A Comparison of Ray Pointing Techniques for Very Large Displays

Ricardo Jota[1], Miguel A. Nacenta[2], Joaquim A. Jorge[1], Sheelagh Carpendale[2], Saul Greenberg[2]

[1]INESC-ID
Technical University of Lisbon, Lisboa, Portugal

[2]Department of Computer Science
University of Calgary, Alberta, Canada

jota.costa@inesc-id.pt, jaj@inesc-id.pt,{miguel.nacenta,sheelagh.carpendale,saul.greenberg}@ucalgary.ca

ABSTRACT
Ray pointing techniques such as laser pointing have long been proposed as a natural way to interact with large and distant displays. However we still do not understand the differences between ray pointing alternatives and how they are affected by the large size of modern displays. We present a study where four different variants of ray pointing are tested for horizontal targeting, vertical targeting and tracing tasks in a room-sized display that covers a large part of the user's field of view. Our goal was to better understand two factors: control type and parallax under this scenario. The results show that techniques based on rotational control perform better for targeting tasks and techniques with low parallax are best for tracing tasks. This implies that ray pointing techniques must be carefully selected depending on the kind of tasks supported by the system. We also present evidence on how a Fitts's law analysis based on angles can explain the differences in completion time of tasks better than the standard analysis based on linear width and distance.

Author Keywords
Large displays, ray pointing, distant pointing, image-plane, targeting, tracing, index of difficulty, ISO 9241, parallax.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
People often interact with a large digital display by *distant* pointing, or directly pointing at the display with their finger, laser pointer, or other input device. *Ray pointing* is a class of techniques that uses *ray casting* (the intersection of a ray with a surface) to determine where a person is pointing to, i.e., the precise cursor position on the distant display. Ray pointing is advocated as a natural interaction technique with these displays [2,18] as: it does not require any physical surface to operate on (as opposed to mouse-based pointing) [26,25]; it gives access to distant locations without requiring displacement (as opposed to direct-input) [19,29]; it is easily understood by people as it builds upon everyday pointing gestures [29]; and it allows multiple users to interact on the same display without their bodies physically getting in the way [16]. Thus it is no surprise that ray pointing is increasingly used in both commercial and research systems, especially for large horizontal and vertical displays [28,6,19,16]. Even game consoles are shifting towards ray pointing for interaction (e.g., Nintendo Wii).

Digital displays are becoming increasingly large, with many sites now reporting one or more wall-sized displays. The issue is that there are many different ray pointing techniques, and we do not understand the differences between their performance characteristics, especially in terms of how they are affected by these very large displays. While there is prior work on ray pointing, it is limited. Some concentrate on specific type of ray pointing (laser pointing) [25,13], neglecting other alternatives. Others limit pointing activity almost exclusively to targeting tasks [13,11]. Moreover, such targeting is often studied only within small display areas, i.e., where pointing ranges cover only a small part of the visual angle of users [15,13].

In this paper we study four different variants of ray pointing techniques over a large display, where people perform both targeting and tracing tasks. Our variants cover a wide spectrum of ray pointing possibilities: *laser pointing*, *arrow pointing*, *image-plane pointing and fixed-origin pointing*. As will be explained, the first three are commonly used in real-life activities such as art and sports.

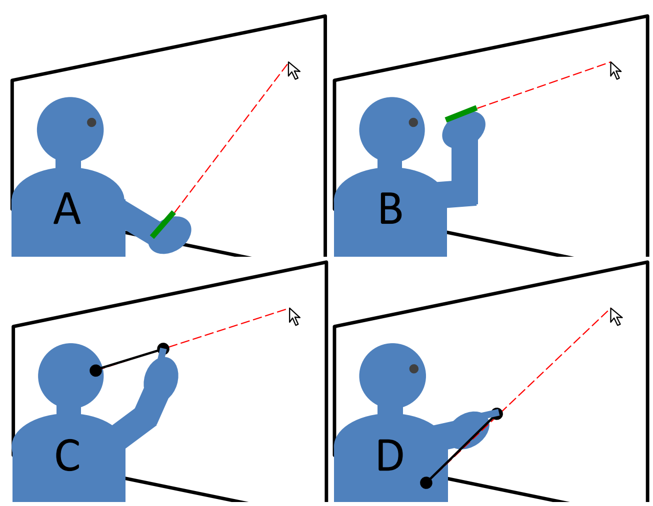

Figure 1. Four variants of ray pointing: A) laser, B) arrow, C) image-plane and D) fixed-origin.

Jota, R., Nacenta, M.A., Jorge, J.A., Carpendale, S., and Greenberg, S. A Comparison of Ray Pointing Techniques for Very Large Displays.

Report 2009-942-21, Department of Computer Science, University of Calgary, Calgary, AB. Canada, T2N 1N4

Specifically, we conducted three studies that tested the four techniques for horizontal targeting, vertical targeting, and tracing tasks on a room-sized display that covers a large part of the user's field of view (>90°). Our study differs from previous work in three fundamental ways:

- We compare several quite different ray pointing variants rather than investigate just a single method;
- We test tracing tasks as well as the usual targeting;
- We specifically consider issues associated with very large displays, e.g., when comparing displayed objects in front of the person vs. at the side, the person will see these objects at quite different distances and angles.

Our results show that techniques based on rotational control (laser pointing and arrow pointing) perform better for targeting tasks, while techniques with low parallax (arrow pointing and image-plane pointing) are best for tracing tasks. This implies that ray pointing techniques must be carefully selected depending on the kind of tasks supported by the large display.

We also found that targeting and tracing tasks are heavily affected by the location on the large display in which they take place. We present evidence on how a Fitts's law analysis based on *angles* vs. the standard *linear width and distance* approach explains differences in completion time of tasks better.

The remaining of the paper is organized as follows. We explain ray pointing fundamentals, and describe the techniques and review previous research in the area. We then describe the experiment and results. We conclude with a detailed discussion of the implications of the findings.

RAY POINTING FUNDAMENTALS

We now describe the particular ray pointing techniques we use and how we implemented them in our study.

We define *generic ray pointing* as any cursor-movement technique that determines the position of the cursor through the intersection of a ray with a distant object or surface (see Figure 2, left). For our purposes, the distant object or surface is a large display.

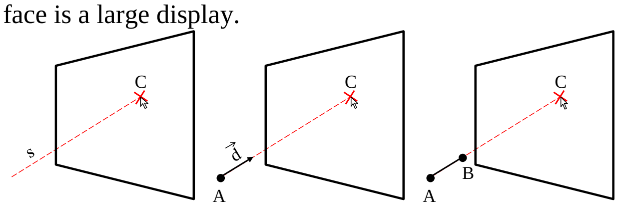

Figure 2. *Left)* Ray pointing specifies the location of the cursor through the intersection of a ray (s) with the display surface. *Center)* the ray (s) can be specified through a point (A) and a direction, or *Right)* through two points (A and B)

Regular Laser Pointing

The most common ray pointing variant is *laser pointing*. Here, the ray is specified directly by the position and direction of a physical device (Figure 1a). The device might or might not be an actual laser; in fact, the only requirement is that the computer system has a way of determining the intersection of the ray with the screen surface. For example, vision technology or special markers on the hand recognizes finger postures as a pointing device [28].

Laser pointing has been proposed and implemented for cursor control many times (e.g., [2,18,6,1,25,5,27]). It is often referred to as distant pointing, remote pointing or virtual pointing. In our study, we implement a laser pointer via an infrared-marked wand tracked in 6DOF.

Arrow Pointing

Arrow pointing is a variant of laser pointing, where we constraint the use of the laser pointer to be somewhat aligned with the user's eye (Figure 1.B). This mimics the real life way people aim when great precision is required (e.g., when using bow and arrow, or playing darts). Our implementation is identical to laser pointing, except now people are instructed to constrain their use of the wand by looking down its shaft at the screen, i.e., as if it were an arrow.

Image-Plane Pointing

An alternative ray pointing technique comes from the visual arts. Painters are often taught to place their thumb at arm's length between their eye and a painting to estimate the sizes and positions of painted objects. This technique has long been adopted in the field of virtual reality, where it is referred to as *image-plane manipulation*, *occlusion selection*, or the *crushing heads* technique [22,7,31,4,30, 12].

The mechanism of image-plane based pointing is simple: instead of determining the ray through the position and orientation of a pointing device, the ray is determined through two points in space: the user's eye location, and another point in space that the user can control (e.g., the position of the tip of the thumb, of a pen, or the point of a pointing device – Figure 1.C). The effect is that the user can see the cursor aligned with the thumb (or device) in her field of view, even if they are actually at different depths (Figure 3). To a certain extent, image-plane pointing is similar to direct-input techniques (e.g., direct-touch) in that feedback and input overlap in the visual space of the user.

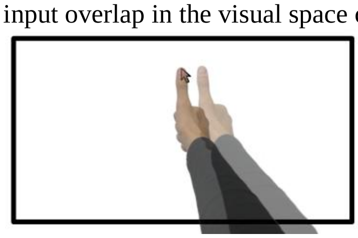

Figure 3. Image-plane pointing seen binocularly and focused on the distant display (the cursor is displayed on the screen).

Image-plane techniques require tracking (or approximating) the eye position, and are usually calibrated so that the dominant eye image aligns the finger or device with the cursor (however, binocular focusing on the distant surface still implies that two separate images of the finger or device are perceived by the user as in Figure 3). In our studies we approximate eye position – the first point of the ray – in real time by placing markers on a hat; a person calibrates the vector between hat and eye before interaction by specifying the position of their dominant eye with another marker. We use the tip of a wand to specify the second point of the ray.

Fixed-Origin Pointing

We can relax image-plane pointing by placing one of the two points of the ray onto any fixed location (instead of the eye). The user still controls the other point, or may even control the two points separately (e.g., one with each hand). The former alternative – which is the one we study – was explored somewhat by Shoemaker and colleagues in *shadow reaching* [23]. Shadow reaching allows the control of a large display through the shadow cast by a person on a large display illuminated from a fixed point behind the person. Because shadows are cast in straight lines, shadow reaching is geometrically equivalent to fixing point A on the location of the light source and using the pointing gesture of the person (usually the finger) as point B. Shoemaker et. al. also speculate using a virtual light source that would move with the user at a certain fixed distance. The latter alternative (allowing the user to control both points) is akin to pointing using the position and direction of an elastic band held with two hands.

We tested only *fixed-origin pointing*, where the origin point of the ray is fixed in space. The user controls the other point to specify the ray's direction. We use an origin point near the bellybutton of the user so that the required pointing device movements are somewhat similar to shadow reaching, where the light source is located close to the floor and behind a person.

RELATED STUDIES

Previous research in ray pointing falls mostly into two categories: laser pointing studies for interaction on distant displays, and virtual reality techniques for manipulation of objects in 3D environments. At the end of this section we also discuss enhancements to ray pointing techniques.

Laser Pointers for Large Displays

Thanks to the studies in this first category we now know a great deal about laser pointing. For example, MacKenzie and Jusoh [13] and Stuerzlinger and Oh [25] showed that laser pointing targeting performance is poor compared to the mouse (and around 1.4b/s or 3.0b/s respectively). Peck [21] parameterized the jitter of a laser pointer spot in terms of angle, and suggests that grip affects it. Myers and colleagues [14] studied the effect of different grips and postures, and found reduced jitter with a PDA-pointer held with two hands. Kopper and colleagues [11] proposed several models for pointing tasks on very large displays that take distance into account.

Most studies of pointing for large displays, with the exception of [9], only test laser pointing techniques. Our study compares a broader range of ray pointing techniques. At the same time, we pay special attention to the effects of very large displays in performance (similar to [11]).

Pointing in VR

The variety of pointing techniques studied in the Virtual Reality literature is broader, since image-plane techniques are easy to implement (the required tracking of head or eyes is already present). In general, studies comparing image-plane selection to ray casting (laser pointing) for manipulation of 3D objects in 3D spaces have found that the image plane method is generally faster [12, 3, 4, 31, 30]. This led Hill and Johnson [7] to propose an interaction manipulation framework based on image-plane techniques. However, most of the above-mentioned studies concern 3D tasks, which can be radically different to the 2D tasks that are our concern. It is not yet clear whether image-plane techniques will provide performance advantages for pointing to large 2D surfaces.

Laser Pointer Enhancements

Several enhancements have been proposed that modify or improve the operation of ray pointing and distant pointing interaction. For example, laser pointers are often filtered [6, 28], its CD gain altered [10], alternative modalities blended into the action [27], and snapping mechanisms added [27, 15]. We know that some of those mechanisms may improve pointing (e.g., fitering) but these also imply trade-offs (e.g., filtering implies delay [20], and *semantic snarfing* [15] makes it harder to operate with empty space). We chose not to alter the basic elements of pointing partly because these modifications can introduce a large number of parameters that can obscure fundamental effects we are after.

RAY POINTING PERFORMANCE FACTORS

There are many possible factors that might affect ray pointing performance (e.g., grip, number of hands and filtering [14, 6]). In our study we concentrate only on *control type* and *parallax*, as described below.

Table 1. Technique classification according to the factors.

Parallax \ Control	*Rotational*	*Positional*
None	Laser pointing	Image-plane
Some	Arrow pointing	Fixed-origin

Control Type

As explained previously, the ray of ray pointing can be specified through two points, or through a point and a rotation. Although geometrically equivalent, our four *control types* (Table 1) result in different types of end-user movement. For example, people that use laser pointing and arrow pointing specify the position of the cursor mostly through the rotation of the device (we call this the *rotational control type*), whereas image-plane and fixed-origin techniques only require the specification of a single position on space and the orientation of the limbs or the device is mostly irrelevant (we call this the *positional control type*). Within these categories, we studied our four previously described methods, chosen as they represent design points in the design space defined by these two factors (Table 1).

Parallax

Our other factor of interest is visual parallax. We define visual parallax as the distance between the center of rotation used to specify the pointing direction (usually a device) and the point of view of the user. In real-life aiming activities, parallax is usually avoided if precision is important. For example, sharp shooters align themselves in the direction of their weapons so that the line of view coincides with the shooting direction.

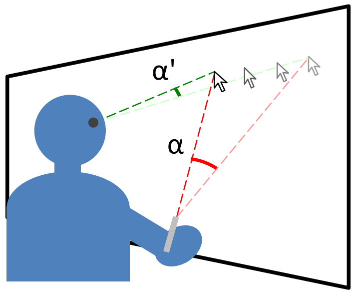

Figure 4. Parallax causes that α ≠ α'.

Our four techniques vary how people perceive parallax. The image-plane technique is, by definition, devoid of parallax. Arrow-pointing transforms laser pointer into an almost parallax-free technique, as the person aligns the pointing device with her line of sight.

EXPERIMENT 1: HORIZONTAL TARGETING

Our first experiment tested *targeting* in the *horizontal* dimension. We were interested in testing targeting separately on this dimension because large displays (e.g., room-sized displays) tend to be much broader than tall, which implies that any effects due to the size of the display and the obliquity of distant areas would be most evident in these tasks, especially if the participant is close to the screen (Figure 5).

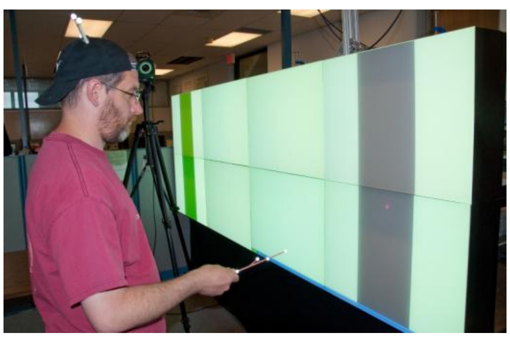

Figure 5. Horizontal task with two distant targets as seen from behind the participant (approximation).

Apparatus

We used a large wall display (292cm x 109cm) composed of 4x2 modular back-projected displays, each one with a resolution of 1024x768px (for a total of 4096x1536px). The modular displays are adjacent to each other with just-perceptible but very narrow image seams (under 2mm). The displays rest on a table 76cm high so that the participant's head lines up approximately with the center of the top row of displays (see Figure 6).

To accentuate the effects of large display widths, we asked participants to stand on a location approximately 73cm from the display and 36cm from its right edge – see Figure 6). From this point of view, the display covered approximately 100º of the user's horizontal field of view, and 68º vertically.

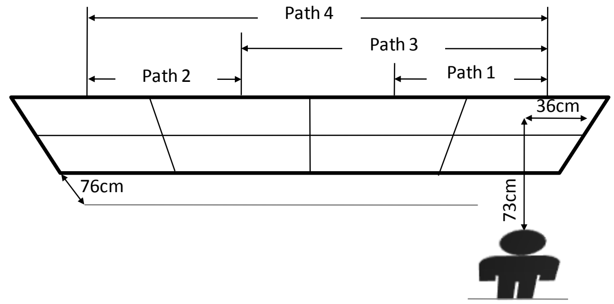

Figure 6. Experimental setup and horizontal paths

We implemented the ray pointing variants using a 25cm wand and a cap equipped with reflective markers, whose positions were tracked by a VICON® motion capture system. The position of the dominant eye of the user was updated in real time by suing the position and orientation of the cap and the calibration data obtained before each block that involved the image-plane technique. Participants selected targets by clicking a mechanical button held in their non-dominant hand (we used a separate button, as pressing a button on the wand could affect its stability).

Experimental software ran on a Core 2 Quad PC running Windows XP. Software was built on the .NET platform and used WPF for presentation. Both image and input were refreshed at a rate well above interactive rate (approx 50Hz refresh rates for display and input).

Task

The horizontal task follows the ISO 9241-9 one-direction tapping test recommendation [8]. Participants had to alternately move the cursor onto each of the target positions that composed a path, and click within its boundaries. The targets were vertical bands that covered the whole height of the display (see Figure 5 and Figure 6).

Targeting tasks varied in the width of the targets (100, 200 and 400px – 7.1, 14.2 and 28.5cm), the distance between target centers (1024, 2048 and 3072px – 73, 146 and 219cm), the position of the targets along the screen, and the direction (left-right or right-left). Figure 6 shows a diagram with the four different paths, which multiplied by three different widths and two directions result in 24 distinct targeting tasks. Visual feedback of errors was provided in the form of color changes of the target.

Participants

Twelve participants recruited from a local university (4 female, 8 male; 24 to 36 years old) took part in the study for $15 remuneration. All participants were right-handed.

Procedure and Design

After signing a consent form each participant provided some basic demographic information, was tested for eye dominance (to determine the dominant eye for the image-plane technique), and received instruction in the four ray pointing techniques.

Technique order was counterbalanced across subjects using a random Latin square design. Participants underwent a block of training for each technique (24 individual trials per technique involving all distances, positions and target widths), and then, in the same order, two separate blocks of actual trials for each technique with three repetitions per individual task. Tasks were presented in order of increasing distance between targets and decreasing target width. After the end of the each technique's trials of the second block, the participants were asked to rate the perceived workload through a NASA TLX questionnaire [17].

At the end of the experiment, participants were asked to rank the techniques according to speed, accuracy, physical effort, and general preference. The entire experimental procedure took approximately 1hour.

Measures and Analysis Methodology

For each trial we measured completion time, location of the cursor during the click, and whether it missed the target (error). We designed the experiment and the analysis to conform to the ISO 9241-9 recommendations [8] as well as the Fitts's study guidelines provided in [24].

As Soukoreff and MacKenzie recommend, we planned error and completion time comparisons as well as throughput comparisons. This requires the calculation of the index of difficulty of each task according to Fitts's law:

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

D is the distance between targets and W the width of targets. However, in a very early stage of the research we realized that targeting tasks have different difficulties depending on their location on the display and the direction of targeting. Consider Figure 5: at the very least, targeting into the farthest region of the display (a distant target) should be harder than targeting onto the near target. Following reasoning parallel to Kopper and colleagues [11], we anticipated that the standard Fitts's model would not capture targeting time differences that can be derived from the geometrical relationships between the person, the display, and the target. Therefore we performed two regressions on the data, one with the standard (linear) version of Fitts's index of difficulty (ID_{Linear}), and one with a variant of the formula that substitutes D and W for the subtended angles of D and W from the location of the user (see Figure 7):

$$ID_{Angular} = log_2\left(\frac{\delta}{\omega} + 1\right)$$

The subtended angles are calculated through standard trigonometric procedures with the generic formula:

$$\alpha = \operatorname{atan}\left(\frac{x_1}{d_p}\right) - \operatorname{atan}\left(\frac{x_2}{d_p}\right)$$

x_1 and x_2 correspond to the horizontal coordinates of the extreme points of the linear distance whose angle we are calculating (see Figure 8). In our experimental setup, ID_{Linear} and ID_{Angular} calculations proved substantially different from each other because of the large size of the display and the position of the user. Figure 9 plots the ID_{Linear} of all tasks against their ID_{Angular}.

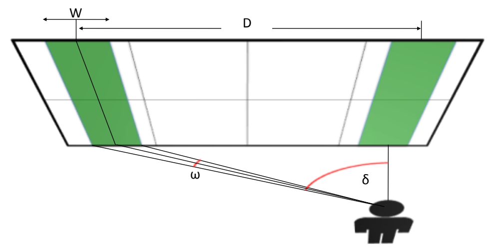

Figure 7. Geometrical relationships between D, W, δ and ω.

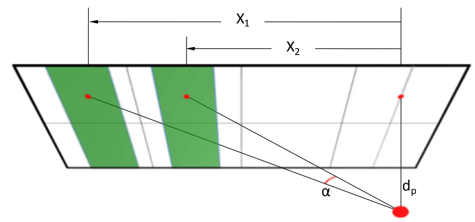

Figure 8. Calculation of subtended angle between x_1 and x_2

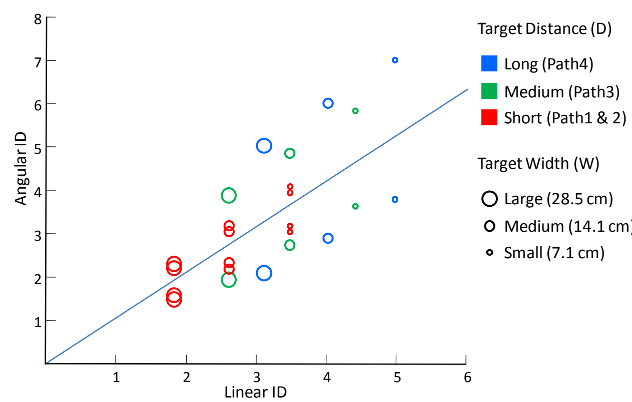

Figure 9. Relationships between linear IDs and angular IDs.

If, as we hypothesize, ID_{Angular} predicts performance significantly better than ID_{Linear} across participants, it would make sense to use this instead to calculate throughput. In either case, for the throughput calculation we apply the effective width corrections as argued in [24]. The calculation of the angle was done using a point 73cm in the direction perpendicular to the top right modular display, which approximates the position of the head of the user.

Results

We begin with our analysis of fit of the linear and angular models, follow by the performance analysis, and end with a summary of the subjective measures results. We performed analysis on throughput, time and error for all tasks. Due to space restrictions, we omit reporting those analyses that are redundant.

Analysis of fit

We did a per-participant, regression analysis of trial completion time for each technique. Using ID_{Linear} as a predictor variable shows an average $R^2 = 0.33$. For ID_{Angular}, the average is $R^2 = 0.61$. That is, using the standard *ID* accounts for only 33% of the variance when used to predict the duration of a horizontal targeting movement. The angular model is much better, as it accounts for roughly twice that proportion. In every single case (all techniques, all participants), the fit of the angular model was superior to the linear model. All p-values of the regression's ANOVA for both models are under 0.001, indicating that the probability of getting these results due to chance is extremely low.

The same two-fold improvement is apparent when we try to fit all data (including differences between participants and between techniques) to particular IDs. With the linear model, the fit is 20% ($R^2 = 0.20$). The same regression with the angular model results in an average fit of 38% ($R^2 = 0.38$).

Analysis of performance measures

On average, the fastest technique was laser pointer (μ = 1015ms) followed by arrow pointing (μ = 1057ms), image-plane pointing (μ = 1139ms) with fixed-origin pointing as the slowest (μ = 1168ms). A repeated-measures ANOVA of throughput (which amalgamates accuracy and speed measures) with technique and task as factors shows a strong effect of technique ($F_{3,33} = 9.99$, $p < 0.001$, $\eta^2 = .47$), where the ordering of average throughput is the same as for completion times (μ_{laser} = 4.05b/s, μ_{arrow} = 3.8b/s, $\mu_{\text{image-plane}}$ = 3.6b/s, $\mu_{\text{fixed-origin}}$ = 3.4b/s). Post-hoc tests (corrected for multiple comparisons) show strong statistical differences between laser and the two worst performing techniques (image-plane and fixed-origin – both $p < 0.004$), while it shows differences that approach statistical significance between laser and arrow ($p < 0.02$) and arrow and fixed origin pointing ($p < 0.022$).

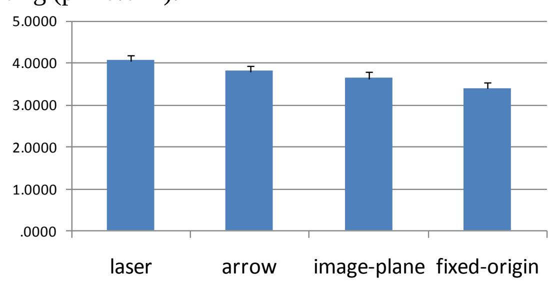

Figure 10. Throughput values (in bits/s) for the horizontal targeting task. Error bars represent standard error.

For error analysis, we performed a non-parametric paired-samples test (Friedman) which shows a significant effect of technique on number of errors ($\chi^2(12) = 10.4$, $p < 0.015$). Fixed-origin pointing was the technique with the largest number of errors (6.3%) followed by arrow pointing (5.15%), laser pointing (4.1%) and image-plane pointing (3.5%).

To summarize, the results of this section indicate a general advantage of laser pointing over image-plane and fixed-origin pointing, whereas arrow pointing finds itself somewhere in between these two groups.

Analysis of subjective measures

Non-parametric paired-measures tests of the subjective workload NASA TLX questionnaires only yielded significant differences between techniques in the *physical demand* ($\chi^2(12) = 9.4$, $p < 0.024$) and *effort* questions ($\chi^2(12) = 8.3$, $p < 0.039$). The mean answers to these questions are shown in Figure 11). Consistent with performance measures, participant preference rankings favored laser and arrow and placed fixed-origin pointing as the least-liked (see Table 2).

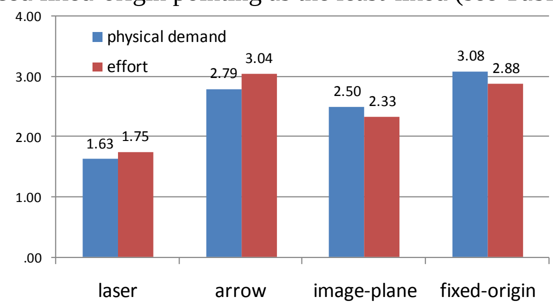

Figure 11. Average physical demand and effort responses (out of a 7-point Likert scale) in the horizontal targeting task (lower means less effort and less physical demand respectively).

Table 2. Preference ranks for the horizontal targeting task.

	Best			Worst
	1	2	3	4
Laser	7	3	2	0
Arrow	4	2	3	3
Image-Plane	1	5	4	2
Fixed-Origin	0	2	3	7

Experiment 2: Vertical Targeting

Our second experiment tests vertical targeting tasks with a double purpose: to generalize the performance results of the horizontal task to vertical movements, and to investigate the effects of the different aspects of parallax in performance. Our techniques were chosen to vary in the amount of parallax (laser pointing and fixed-origin pointing have large parallax, whereas arrow pointing and image-plane pointing have little or no parallax). However, parallax does not affect the horizontal targeting task because the direction of movement is perpendicular to the direction of parallax (i.e., horizontal movement angles are similar for hand and eye, whereas vertical movement angles are very different). Even so, we hypothesized that parallax might affect a vertical task since the targets cover different angles from the eye or from the pointing device.

Apparatus

The apparatus and the location of the participant with respect to the screen was the same than in experiment 1. The task was performed on only two rightmost display modules (see Figure 12).

Task

The vertical task was equivalent to the horizontal task, but in the vertical direction. Pairs of targets were as wide as the modular displays in front of the user, and located at different heights.

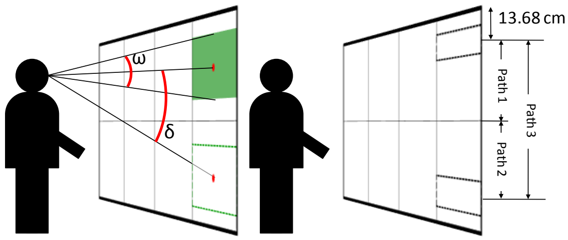

Figure 12. Vertical task angles and paths.

Targeting tasks varied in the height span of the targets (50, 100, and 200px – 37, 73 and 146cm), the position of the targets along the screen (centered at 13.7, 54.7, and 95.8cm from the top of the display) and the path. Figure 12 (right) shows diagrams of the three different paths. This results in 18 distinct targeting tasks when combined with three different widths and two directions (3x3x2).

Participants

A different set of 12 participants (5 female, 7 male; 20 to 40 years old) took part in the study for $15 remuneration. Experiment 2 and 3 shared the same participants.

Procedure and Design

The general procedure was identical to experiment 1, except that participants performed all tasks from experiment 2 and then all tasks from experiment 3 within the same session before they filled the questionnaires (ranking and NASA TLX).

For the vertical targeting task participants carried out three blocks of trials, the first of which was considered training. Each block contained five repetitions of each of the different targeting tasks with each of the techniques. The tasks were presented in increasing order of distance between targets and decreasing target height.

Measures and Analysis Methodology

For each trial we measured completion time, location of the cursor during the click, and whether it missed the target (error).

Since the display we used for our experiment is much broader than tall, we did not expect to find important differences in fit between the angular and linear models; nevertheless we ran regressions with both and used the better fitting model to calculate throughput. For the angular model calculations we used again a fixed point in space that approximates the position of the user's eyes.

Results

Analysis of fit

We run regressions of the data with the angular and linear models and found that the linear model to have a slightly (but statistically significant) better fit than the angular model ($R^2_{Linear} = 0.37 > R^2_{Linear} = 0.33$). Calculations of throughput were thus based on the linear model.

Analysis of performance measures

The averages of task completion times are identical to those found in the horizontal tasks (μ_{laser} = 391ms, μ_{arrow} = 421ms, $\mu_{image\text{-}plane}$ = 453ms, $\mu_{fixed\text{-}origin}$ = 453ms). For throughput, the repeated-measures ANOVA shows a strong effect of technique as well ($F_{3,33} = 8.5$, $p < 0.001$, $\eta^2 = .43$). For the vertical task, arrow had the highest average throughput (μ = 3.89b/s), followed by laser (μ = 3.82b/s), image-plane (μ = 3.48b/s), and fixed-origin pointing (μ = 3.47b/s). The throughput rankings are in slightly different order because throughput depends of both speed and errors.

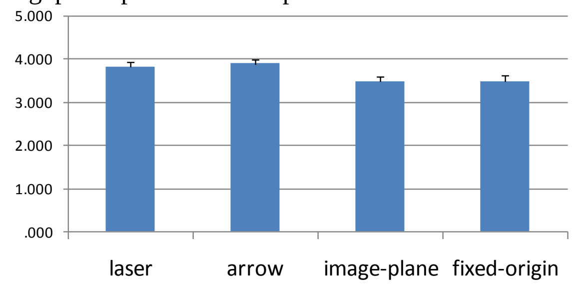

Figure 13. Throughput (in bits/s) for the vertical targeting task. Error bars represent standard error.

The post-hoc tests show statistically significant differences between laser and image-plane (p<0.001), laser and fixed-origin (p<0.003), arrow and image-plane (p<0.003) and arrow and fixed-origin (p<0.009 – this last one is only marginally significant with the Bonferroni adjustment for multiple comparisons $\alpha = 0.05/6 = 0.0083$).

The Friedman test for errors shows statistical differences ($\chi^2(12) = 9.14$, $p < 0.027$). For vertical tasks, fixed-origin pointing has again the largest error rate (3.6%), followed by laser (3.4%), arrow (3.2%) and with image-plane again as the most accurate (2.1%).

In summary, performance in the vertical task is similar to the horizontal task, except that post-hoc tests show more power: we can completely separate techniques in two groups of performance, with laser and arrow outperforming the rest.

Analysis of subjective measures

Because experiments 2 and 3 were grouped, the post-study questionnaire will be discussed as part of experiment 3.

Experiment 3: Tracing

Finally, we examined the performance of our techniques with a tracing (*aka steering)* task.

Apparatus

The apparatus was identical to experiment 2, although tracing tasks took place across the whole width of the display.

Task

The screen presented a rectangle (a 'tunnel') with a square at one end (Figure 14). Participants were instructed to enter the rectangle through the non-square entrance, and reach the square while remaining within the rectangular tunnel. When the square was reached, it disappeared and a square appeared at the other end starting a new trial (Figure 14).

The tunnels were always 384x96 pixels (27.4x6.8cm), which results in a tracing index of difficulty of 4 (length/width). However, they appeared in four different locations (centered on the modular displays at the corners of the display wall) and with two different orientations (horizontal and vertical), for a total of eight different tracing tasks (see Figure 14). We did not consider different directions (e.g., top-bottom and bottom-top) as different tasks because the angles covered are the same, i.e., they are independent of the tracing direction. For error feedback, we changed the tunnel's outline color when the subject moved the cursor outside of it.

Figure 14. Tracing task tunnels and positions.

Participants

Participants were the same than in experiment 2.

Design

Participants carried out three blocks of tracing tasks for each technique, with the first block for training. For each block, participants performed six repetitions of each of the eight tasks for each technique. Tasks were presented always in the same order (vertical to horizontal, top to bottom, left modular display to right modular display).

Measures and Analysis Methodology

For each trial we measured completion time and average distance to the longitudinal line of the tunnel.

We tested the model fit with linear and angular measures (analogous to those in experiment 1). For the calculation of *ID* in tracing we used the formula suggested by the ISO 9241-0 standard:

$$ID = \frac{D}{W}$$

D is the length of the tunnel and W is the width that the cursor can move transversally without leaving of the tunnel (our cursor was considered to be of zero diameter).

For the angular calculations, we used approximations based on the perceived angles of width and length at the center of each tunnel. Angles were calculated from the same fixed point used for experiment 1. Angular IDs for the eight tasks range from 1.27 to 12.60b/s. As for experiment 1, we would use the *IDs* of the model that offered the best fit for the throughput calculation, and apply the corresponding adjustment for accuracy.

Results

Analysis of fit

The angular model shows an average fit of the data comparable to the fit for the horizontal targeting task ($R^2 = 0.50$). Since we only tested one linear index of difficulty, comparing the fit of both models is equivalent to testing whether the angular regression is significant. The ANOVA test measure of the angular regression does exactly that. This result held in all but one of the 48 regressions (12 participants x 4 techniques), with $p < 0.002$. An omnibus regression using all participants and all techniques simultaneously produced a $p < 0.0001$, and $R^2 = 0.37$.

Analysis of performance measures

The repeated-measures ANOVA of throughput calculated from the angular indexes of difficulty shows a strong main effect of technique ($F_{3,33} = 12.227$, $p < 0.001$, $\eta^2 = .53$). Image-plane had, on average, the highest throughput (μ = 107b/s), followed by arrow (μ = 83b/s). Laser and fixed-origin exhibited much lower performance (μ_{laser} = 56b/s, $\mu_{fixed\text{-}origin}$ = 65b/s). Post-hoc tests statistically differentiate image-plane from laser and fixed-origin (both $p < 0.002$), and arrow from laser ($p < 0.001$). Note that the throughput values for tracing are not necessarily comparable with those of targeting, and that the adjustment for accuracy suggested in [24] greatly increases the effective angular *ID*s.

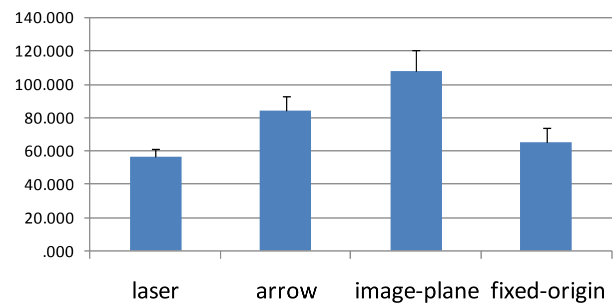

Figure 15. Throughput (in b/s) for the tracing task. Error bars indicate standard error.

A repeated-measures ANOVA of the average deviation with respect to the middle of the tunnel also shows a strong effect of technique ($F_{3,33} = 4.9$, $p < 0.006$, $\eta^2 = 0.30$). The most accurate technique was image-plane (μ = 18.4px) followed by the rest in a very tight group (μ_{arrow} = 21.7px, μ_{laser} = 22.8px, $\mu_{fixed\text{-}origin}$ = 22.8px). Post-hoc tests only show statistically significant (or marginally significant) differences between image-plane and the rest.

Analysis of subjective measures

As the post-study questionnaire encompassed experiments 2 and 3, the subjective measures reflect the opinions of subjects in both tasks (vertical targeting and tracing).

The measures of the NASA TLX only showed statistical differences between techniques for the *physical demand* question ($\chi^2(12) = 7.8$, $p < 0.024$). Interestingly, subjects judged the arrow as the least physically demanding technique, whereas the horizontal targeting task arrow was considered the second most demanding.

Table 3. Preference ranks for experiment 2.

	Best			Worst
	1	2	3	4
Laser	**2**	**2**	**5**	**3**
Arrow	**5**	**5**	**2**	**0**
Image-Plane	**1**	**5**	**1**	**5**
Fixed-Origin	**4**	**0**	**4**	**4**

The preference rankings (Table 3) also show very different results than experiment 1.

DISCUSSION

We discuss the findings around three topics: the ray pointing technique variants, the effect of the large display, and the limitations of the experiments.

Differences Between Ray Pointing Variants

Our targeting experiments revealed differences in targeting performance of up to 10%. Arrow and laser (both rotational techniques) performed better than their positional counterparts, even for vertical targeting for which we had hypothesized that parallax would have a strong influence. Counter to our intuitions, this suggests that parallax is *not* important for targeting tasks.

The advantage of rotational techniques over image-plane pointing is somewhat surprising because it contradicts the mounting evidence from VR studies that found image-plane to be more efficient. This evidence implies that targeting tasks are different for 3D and 2D environments, and that practitioners and designers alike should be careful not to extrapolate results from one field into the other, regardless of the apparent similarity of the tasks.

When looking at the tracing task, the story takes a different turn: laser pointer is no longer the leading technique. Instead, image-plane and arrow (the two parallax-free techniques) perform best, which suggests that parallax is a critical factor for this kind of task. Designers of large display systems should therefore consider how their selection of ray pointing technique fits with the expected task (e.g., steering vs. selecting). The fact that the same technique can produce such different results also opens interesting questions. Will users naturally adopt the most advantageous use mode of the device? Alternately, how can we design a pointer that encourages the best use of the device?

Overall, our results suggest that arrow is a good choice for both task types; it performs close to laser pointing in selection tasks, and better in steering tasks. Arrow is also relatively cheap to implement compared to image-plane (which requires some degree of head-tracking). However, participants did find arrow-pointing tiring (at least for the horizontal targeting task). Longer term studies should test whether physical effort is really a serious issue.

We anticipate that the magnitude of differences found in our experiments will be relevant in many cases (e.g., when performance is very important or errors are very costly). More extreme effects are also possible for larger displays and smaller targets. Nevertheless, designers should always put the magnitude of the effects into context, especially if other tradeoffs are in effect (e.g., system comfort, cost).

Ray Pointing In Very Large Displays

Our regression analysis in the horizontal targeting and the tracing experiments showed strong support for the use of a model based on angular *vs.* linear distances in displays that cover a wide angle. This is relevant for the design of interfaces of very large displays. For example, an angular model predicts large differences between tasks that go from center to left of the screen and vice versa, whereas the linear model does not.

We are not the first to suggest an angular adaptation of ID calculation, but our data strongly supports a model that retains the simplicity and the number of variables of the original, with superior fit. The closest previous work by [11] suggests complex models with extra variables, and which provide only an incremental increase of fit with respect to our angular model.

Although previous research shows much higher R^2 values than ours, this is because they average trial completion times before fitting the model. Importantly, that fit measure does not account for the natural differences between repetitions of the same task. Our method is certainly not the only way to perform regression analysis, but we believe that fitting more data helps put the fit differences into perspective (a small increase in the fit of a model will result in an even smaller – or even inferior – fit measure when the data includes more of the real variability).

Limitations

Any experiment is necessarily limited in the amount of conditions and factors that it can test, and ours is not an exception. Most notably, our experiment only tested one distance from the display. It is possible that other models better reflect distance variability, but this remains an open question.

Our evidence on the factors that cause technique performance differences is also not definitive. For example, establishing a strong causal link between parallax and poor tracing performance will require further research.

It is also possible that the targeting results for the laser variants are exaggerated by the familiarity of the device in comparison with the positional techniques. We added significant amounts of training to the experiment to avoid this problem, but this issue also requires further research.

Finally, there are a number of small experimental manipulations that should be studied to paint a more complete picture of ray pointing performance. For example, do variant differences interact with the location of the button click, or with handedness?

Lessons for Practitioners

We summarize the implications for researchers and designers in four main points:

- If targeting is important, choose laser-style ray pointing
- For modeling targeting and tracing tasks that span a large angle from the point of view of the user, use a performance model that takes angles into account
- For tracing tasks choose parallax-free pointing
- Training people to use a pointing device from a location close to the eye (arrow pointing) provides good performance for both targeting and tracing tasks

CONCLUSION AND FUTURE WORK

As very large displays become more common, it becomes increasingly important to provide input mechanisms that are appropriate for interaction at a distance. Ray pointing techniques allow convenient interaction at a distance. We looked at two factors that influence ray pointing: control type and parallax. We tested four ray casting variants and two tasks (targeting and tracing), and found evidence that control type affects targeting and parallax affects tracing. Furthermore, we present an angular adaptation of ID calculation. This adaptation fits our data better than a linear model, suggesting that perceived angle must be taken into consideration when designing for large displays.

ACKNOWLEDGMENTS

Thanks to our sponsors NSERC, iCORE and SMART Technologies (Canada) and Ricardo Jota's support from FCT Portugal (SFRH/BD/ 17574/2004). Thanks also to Alex Hill for insightful discussions, Eileah Trotter for helping run the studies, and to all participants for their effort.

REFERENCES

1. Ahlborn, B.A., Thompson, D., Kreylos, O., Hamann, B., and Staadt, O.G. A practical system for laser pointer interaction on large displays. *Proc. VRST'05*, 106-109.
2. Bolt, R. A. 1980. "Put-that-there": Voice and gesture at the graphics interface. *Proc.* SIGGRAPH '80, 262-270.
3. Bowman, D.A. and Hodges, L.F. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. *Proc. Interactive 3D graphics '97*, 35.
4. Bowman, D.A., Johnson, D.B., and Hodges, L.F. Testbed Evaluation of Virtual Environment Interaction Techniques. *Presence: Teleoperators & Virtual Environments 10*, 1 (2001), 75-95.
5. Cavens, D., Vogt, F., Fels, S., and Meitner, M. Interacting with the big screen: pointers to ponder. *Ext. Abs. CHI '02*, 678-679.
6. Davis, J. and Chen, X. Lumipoint: multi-user laser-based interaction on large tiled displays. *Displays 23*, 5 (2002), 205-211.
7. Hill, A. and Johnson, A. Withindows: A Framework for Transitional Desktop and Immersive User Interfaces. *Proc. 3DUI'2008*, 3-10.
8. ISO. Iso 9241-9:2000(e): Requirements for non-keyboard input devices (iso 9241-9). 2002.
9. Jota, R., Pereira, J.M., and Jorge, J.A. A comparative study of interaction metaphors for large-scale displays. *Ext. Abs. CHI'09*, 4135-4140.
10. König, W. A., Gerken, J., Dierdorf, S., and Reiterer, H. Adaptive pointing: implicit gain adaptation for absolute pointing devices. *Ext. Abs. CHI'09*, 4171-4176.
11. Kopper, R., Bowman, D., and Silva, M. A Human Motor Behavior Model for Distant Pointing Tasks. Virginia Tech. Tech.Rep. (2008). http://eprints.cs.vt.edu/archive/ 00001056/01/chi2009-DistantPointing-Final.pdf
12. Lee, S. Evaluation of pointing techniques for ray casting selection in virtual environments. *Proc. SPIE*'2003, 38.
13. Mackenzie, I. and Jusoh, S. An Evaluation of Two Input Devices for Remote Pointing. In *Engineering for Human-Computer Interaction*. 2001, 235-250.
14. Myers, B.A., Bhatnagar, R., Nichols, J., et al. Interacting at a distance: measuring the performance of laser pointers and other devices. *Proc. CHI'02*, 33-40.
15. Myers, B., Peck, C., Nichols, J., Kong, D., and Miller, R. Interacting at a Distance Using Semantic Snarfing. *Proc. Ubicomp 2001*, 305-314.
16. Nacenta, M. A., Pinelle, D., Stuckel, D., and Gutwin, C. 2007. The effects of interaction technique on coordination in tabletop groupware. *Proc. GI'07*, 191-198.
17. NASA Human Performance Research Group. Task Load Index (NASA-TLX) 1987. NASA
18. Olsen, Dan., and Nielsen, T. Laser pointer interaction. *Proc. CHI 2001*, 17-22.
19. Parker, J. K., Mandryk, R. L., and Inkpen, K. M. 2005. TractorBeam: seamless integration of local and remote pointing for tabletop displays. In Proc. GI'05, 33-40.
20. Pavlovych, A. and Stuerzlinger, W. The tradeoff between spatial jitter and latency in pointing tasks. *Proc. EICS*'09, 187-196.
21. Peck, C.H. Useful parameters for the design of laser pointer interaction techniques. *Ext.Abs. CHI '01*, 461-462.
22. Pierce, J.S., Forsberg, A.S., Conway, M.J., Hong, S., Zeleznik, R.C., and Mine, M.R. Image plane interaction techniques in 3D immersive environments. *Proc Interactive 3D graphics 1997*, 39.
23. Shoemaker, G., Tang, A., and Booth, K.S. Shadow reaching: a new perspective on interaction for large displays. *Proc. UIST'*07, 53-56.
24. Soukoreff, R.W. and MacKenzie, I.S. Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. *Int. J. Hum.-Comput. Stud. 61*, 6 (2004), 751-789.
25. Stuerzlinger, W and Oh, J. laser pointers as collaborative pointing devices. Proc. of Graphics interface 2002.
26. Teather, R.J. and Stuerzlinger, W. Assessing the Effects of Orientation and Device on (Constrained) 3D Movement Techniques. *Proc. 3DUI*, 43-50.
27. Tse, E., Hancock, M., and Greenberg, S. Speech-filtered bubble ray: improving target acquisition on display walls. *Proc Multimodal Interfaces 2007*, 307-314.
28. Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large, high resolution displays. *Proc. UIST'2005*, 33-42.
29. Voida, S., Podlaseck, M., Kjeldsen, R., and Pinhanez, C. A study on the manipulation of 2D objects in a projector/camera-based augmented reality environment. *Proc. CHI'*05, 611-620.
30. Ware, C. and Lowther, K. Selection using a one-eyed cursor in a fish tank VR environment. *ACM Trans. Comput.-Hum. Interact. 4*, 4 (1997), 309-322.
31. Wingrave, C., Tintner, R., Walker, B., Bowman, D., and Hodges, L. Exploring individual differences in raybased selection: strategies and traits. *Virtual Reality, 2005*, 163-170