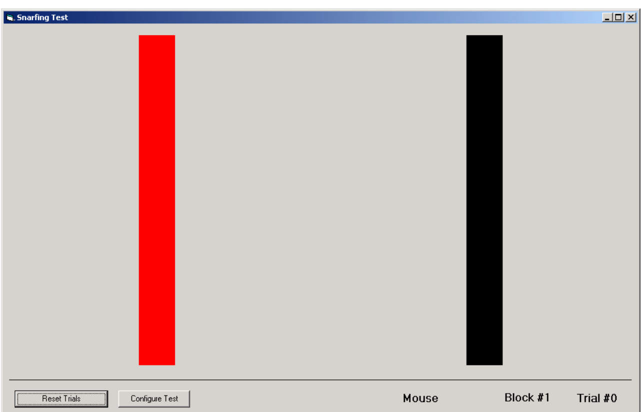

Figure 7. Screen for experiment 2.

The test consisted of tapping back and forth between the two targets. The targets had one of three widths: 27, 53, and 107 pixels, which were about 1.5, 2.75, and 5.5 inches on our SmartBoard. The targets were also separated by one of three distances: 107, 213, and 427 pixels, which were about 5.5, 11, and 22 inches on our SmartBoard. Figure 7 shows the most difficult condition, with the smallest bars and the widest separation.

We tested four devices, including a conventional mouse that served as a point of comparison with other experiments. A table was placed in front of the participants so they could use a mouse in the conventional way, except that they watched the large SmartBoard display instead of a monitor (see Figure 8).

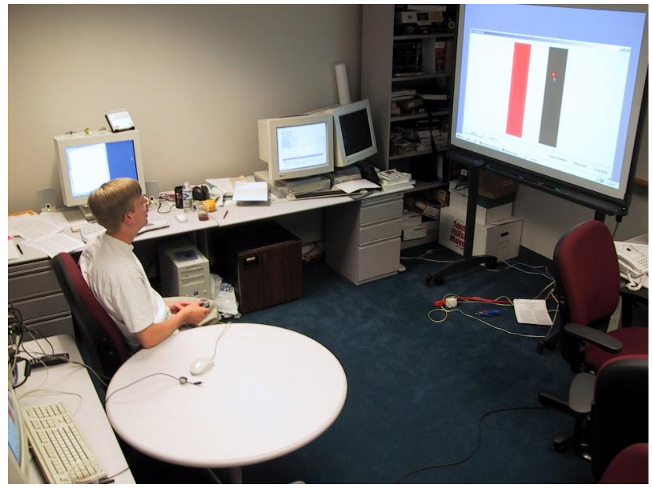

Figure 8. Configuration for experiment 2 with a participant using the laser pointer. The participant is seated about 5 feet from the projected screen. For the SmartBoard condition, the participant stood at the SmartBoard and tapped on the bars with one or two fingers.

We used the Symbol SPT 1700 as the laser pointer device in the second experiment, because it was the most stable of the pointers in the first experiment. Another advantage of the SPT 1700 was that we could use the buttons on the device to signal when the laser pointer was at the desired location. This was equivalent to pressing the mouse button. Other laser pointer interaction techniques [2, 9] required the user to dwell at a point to signal a press. Our first experiment showed that the dwell time to signal a location needed to be at least one second, which would make it inappropriate for use in a timed test. The participant could press any physical button on the Symbol to emulate pressing the mouse button (see Figure 1-c). Users were seated and held the SPT 1700 in their dominant hand and pressed the buttons with the same hand. We restricted the participants to only one hand because all of the other devices were used with only one hand. The users were allowed to rest the back of the device on their lap or on the table, as they felt comfortable. In this condition, the laser beam was always on, so the times for acquiring and leaving the target measured in the first experiment were not relevant.

In the third condition, the participants tapped directly on the touch-sensitive SmartBoard. For this condition, participants stood in front of the SmartBoard and used their dominant hand to tap on the board with one or two fingers.

Finally, the participants used a technique called "semantic snarfing" [8] where the picture from the main screen was captured onto the small screen of a PDA and the interaction occurred on both the small and big screens. We found that this is a good way to compensate for the poor resolution of laser pointer interaction: the laser can be used to signify the general area of interest, which is then copied to the small device for detailed work. Later, the results are sent back to the big device. In our snarfing condition, the participants were seated in front of the SmartBoard, holding a Compaq iPaq on which a smaller version of Figure 7 was displayed. We used the iPaq sideways so we could get a greater distance between the targets. On the iPaq, the width of the targets ranged from 9 to 36 pixels (3/32 to 3/8 inches), and they were from 36 to 144 pixels (3/8 to 1.5 inches) apart.

In our pilot tests, we also tried using the remote control device that came with our video projector (an Epson Powerlite 710C). The remote control has a built-in joystick that moved the computer's cursor. Pressing down on the joystick signals a mouse click. However, our pilot participants found this very hard to use, and it seemed to be about 2 to 4 times slower than the other devices. Since it was also very much like the *RemotePoint* device in the MacKenzie study [6], we decided to omit it from our main study.

This was a within-subjects experiment, so all participants used all four devices. For each device, each participant performed four blocks of trials. Within each block, participants completed twenty trials of clicking back-and-forth among the target, for each of the nine possible configurations of targets (3 widths × 3 positions). Each participant performed 4×4×3×3×20 = 2880 clicks during the experiment.

Participants were randomly assigned to one of four groups to spread the learning effects of the experiment equally among all devices. Each group saw the devices in a different order based on a Latin Square.

Sixteen participants between the ages of 19 to 28 took part in this experiment, 4 women and 12 men. The participants were paid for participating. After the experiment, the participants filled out a questionnaire on their proficiency with using a laser pointer and a PDA. The mean score of the users were, on a scale of 0 to 7, 2.375 for laser pointer and 2.94 for the PDA. Seven of the sixteen participants in this experiment owned a handheld computer and thirteen had some experience with one.

Results of Experiment 2

We recorded three metrics of our participant's performance: movement time, error rate, and throughput. Throughput is a measure of pointing device performance that is independent of task difficulty and has the units of bits per second (bps). We calculated effective throughput according to the method used in [6]. The average of each of these measures across all trials is shown in Table 2 and Figure 9. All devices were significantly different from all others for each of the three metrics ($p < 0.002$ for all comparisons).

We found that the laser pointer was the worst device in terms of both movement time and throughput. Only semantic snarfing was worse than the laser pointer in the error metric. The SmartBoard device was the best in terms of speed and errors, with more than 50% fewer errors in 16% less time than the next best device in those categories. The throughput of the Snarfing device is probably high in spite of the large error rates probably because users tapped repeatedly in the same location, which decreased the effective width. By way of comparison, the throughput reported in the original Fitt's tapping task for the best 8 out of 16 was between 10.3 and 11.5 [3].

Table 2. Results of Experiment 2

Device	Movement Time (ms)	Error Rate	Throughput–bps (std. dev.)
Mouse	628	4.34%	6.98 (1.50)
SmartBoard	473	1.94%	11.80 (3.93)
Laser Pointer	930	6.26%	5.08 (1.47)
Semantic Snarfing	562	8.44%	13.13 (5.99)

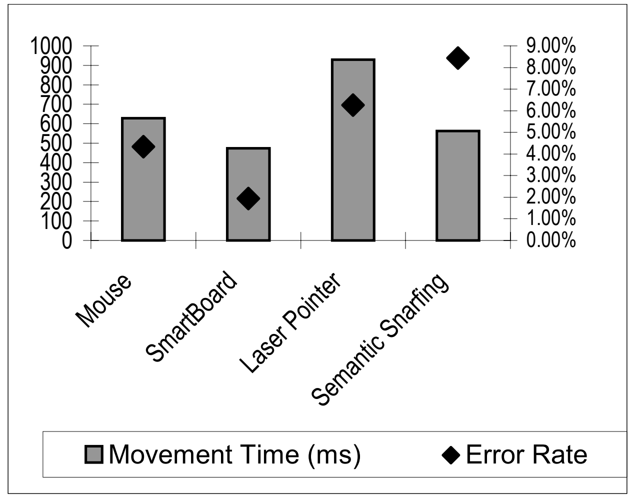

Figure 9. Results from Experiment 2.

On the questionnaire at the end of the experiment, participants were asked to rank the four devices in order of their perceived performance. Participants correctly identified that the SmartBoard performed the best, giving it an average ranking of 1.28 with 12 first-place votes. The mouse was second with an average ranking of 2.00 and 3 first-place votes. Participants slightly preferred snarfing to using the laser pointer, giving them average rankings of 3.19 and 3.53, respectively. The snarfing device also received one first place vote.

Participants were also asked to rate the devices for ease of use on a scale from 1 to 5, with 5 being easiest. Again, the SmartBoard received the highest marks with an average rating of 4.75. The mouse, snarfing, and laser pointer devices followed in that order, with ratings of 4.13, 3.5, and 2.56 respectively.

Discussion of the Second Experiment

Overall, the laser pointer performed the worst of any pointing device in this experiment. It finished last in every comparison except for error rate, in which the snarfing device was worse. Based on the throughput measurements, a user would be able to accomplish only about half as much work using the laser pointer as they would with the SmartBoard in the same amount of time.

The results of the snarfing device in this experiment were also interesting. It was quick (10% faster than the mouse and just 16% slower than the SmartBoard) but suffered from a higher error rate because the interface was scaled down, which made some of the targets very small and difficult to hit reliably. Almost all of the errors made with the snarfing device—95% of them—were made with the two smaller targets. This is 11% more than the laser pointer, which we might also expect to perform badly with smaller targets. In a pilot study, we experimented with an unscaled image of the interface on the iPaq that the user had to scroll, but this performed much worse and the participants did not like scrolling. To improve the snarfing device, we need to find a better way to control scrolling and zooming.

An interesting finding is that the SmartBoard and snarfing techniques both performed better than the mouse in movement time. This suggests that direct tapping interfaces can perform better than indirect devices such as the mouse.

Comparing our results to MacKenzie's [6], we had approximately the same values for the mouse (he tried two mice with times of 598 and 666 compared to our 628). His time was 930 ms for the GyroPoint air mouse that used a gyroscope, and 1633 ms for the RemotePoint that had a joystick. Thus, these devices fared the same or worse than the laser pointer, and much worse than the snarfing technique for interacting at a distance.

FUTURE WORK

Since the laser pointer performed poorly, as predicted by other laser pointer research, but the snarfing technique performed relatively well when targets were not too small, we believe there is a good synergy between these two techniques. If the laser pointer were used to indicate a portion of an interface to snarf, and the user did fine-grain interaction with a handheld touch-screen, such as the snarfing device, we expect the overall pointing interaction would be improved. Also, this may eliminate the need for scrolling on the snarfing device. We are currently exploring this type of interface.

We are very interested in exploring other new ways to combine interactions across multiple devices, in what are called "multi-machine user interfaces." One example is using a laser pointer for remote pointing built into a Palm device for local interaction is. Many other combinations are possible. Now that we know some of the parameters of laser pointer interaction, we will try to develop new laser pointer-only interaction techniques as well as multi-modal interaction techniques that incorporate a laser pointer, and which can best take advantage of the laser pointer's properties.

CONCLUSIONS

Prior techniques for using laser pointers with computers have been largely unsuccessful. Our studies show why conventional interaction techniques designed for a mouse or a stylus are doomed to fail when used with a laser pointer. The beam is too unsteady, users cannot turn the beam on or off where they want, and there is no mouse button. However, it is quite natural and convenient to be able to point to things across the room using a laser pointer. As computerized devices become increasingly ubiquitous and interacting at a distance becomes more common, it will be important to provide interaction techniques that are quick, convenient, and accurate. Hopefully, the data from the studies reported here will be helpful for designing those systems.

ACKNOWLEDGEMENTS

For help with this paper, we would like to thank Al Corbett, Scott MacKenzie, and Jake Wobbrock.

This research is supported by grants from DARPA, Microsoft and the Pittsburgh Digital Greenhouse, and by equipment grants from Symbol Technologies, Palm, Hewlett-Packard, Lucent, IBM and SMART Technologies, Inc. This research was performed in part in connection with contract number DAAD17-99-C-0061 with the U.S. Army Research Laboratory. The views and conclusions contained in this document are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory or the U.S. Government unless so designated by other authorized documents. Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

REFERENCES

1. Card, S.K., Moran, T.P., and Newell, A., *The Psychology of Human-Computer Interaction.* 1983, Hillsdale, NJ: Lawrence Erlbaum Associates.
2. Eckert, R.R. and Moore, J.A., "The Classroom of the 21st Century: The Interactive Learning Wall." *SIGCHI Bulletin*, 2000. **23**(2): pp. 33-40.
3. Fitts, P.M., "The information capacity of the human motor system in controlling amplitude of movement." *Journal of Experimental Psychology*, 1954. **47**: pp. 381-391.
4. Horn, G.A.V., "Proxima's new Ovation+ projection panels do up multimedia." *Byte (on-line)*, 1995. http://www.byte.com/art/9501/sec12/art9.htm.
5. Kirstein, C. and Muller, H. "Interaction with a projection screen using a camera-tracked laser pointer," in *Multimedia Modeling; MMM '98 Proceedings.* 1998. pp. 191 -192.
6. MacKenzie, I.S. and Jusoh, S. "An Evaluation of Two Input Devices for Remote Pointing," in *Proc. of the 8th IFIP Working Conf. on Engineering for HCI - EHCI 2001.* 2001. Heidelberg, Germany: Springer-Verlag. pp. To appear. http://www.yorku.ca/mack/ehci01.html.
7. Myers, B.A., "Using Hand-Held Devices and PCs Together." *Communications of the ACM*, 2001. **44**(11): pp. 34-41. http://www.cs.cmu.edu/~pebbles/papers/pebblescacm.pdf.
8. Myers, B.A., Peck, C.H., Nichols, J., Kong, D., and Miller, R. "Interacting At a Distance Using Semantic Snarfing," in *ACM UbiComp'2001.* 2001. Atlanta, Georgia: pp. 305-314.
9. Olsen Jr, D.R. and Nielsen, T. "Laser Pointer Interaction," in *ACM CHI'2001 Conference Proceedings: Human Factors in Computing Systems.* 2001. Seattle, WA: pp. 17-22.
10. Symbol Technologies, "SPT 1700 Pocketable Computers," 2001.http://www.symbol.com/products/mobile_computers/mobile_palm_pi_hdwr_spt1700.html.
11. Weiser, M., "Some Computer Science Issues in Ubiquitous Computing." *CACM*, 1993. **36**(7): pp. 74-83. July.
12. Winograd, T. and Guimbretiere, F. "Visual Instruments for an Interactive Mural," in *ACM SIGCHI CHI99 Extended Abstracts.* 1999. Pittsburgh, PA: pp. 234-235. http://graphics.Stanford.EDU/projects/iwork/papers/chi99/.