Laser Pointers as Collaborative Pointing Devices

Ji-Young Oh, Wolfgang Stuerzlinger

Department of Computer Science, York University
http://www.cs.yorku.ca/~wolfgang

Abstract

Single Display Groupware (SDG) is a research area that focuses on providing collaborative computing environments. Traditionally, most hardware platforms for SDG support only one person interacting at any given time, which limits collaboration. In this paper, we present laser pointers as input devices that can provide concurrent input streams ideally required to the SDG environment.

First, we discuss several issues related to utilization of laser pointers and present the new concept of computer controlled laser pointers. Then we briefly present a performance evaluation of laser pointers as input devices and a baseline comparison with the mouse according to the ISO 9241-9 standard.

Finally, we describe a new system that uses multiple computer controlled laser pointers as interaction devices for one or more displays. Several alternatives for distinguishing between different laser pointers are presented, and an implementation of one of them is demonstrated with SDG applications.

Keywords: Single display groupware, hardware for collaboration, laser pointer, input devices.

1 Introduction

General computer systems are designed to support interaction with only one user at a time. When multiple users want to use a computer collaboratively, they are currently forced to take turns to control the system.

Single Display Groupware (SDG, see e.g. [8]) is defined as a system that can support collaborative work between people that are physically in relative proximity. Since users are sharing a display in the same room, they are free to interact face to face, which facilitates collaboration. Ideally, a SDG system should provide multiple *independent* input streams so that several users can control a display at the same time without waiting for their turn to use a the single input device. Even though the SDG software *may* still enforce taking turns, there are several scenarios and applications where allowing at least a degree of parallel interaction can be very beneficial.

As not all collaborators can stand in the same place in a SDG system, *remote* interaction devices become almost a necessity. Consider a planning session for rescue efforts with ten or more participants scattered around a room with an SDG system. If there is only one interaction device that is wired to the computer, the device can only be used in a limited range. If participants in the back of the room want to contribute, they have to come to the front to use the interaction device. Even worse, since there is only one input device available, people are forced to go through the process of acquiring the device before they can contribute their ideas. A system with multiple remote devices does not suffer from this drawback.

Another scenario is a presentation, where interested members of the audience are given the ability to control the slides in the question and answer part of the talk. Again, remote input devices facilitate participation by all people in the audience. Yet another illustration is a brainstorming session in any of a large number of design applications. Here designers can quickly and effortlessly try out their ideas if multiple remote interaction devices are available.

Last, but not least, games become much more compelling when one display is shared between multiple players, while still allowing independent input from each player. Inkpen *et al.* [10] beautifully illustrated the benefits of multi-user SDG in an experiment with children. A pair of children played a paper-based puzzle game on a computer with one display and one mouse, and then on a computer with one display and two mice. The result shows that children are much more engaged and active when two mice are available compared to the single mouse condition.

There are other cases when it can be advantageous to share a display between multiple users, each having an input device. Common to all above applications is that extremely precise pointing and selection are usually not required as users are sitting or standing at varying distances from the display. Consider that a user in the back of the room can see only a certain level of detail due to fundamental limitations of the human visual system. If we assume that a remote pointing device is used, similar limitations in pointing accuracy apply due to hand jitter (see e.g. [1]). Consequently, a SDG application needs larger icons, buttons, and labels to be usable for *all* participants. Note, however, that a user right in front of the display may obtain high accuracy by working directly on or close to the display surface. For this reason we believe that fine manipulation tasks are counterproductive in SDG settings, unless fine positioning is *the* focus of the meeting.

We argue that if precise alignment, etc. is required, the SDG application should provide helpful manipulation techniques such as automatically snapping objects together for ease of alignment or providing 'intelligent' objects that

place themselves according to context in a 3D environment [13].

Laser pointers have been favored in meeting rooms as participants can directly point to objects on the display from any distance. We build on this and also address one of the main limitations of most current SDG hardware, namely that only *one* user can interact with the system at any given time. A camera, pointed at the screen, is used to detect the position of the laser spot. This position is then communicated to the system for further processing.

This paper discusses issues related to using a laser pointer as an input device. Then a usability test is presented to address if the device is suitable for SDG applications. Finally, we propose an efficient method to detect and identify *multiple* laser pointers.

2 Previous Work

Kirstein and Muller [2] presented a system that uses a laser pointer as a pointing device. Their system acquires video frames at 20 fps. They report that they are able to detect the laser spot in only 50% of the frames.

A simple performance test of the laser pointer as an input device was presented in the Pebbles project [1]. Since the button press on a laser pointer is not communicated to the computer in their system, they ask the user to turn the laser pointer on or off to select objects. However, the results of the user test clearly show that users cannot reliably turn on or off the laser pointer at the position they intend. To combat this, the authors suggest using the first and last dwelling positions as selection events. Still the problem is that detecting precise dwell positions takes about 2.5 seconds. The paper concludes that the laser pointer is inappropriate for selecting precise positions on a screen. Consequently, the authors claim that the laser pointer is not suited for selecting targets such as buttons or menus. Instead they suggest that the laser pointer be used only to select a region of an object or a menu system and to copy the selected object to a hand-held device for further manipulation.

Olsen [3] proposes an inexpensive interaction technique by introducing a set of window events for the laser pointer such as laser-on/off, and laser-move/dwell. To make the user aware of errors, and noise of the laser pointer, it displays a cursor corresponding to the detected position of the laser spot on the screen. Winograd and Guimbretiere [9] propose a new kind of interaction techniques for large displays, which are based on "gesture and sweep" paradigm instead of the usual "point and click".

Chen and Davis [4] describe a system that can provide multiple laser pointer inputs with multiple cameras. The system identifies the strokes created by the laser pointers on the screen in the following way. A Kalman filter is used to smooth the readings, which in turn are used to predict where the laser pointer will appear in the future. In each frame all points that potentially correspond to the current position of each stroke are collected. Then, the system collects all points that are within a certain range of the prediction and chooses the closest one as the candidate. As no

physical identification to distinguish strokes exists, wrong points can be selected as a part of the stroke and consequently the system cannot keep track of individual users.

Very few of the mentioned papers discuss the practical tradeoffs concerned with the detection of laser spots. None of the previously mentioned approaches can reliably support multiple users interacting with the system.

3 The laser pointer as a pointing device

In this section, we discuss several issues related to laser pointers. We mention laser spot detection, mechanisms to indicate selection, and a computer controlled laser pointer.

Laser spots appear whenever a laser pointer is held towards a screen. Similar to previous work the basic system utilizes a camera connected to a computer pointed at the screen to detect these laser spots. Note that depending on the setup, the camera can either be in front of the reflective screen or in the case of a transparent display, on either side of the screen. We recommend a back-projected screen where the camera is situated behind the screen, because it allows users to move more freely. Due to the space requirements of back-projected screens, one or more mirrors are often necessary to fold the optical path. We recommend using the same mirror to fold the optical path of the camera, too. Note that in this case, care has to be taken to avoid the direct reflection of the projection system, which usually saturates the image.

As in previous work, we perform a threshold operation to identify bright pixels for every frame acquired by the camera, and obtain the center of the laser spot from the weighed average of the bright pixels, which results in sub-pixel accuracy. This is also based on the observation that laser spots usually form the brightest pixels in the image. One reason for this is that many cameras are most sensitive in the red region of the spectrum. Consequently, a bright red laser spot often saturates the pixels in the video image.

One important advantage of using laser pointers is that there is no need to display a cursor, as the laser spot itself provides visual feedback of the selection.

3.1 Issues with laser spot detection

However, there are some issues that make it hard to perform reliable detection of laser pointer spots with this simple method that relies on contrast.

First, the camera must be focused on the screen. This can be difficult for large screens, especially if the camera is too close to the screen or mounted at a significant angle relative to the screen. This will cause a blurry and consequently dimmer image of the laser pointer spot in some regions of the screen. This also happens in the corners of a screen, if the screen is large relative to the distance of the camera as the brightness of the image diminishes with relative angle to the portion of the screen. The only solution here is to position the camera far enough from the screen and to make sure that the whole image is in reasonable focus.

The second issue is that the projected image and the image in the camera need to be registered. In other words, we

need to know for each pixel in the camera image, what the corresponding point on the display is. For a camera with good optics, a perspective mapping is enough. With cheaper optics, the barrel distortion of the image may also need to be compensated.

On the topic of image quality, the camera image as a whole might be too bright. One way to reduce the brightness is to reduce the aperture of the camera, but this changes the image geometry. Our preferred alternative is to reduce the shutter/exposure time of the image. With a relatively long exposure time the image of a fast moving laser spot is in general a blurred and dimmed trail. Consequently, we recommend the shortest shutter time possible.

A related topic is low contrast images due to a high level of incident illumination. The best solution is to make sure that there is not much light directly falling onto the screen. This can often be accomplished with appropriate baffles or similar devices. If there is a lot of ambient (undirected) illumination the only solution is to reduce the light level, which may necessitate using directed lighting for work surfaces. Another solution would be brighter laser pointers, but we hesitate to recommend this due to eye safety concerns. Last of all, we could use a more sensitive camera, but such devices are usually also much more expensive.

Users standing at an oblique angle to the screen also cause problems because the laser spot becomes progressively larger and dimmer. We observe this only to be a problem at angles greater than 85 degrees from the normal of the screen surface.

Another issue with camera based systems is latency. Latencies occur due to the time for the transmission of the image from the camera to the computer and the time for the computation to detect the laser point. Using grayscale images instead of color images can reduce the transmission time. Another alternative is to reduce image resolution, but this also reduces the effective resolution of the input device. Finally, we could use compression, but here compression artifacts are a concern.

Due to hand jitter, it is practically impossible to use a remote pointing device to indicate a small target for an extended period. Therefore, only lightweight remote pointing devices are recommended in practice to avoid fatigue issues. Another choice is to use supporting surface, such as e.g. the work surface in the case of a mouse. But in many SDG systems such a surface may not be available everywhere in a meeting environment.

To reduce the effects of jitter on the laser pointer position we use a Kalman filter to smooth out the readings similar to the approach described in [6]. With the described techniques we are able to reliably track a laser pointer with an NTSC camera at 60 Hz. If configured correctly a side benefit of the Kalman filter is that it smoothes out jitter due to the even and odd fields of the NTSC signal.

3.2 Selecting objects with laser pointers

For a person with a laser pointer there are multiple possibilities to indicate selection of an object. As discussed in the paper about the Pebbles project [1] one alternative is to use the on/off of the laser pointer as an indication of selection. In other words, whenever a laser spot appears this indicates selection. However, the authors report that the user cannot reliably make the laser spot appear where they intend with this method.

Another option explored in the Pebbles project is to use a threshold on the dwell time (e.g. 2.5 seconds). As reported there, this approach is problematic due to hand jitter and has the additional disadvantage that it is time-consuming to select an object. The authors of [1] abandoned this approach and used a handheld computer with an integrated laser pointer for their future work.

The last alternative is to directly use the button on the laser pointer without switching the laser pointer on or off. In our experience, most first-time users for the laser pointer believe that this is the natural way of doing things. To achieve this, it is sufficient to modify a laser pointer so that the button and laser power can be controlled separately. For all hand-held laser pointer devices an important question is if the button event should be transmitted on the up or down motion. Most desktop-based devices use the up-event to designate a click. For laser pointers this does not work, as each button press will cause a small 'dip down' in the position of the laser spot due to the force exerted onto the laser pointer case as the user holds the laser pointer in the air. Consequently, it is better to record the position for the button down event and to use this for selection purposes.

One way to separate power and button control is to mount a little power switch on the laser pointer and use the button to interrupt the laser power while it is pressed. However, with this method the laser may stay on for extended periods, which significantly shortens the lifetime of the laser diode. More seriously, a laser pointer that is always on may cause safety problems. A better realization is to put the laser pointer under computer control.

3.3 A computer controlled laser pointer

For a computer controlled laser pointer, there are two important issues: sensing button presses and switching the laser on or off from the computer.

There are several ways to connect a laser pointer to a computer. The simplest and most reliable approach is to use a cable. However, from our experience with Virtual Reality (VR) hardware, we know that cables are a major nuisance. People are limited in their range of operation or trip over cables, the cable often gets in the way, or the weight of the cable itself causes problems.

Wireless alternatives are to use infrared, ultrasonic or radio transmission. If the same mechanism is used to transmit button presses and the power signal, there is the potential for cross talk, which complicates things. If different mechanisms are used this issue is eliminated.

The alternative we recommend is to use infrared for the 'power-on' signal and to use ultrasound for transmitting the button press. As for the button press, the simplest alternative is to use an ultrasonic emitter, which is activated as

long as the button is pressed. For the 'power-on' signal infrared LED's should be mounted near the screen (preferably at the bottom or the top) to transmit the 'power-on' signal to the laser pointer. A little infrared sensor mounted on the front of the laser pointer then closes the circuit to the batteries as long as the signal is received. This is an application of the technology used to control so-called shutter glasses for VR. The benefit of using infrared is that the laser pointer may turn itself automatically off if the pointer is not pointed in the direction of the screen

4   User test of the laser pointer

In this paper, we are considering a laser pointer as an input device for a large-scale display. To see if a laser pointer can be used in practice, we have to assess its performance from the user's perspective. The best comparison point is the mouse, yet we have to address the issue of different screen sizes. Consequently, we choose to perform a baseline comparison that compares the laser pointer with a mouse, where the mouse is used on both a small and a large display.

4.1   The task and measurement

We have chosen to follow the methodology proposed by the ISO 9241 standard [5], which is based also on MacKenzie *et al.* [7] work. In this study we perform a "pointing and selecting" task (see the ISO standard) with the laser pointer and the mouse. A sketch of the display for the task is shown in Fig. 1. The task is designed to exercise many different directions of movement. The numbers marked on the circles define the order of the selections the participant has to perform. Selecting the target marked '0' starts a block of trials. Each trial starts after the selection of the current target, and ends at the selection of the next target. Consequently, there are 15 trials in each block. The movement time is measured on a per trial basis.
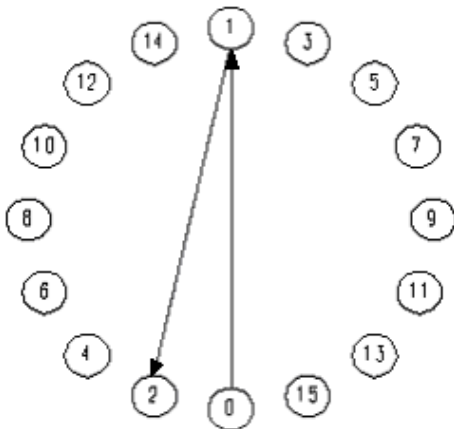

Fig. 1. Sketch of input device evaluation task.

As explained above we use the 'button-down' event for the laser pointer. For the mouse, we still use the 'button-up' event as in standard desktop applications. A cursor of the laser pointer on the screen is the laser spot itself, and no graphical cursor is displayed. An error is recorded when the participant misses a target.

According to the ISO standard, different indices of difficulty should be included in the test. The index of difficulty (ID) is a measure of the precision required from the user in a task. It is defined in bits as:

$$ID = \log_2(1 + \frac{D}{W}),$$

where D is the distance between targets, and W is the width of the target. In our user study, we asked participants to perform tasks at three different ID's.

Throughput as defined by ISO standard, is then calculated from the measured movement time (MT) in seconds and the recorded user events:

$$Throughput = \frac{ID_e}{MT},$$

where ID_e is the effective index of difficulty. Throughput is measured in bits per second (bps). The effective index of difficulty is defined as:

$$ID_e = \log_2(1 + \frac{D}{W_e}) \text{ with } W_e = 4.133 \times SD_x,$$

where W_e is the effective width of target and SD_x is the standard deviation in the selection coordinates measured along the axis of approach to the target. The use of the effective width and not the real width accounts for the fact that users sometimes miss the target by a small amount.

4.2   Installation and procedure of the test

We mounted a class I (output < 1mW) laser pointer into a small case and installed two buttons on top of this case. Only the foremost button was used in this test. See Fig. 2 for a picture, the head of the laser pointer is visible on the left side.
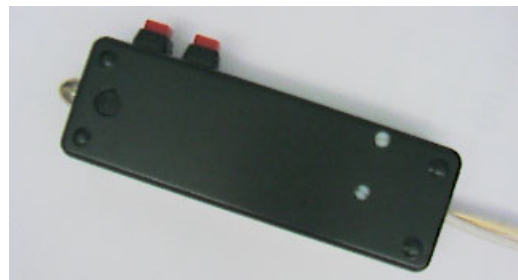

Fig. 2. The laser pointer used in the baseline comparison.

The computer controls the power of the laser pointer via a small circuit consisting mainly of an optical coupler mounted into a control box. For simplicity and speed, we choose to use the parallel port of the computer for I/O. The laser pointer itself is connected to the control box via lightweight telephone wire.

Depending on the condition the task either displayed on a back-projected 6' by 4' screen (using a 700 lumens projector) or a 19" desktop screen. For this study, we used a NTSC camera. The camera was installed behind the 6' by 4' screen. The camera captures the grayscale image of the screen at 60 Hz. A PC running Windows NT acquires the

images. In all conditions participants were seated and we asked them not to move the chair during the test. To ensure that the apparent size and viewing angle of the screen were consistent during the test, we positioned the 19" screen so that the screen corners lined up with the corners of the 6'x4' screen from the participant's point of view. The test was performed under normal office lighting conditions.

Prior to the test, we gave a practice session to minimize learning effects and to familiarize the participant with the use of the laser pointer. In this practice period, each participant was required to repeat each task with each device at least three times and we gave them the option of further practice. For the test, participants were instructed to hit the targets as fast as possible. Between the trial blocks, we allowed participants to rest whenever they expressed a need to do so. After the test, participants were asked to answer questions assessing the comfort and ease of use, as specified in ISO9241-9 [5]. In our user study, participants took an average of 50 minutes for the total test including practice, test, and questionnaire fill-up time.

4.3 Design

The experiment is a (3 device conditions) x (15 trials) x (10 blocks) x (3 index of difficulty) x (12 participants) factorial design. Twelve participants (six male, six female, age ranging from 26 to 45, average 26.3) were recruited from graduate and undergraduate students at the local university. Each of participants performed a total of 1,350 trials according to the test configuration specified in table 1. We counterbalanced the device conditions to combat learning effects.

Table 1. Configuration of the experiment

Factor	Values
Device	Mouse with 19" screen Mouse with 6'x4' screen Laser pointer with 6'x4' screen
Index of difficulty	2.58 (300,60)* 3.17 (400,50) 3.75 (500,40)
Block	10 blocks of trials
Trial	15 trials in each block

* The pair of numbers in brackets indicates distance between targets and diameter of a target in pixel, respectively.

The above design was chosen based on the results of a small pilot study and the fact that we were primarily interested in a baseline comparison. Based on our observations in the pilot study the three indices of difficulty mentioned above correspond to an easy, medium and reasonably difficult task. Furthermore, we designed the experiment to keep the time per participant to less than one hour to limit fatigue effects. This imposed a maximum on the number of trials and on the different factors we could investigate in the test.

4.4 Results

Learning effect

First, we analyzed the results according to learning effects. A quick glance at the graph of average trial movement times grouped by block (see Fig. 3) shows that there are still learning effects in the first few trials. A more detailed analysis with a repeated-measures ANOVA by blocks shows that the blocks from 5 to 10 form one group ($F_{2,9}$=3.82, p<0.01).

We also observe that the movement time of the laser pointer is increasing towards the end of the user test. As far as we can tell this is because participants felt tired from holding the laser pointer in the air throughout the test. However, the difference between the blocks of laser pointer trials is not significant ($F_{9,350}$=0.97, p>0.05).
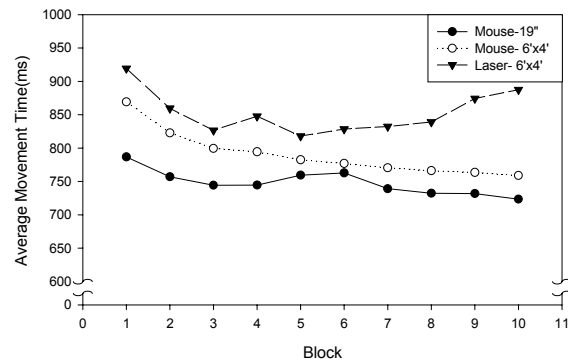

Fig. 3. Average movement time of device by block.

To exclude potential issues with learning effects, we choose to ignore blocks 1 to 4 and all results reported from now on are for block 5-10 only.

Movement time

The movement time by device is shown in Fig. 4 as a box plot. Boxes indicate the 25^{th} to the 75^{th} percentile range. The bars above and below the box indicate the 90^{th} and 10^{th} percentile. The line in the middle marks the median.
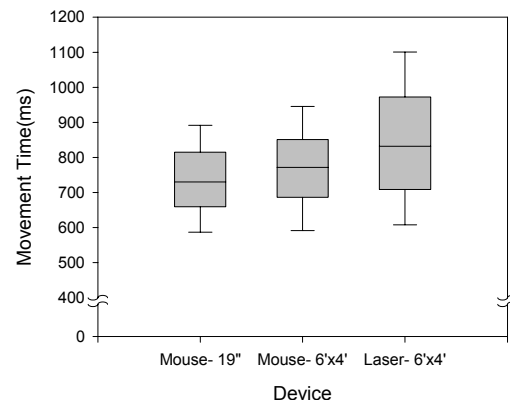

Fig. 4. Movement time by device.

The average movement time per trial for the laser pointer is 846.78 ms, the mouse with 6'x4' screen is 769.72 ms, and the mouse with 19" screen is 741.06 ms. The movement

time of the laser pointer is significantly different from both mouse conditions at $F_{2,645}=28.46$, $p<0.001$. The difference between the two mouse conditions is statistically not significant. Consequently, we can conclude that the size of the screen is not important, but that there is a difference between the devices.

Analysis of the movement time according to the index of difficulty yields the unsurprising result that more difficult tasks take longer. The movement time of the laser pointer is significantly different from the rest of the device conditions with the first index of difficulty at $F_{2,213}=7.15$, $p<0.001$.

Error Rate

The average error rate for the laser pointer is 32.16%, the mouse with 6'x 4' screen is 20.46%, and the mouse with 19" screen is 20.31%. The difference of error rate between the mouse with 19" screen 6'x4' screen is not significant. The difference between the laser pointer and the mice conditions is significant at $F_{2,645}=19.71$, $p<0.001$.

A more detailed analysis of the error rate reveals that the error rate for the laser pointer increases stronger with increasing index of difficulty compared to the mouse. In other words the smaller the target the harder it is to select it with the laser pointer. In the most difficult condition, the error rate is approximately 40%.

Throughput

The average throughput of the laser pointer is 3.04 bps, the mouse with 6'x4' screen is 3.98 bps, and the mouse with 19" screen is 4.09 bps. The throughput of the laser pointer is significantly different from the mouse with 19" screen and with 6'x4' screen at $F_{2,645}=67.50$, $p<0.001$. In other words, the average throughput of the laser pointer is about 75% of the mouse. Fig. 5 shows the distribution of throughput of each device.
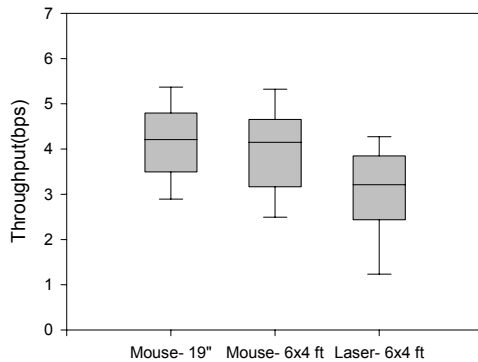

Fig. 5. Throughput by device.

Qualitative evaluation

After the test, participants answered the thirteen questions listed in the ISO standard. Those questions are designed to evaluate the comfort and difficulty of using device.

As shown in table 2, the averages of general comfort and overall difficulty of laser pointer are slightly lower than the mouse, but the difference is not statistically significant. The results for all other questions are not significant, as well.

Table 2. Evaluation of general comfort and overall difficulty for different conditions.

	Mouse-19"	Mouse-6'x4'	Laser-6'x4'
General comfort	3.92	3.83	3.42
Overall difficulty	4.67	4.58	4.17

Furthermore, we asked participants to freely comment on the devices. The most frequent comment was that the button on the box was hard to press and that this caused them difficulties in selecting targets. Another comment was that the contrast of the laser spot against the display was not high enough. In fact, we observed that participants sometimes completely lost track of the laser spot in fast movements. Both factors cause fatigue and discomfort, which may explain the rising times towards the end of the study.

On the positive side, participants commented that they would really like to use this device for giving presentations, as it is free of the desk surface, yet provides optical feedback of the current position via the laser spot. Asked about their first choice for a remote pointing device, participants clearly expressed a strong preference for the laser pointer.

4.5   Discussion

The most striking result of our baseline comparison is that the throughput of the laser pointer is significantly below that of the mouse, at approximately 75%. Since the laser pointer is held in the hand, whereas the mouse rests on the desktop, this result is not very surprising. In fact, we consider 75% compared to the mouse a relatively good result, considering the absence of a stabilizing surface. Since laser pointers have strong advantages over mice in SDG environment such as mobility and the ability to interact with the display from everywhere in a room (even *standing*), we see this result as an indication that laser pointers are very good input devices for SDG systems.

The size of the screen slightly affects mouse performance, but no statistical difference exists. Thus, we can say that the size of the display does not matter for mouse performance.

A side benefit of the laser pointer is that the bright spot caused by the laser pointer serves directly as the cursor. While it is physically smaller than a mouse cursor, it is still very noticeable. With our first prototype the laser spot was sometimes too dim during fast movements, which caused problems. Users effectively had to perform a visual search for the laser spot after a movement, which decreased performance.

If we compare our results with MacKenzie's experiments [7] we can see that the error rate is much higher. We attribute this to the fact that we instructed our participants only to complete the task as fast as possible – as opposed to achieving maximum speed while still hitting all the targets. This naturally led to a much higher error rate in our study.

The biggest technical problem in our baseline comparison was the button, which required a relatively large amount of force to be pressed. From our observations we believe that the button press itself took longer compared to a mouse button press, but we did not attempt to measure this. Sometimes participants involuntarily performed repeated button presses on the same target, which increased the error rate for the laser pointer by a significant amount

4.6 Modifications to the original design

Based on the user comments and our observations from the user study we modified our design for the laser pointers. The first modification was to use a common laser pointer case as enclosure, which also makes the device lighter. An added benefit is that the button in such laser pointers is connected to a microswitch similar to the ones used in mice. Furthermore, we replaced the laser diode with a model that conforms to class IIIa (output < 5mW), which increases the brightness of the laser pointer. We also took this modification as an opportunity to replicate the circuitry for one pointer to support multiple laser pointers, see the next section. Fig. 6 shows our current prototype with three laser pointers.

We also replaced the NTSC camera with a Firewire camera (Pyro Webcam) that can acquire images with 640x480 resolution at 30 Hz. While the lower frame rate somewhat degrades performance, the image quality is significantly better. Another major advantage of the new camera is that the shutter time can be controlled electronically, which greatly simplifies setup and adaptation to different lighting conditions.

In our experience, the modifications improved both the reliability of laser spot detection and the formfactor for the laser pointer. In addition, the new button allows for practically error free button press detection, which makes it much more responsive from the user's point of view. Although we have not performed a formal user study with this new version, we are confident that results will be better than for our first implementation.

Fig. 6. Three computer controlled laser pointers and the box containing the electronics to control them.

5 Distinguishing multiple laser pointers

The basic idea in tracking multiple laser spots is to use prediction and to assign the laser spot that is closest to each predicted position to the corresponding laser pointer. This has been tried before [4]. However, as mentioned there, this scheme fails when multiple laser spots come in close proximity. Also, it is hard to reliably determine which user is manipulating which laser when laser pointers are turned on or off or leave the screen.

5.1 Identifying different laser pointers

To improve on this we discuss multiple ways of distinguishing between different laser pointers. We consider different colored lasers and blinking the lasers.

While different colored laser pointers make it easy for the users to distinguish them, there are a couple of problems. The biggest problem in using color as an identifier is that laser pointers outside the 650-680 nm range are relatively expensive. Moreover, as bright laser spots often saturate the CCD, it is relatively hard to reliably detect the color of a laser pointer from the image. An alternative is to use multiple cameras with different band-pass filters. This is more expensive and the use of filters makes the images much darker, which amplifies laser spot detection problems.

A better alternative is to use different patterns to blink the laser pointers. This necessitates that the blinking cycles are synchronized with the picture acquisition by the camera. This is easy to realize by modifying the state of the laser pointer directly *after* a frame was recorded. Please note that blinking standard laser diodes may shorten the lifetime of these devices significantly. We recommend using laser diodes that support modulated input signals.

The simplest approach is to power each laser pointer in turn in a cyclic pattern. The biggest disadvantage is that each laser pointer is only powered for a fraction of the time, which dims the laser spot for the user. Another approach is to blink the lasers in a binary or Gray code pattern, which maximizes the number of laser pointers that can be used. Note that the pattern 00…0 cannot be used. The downside to using binary patterns is that some laser pointers will be off most of the time (consider the pattern 0…01) and some will be on most of the time (e.g. 1…10). That means that each laser pointer has different number of measurements, which leads to different tracking behavior. The last approach is to turn each laser pointer *off* for only one cycle of a repeated pattern. Compared to the first alternative, every laser pointer will be on most of the time, and we have the benefit of equal number of measurements per time for each laser pointer. See Table 3 for a comparison of the different alternatives.

Table 3. Comparison between blink patterns of length 3.

Method	Blink-on	Binary	Blink-off
1st laser pointer	001	001	110
2nd laser pointer	010	010	101
3rd laser pointer	100	011	011
4th laser pointer	N/A	100	N/A
…	…	…	…
7th laser pointer	N/A	111	N/A
Percentage on	33%	Varies	66%

Regardless of the pattern used, the cycle length limits the frequency at which laser pointers can be identified reliably. Consequently, for a 60 Hz camera and cycle length three, we can get reliable identification of laser pointers at 20 Hz. Depending on the details of the algorithm, intermediate classifications may be correct or incorrect.

5.2   Proof of Concept

In our implementation, we use the blink-off approach to identify the laser pointers. In the following description we use a pattern length of three, but the algorithms generalize easily to more laser pointers.

For each frame, we predict the position where a laser pointer should currently appear with the help of a Kalman filter. As in [4] we initially assign to each laser pointer the spot that is closest to the prediction for this pointer. Simultaneously, the program records if the laser pointer has been assigned the detected spot.

Depending on the circumstances, at the end of this per-frame process there may be laser spots that cannot be associated with a laser pointer and laser pointers that cannot find a laser spot close to the predicted position. Unassociated laser spots are classified as starting points for a new laser pointer. If there is a laser pointer where no close enough laser spot exists, we assume that it has been turned off. Finally, each laser pointer updates its measurement if a current laser spot was found and record their pattern as on, otherwise they record their pattern as off.

After a number of frames corresponding to the length of the pattern have been recorded (i.e. 3 in our discussion), the method compares the recorded and predefined patterns. If the patterns do not match, we can say that one or more measurements were erroneous. In our current implementation we simply look for another laser pointer with wrong measurements (which must exist) and swap the last readings for these two pointers.
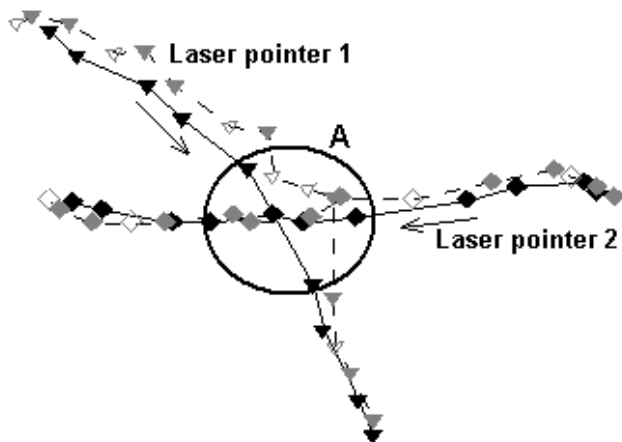

Fig.7. Recorded paths of two laser pointers.

Fig. 7 shows the path for two laser pointers recorded at 30 Hz. Solid lines with black points are measured points and dashed lines with gray points illustrate the prediction. Empty shapes stand for predictions without a correspond-

ing measurement. The first laser pointer is designated by triangles and has the pattern 011, whereas the second is shown with diamonds and has the pattern 110. One can observe misclassifications in area A, but can also observe that the misclassification is quickly corrected. Note also how the prediction error increases after empty shapes due to the higher inaccuracy in prediction.

With this scheme, the classification of the acquired points depends only on the predictions for each frame and their recorded state pattern. Problems arise if a laser spot is not detected reliably e.g. due to very fast movements or when two laser spots coincide. Consequently, we cannot *guarantee* that each laser pointer is detected correctly in every frame. However, we can reliably identify laser pointers correctly at every third frame!

For a wireless implementation of our system, we need to separate signals for different laser pointers. For the buttons, this can be achieved by using transmitters with different frequencies. For the infrared transmission of the power-on signal, we can use the circuitry used in TV remotes to transmit different signals and to distinguish among them at the laser pointer side.

5.3   Extension for tiled displays

Large tiled displays are needed to display, annotate, and edit large amounts of information in a collaborative environment. For example, the i-Land system [11] presents the DynaWall, a wall sized touch sensitive display, as one of the components of the roomware prototype.

Without expensive equipment such as a touch sensitive display, laser pointers are a simple and cheap input alternative that supports multiple/parallel inputs for large-scale displays. Our current implementation uses only one display, but there's no reason why the system cannot be extended to multiple displays. In fact, one way to achieve this has been described in [4].

5.4   Discussion

We created two demonstration applications. One is a simple puzzle game, another is a 3D interior design application. When running the system we observe that users indeed use the opportunity to work in parallel if they are not focusing on a single object.

Surprisingly, although there is no visual identification of each laser pointer, users do not seem to experience problems identifying their laser pointer. A quick 'wiggle' of the laser pointer will usually give enough visual feedback to identify it.

6   CONCLUSIONS

In this paper, we presented a comprehensive discussion of laser pointers as input devices. Several issues regarding reliable detection of laser spots and selection techniques with laser pointers were mentioned. A new form of computer controlled laser pointer device was introduced. This implementation may be more expensive than using multiple

mice. However, the mobility and accessibility from remote distances are strong advantages over using mice.

Furthermore, we presented an initial baseline comparison of laser pointers with a mouse according to the ISO standard. Our results indicate that laser pointer performance is approximately 75% compared to a mouse. We interpret this as a good indication that laser pointers can indeed be used in applications where users may be standing or sitting in chairs. The results of the presented study also show an increased error rate for the laser pointer, which may be due to limitations in the first prototype. Based on this and other issues raised by the user study we discussed modifications to the prototype, which makes the device easier to use.

The refined system also supports multiple laser pointers as interaction devices on one or more displays. Consequently, the system affords seamless and parallel collaboration among several users. Several alternatives for distinguishing between different laser pointers were discussed, and one implementation was demonstrated with SDG applications.

For a wireless version of the system, we do not expect any major functional differences to the results presented here. We are currently working on such a system and will report about it in future work.

In addition, we plan to perform a more controlled user study of the performance of the modified laser pointer system. This study will explore a wider range of indices of difficulty and will also investigate the fatigue factor. Also we plan to inform participants better about the need to hit targets accurately, which should improve the error rates.

Furthermore, we are currently investigating way to make computer controlled laser pointers safer, as a class IIIa laser pointers can *potentially* cause at least temporary eye damage [12]. Finally, we intend to study of the multi-user aspects of our system with SDG applications.

REFERENCES

[1] Myers, B. A., Bhatnagar, R., Nichols, J., Peck, C. H., Kong, D., Miller, R., and Long, A. C., Interacting at a distance: measuring the performance of laser pointers and other devices, *Proceedings CHI'02*, to appear.

[2] Kirstein, C. and Muller, H., Interaction with a projection screen using a Camera-tracked laser pointer, *Multimedia Modeling '98 Proceedings*, pp. 191-192.

[3] Olsen, D. R. Jr. and Neilsen, T., Laser pointer interaction, *CHI'2001*, pp.17-22.

[4] Chen, X. and Davis, J., LumiPoint: Multi-user laser-based interaction on large tiled displays, *graphics.Stanford.edu/papers/multiuser*.

[5] ISO, ISO/TC 159/SC4/WG3 N147: Ergonomic requirements for office work with visual display terminals (VDTs) - Part 9 - Requirements for non-keyboard input devices, *International Organisation for Standardisation*, May 25, 1998.

[6] Azuma, R. T., Predictive Tracking for Augmented Reality, *Ph.D. Dissertation*, University of North Carolina at Chapel Hill, NC, USA.

[7] MacKenzie, I. S., Jusoh, S., An evaluation of two input devices for remote pointing. *EHCI 2001*, Heidelberg, Germany: Springer-Verlag.

[8] Stewart, J., Bederson, B. and Druin, A., Single display groupware: a model for co-present collaboration, *CHI'99 Conference Proceedings*, pp. 286-293.

[9] Winograd, T. and Guimbretiere, F., Visual Instruments for an Interactive Mural, *CHI'99 Abstracts*, 234-235.

[10] Inkpen, K M., Ho-Ching, W., Kuederle, O., Scott, S., Shemaker, G. B.D., "This is fun! We're all best friends and we're all playing", Supporting Children's Synchronous Collaboration, *Proceedings of Computer Supported Collaborative Learning (CSCL) '99*. Dec. 1999. Stanford, CA.

[11] Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R., i-LAND: An interactive Landscape for Creativity and Innovation, *CHI '99*, pp. 120-127.

[12] http://www.sciam.com/askexpert/medicine/medicine38/medicine38.html

[13] Smith, G., Salzman, T., Stuerzlinger, W., 3D Scene Manipulation with 2D Devices and Constraints, *Graphics Interface 2001*, pp. 135-142, June 2001.