# A Comparison of Ray Pointing Techniques for Very Large Displays

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

[1]Inesc-ID, [2]University of Calgary

**Abstract**

Ray-pointing techniques are often advocated as a way for people to interact with very large displays from several meters away. We are interested in two factors that can affect ray pointing: the particular technique's control type, and parallax.

Consequently, we tested four ray pointing variants on a wall display that covers a large part of the user's field of view. Tasks included horizontal and vertical targeting, and tracing. Our results show that (a) techniques based on 'rotational control' perform better for targeting tasks, and (b) techniques with low parallax are best for tracing tasks. We also show that a Fitts's law analysis based on angles (as opposed to linear distances) better approximates people's ray pointing performance.

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

**Index Terms:** H5.2.User Interfaces: Input devices and strategies

## 1 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 [2,20] as: it allows people to interact from the distance (as opposed to direct-input) [21,31]; it does not require any physical surface to operate on (as opposed to mouse-based pointing) [27,28]; it is easily understood by people as it builds upon everyday pointing gestures [31]; and it allows multiple users to interact on the same display without their bodies physically getting in the way [17]. Thus it is no surprise that ray pointing is increasingly used in both commercial and research systems, especially for large horizontal and vertical displays [6,17,21,30]. Even game consoles are exploiting ray pointing for interaction (e.g., Nintendo Wii).

As large and very large displays (i.e., wall-sized displays) become widespread, ray pointing will likely become a primary way to interact with all kinds of interfaces from a distance or to access hard to reach areas of the display. This is why there is already a significant amount of literature devoted to the performance and different modes of ray pointing; for example, different variants of ray pointing are used and studied by Virtual Reality researchers [1,4,24], and there have been some efforts to characterize laser pointing (one of the possible ray pointing variants) for interaction with 2D interfaces [14,15,22,27]. However, in our own design and implementation of large-display interfaces we have found that previous work in the area does not suffice because of several reasons. A) Although the largest diversity of ray pointing alternatives has been studied in VR, ray pointing tasks and setups (e.g. CAVEs, stereoscopic displays) are substantially different than the more common 2D tasks that we are interested in. B) Previous empirical work on 2D tasks is focused almost exclusively on laser pointing or studies only small displays [14,16,27]. C) With a few exceptions [12,15] previous work for 2D environments does not try to provide explanations or general principles of the differences in ray pointing.

[1]Rua Alves Redol 9, Apart. 13069 1000-029 Lisboa, Portugal {jota.costa,jaj}@inesc-id.pt [2]Comp. Sci., 2500 University Dr. NW. Calgary, AB, Canada T2N 1N4. {miguel.nacenta, saul.greenberg, sheelagh.carpendale}@ucalgary.ca

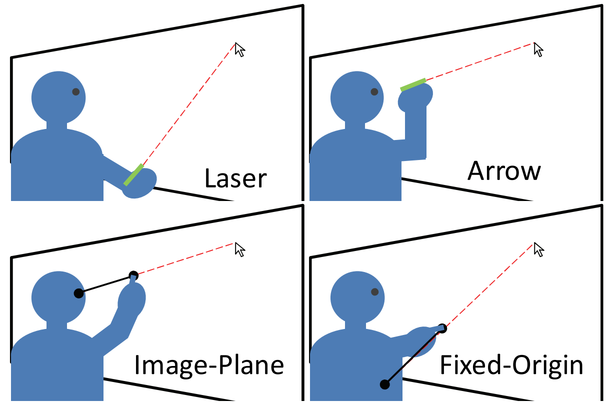


**Figure 1**. Four variants of ray pointing

We build upon this previous research by contributing, via a new experiment: a) tests of two fundamental 2D interaction tasks (targeting and tracing), one of which has not previously been studied in the context of ray pointing; b) comparisons of four different ray pointing variants (laser pointing, arrow pointing, image-plane pointing and fixed-origin pointing) which map to two previously unstudied factors relevant to ray pointing in 2D tasks (parallax and control type); and c) identification of specific issues related to the large size of the displays, like the effect of the location of targets with respect to users.

Our experimental results show that targeting performance is best explained by the control type factor, with rotational control being generally superior to position control, whereas for tracing tasks it is the presence of parallax that better explains differences between variants. The study also contributes rigorous support for the use of an angular formulation of Fitts's law [10,12] for large-displays, as opposed to the traditional linear formulation for both tasks.

These results have important implications for designers and researchers alike. First, designers must now consider image-plane techniques and how the pointing device is held, especially if they want to support tracing tasks (e.g., menu selection, drawing). Second, researchers can now consider parallax and control type as important factors, and have solid evidence to start using angular adaptations of Fitts's law for interaction in large displays.

## 2 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. We consider only monocular absolute 2D pointing.

Jota, R., Nacenta, M., Jorge, J., Carpendale, S. and Greenberg, S. (2010)
A Comparison of Ray Pointing Techniques for Very Large Displays. In Proceedings of Graphics Interface - GI'2010. ACM Press, 8 pages, May 31 - June 2.

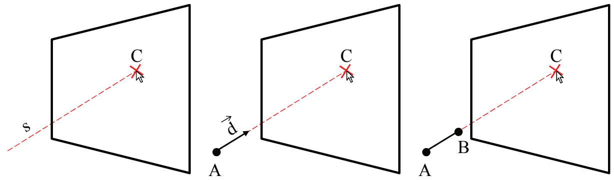


**Figure 2.** ***Left)*** Ray pointing specifies the location of the cursor through the intersection of a ray (s) with the display surface. ***Center)*** the ray (s) can be specified through a point (A) and a direction, or ***Right)*** through two points (A and B)

### 2.1 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 1). 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 [30].

Laser pointing has been proposed and implemented for cursor control in 2D interfaces many times (e.g., [2,5,6,20,27,29]). 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.

### 2.2 Arrow Pointing

*Arrow pointing* is a variant of laser pointing where we constraint the use of the pointer to be somewhat aligned with the user's eye (Figure 1). 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.

### 2.3 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 for the selection of 3D objects, where it is referred to as *image-plane manipulation*, *occlusion selection*, or the *crushing heads* technique [1,4,7,13,32,33].

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). 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.

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.

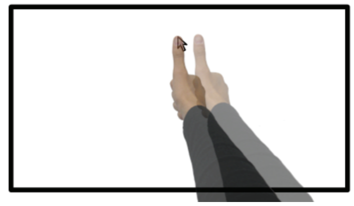

**Figure 3.** Image-plane pointing seen binocularly & focused on the distant display (cursor displayed on screen).

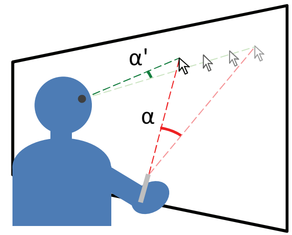


**Figure 4.** Parallax causes α ≠ α'

### 2.4 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). This was explored somewhat by Shoemaker and colleagues in *shadow reaching* [25]. 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.

We tested *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.

## 3 Related Studies

Previous research in ray pointing falls mostly into two categories: laser pointing studies for distant displays, and virtual reality techniques for object manipulation in 3D worlds. In this section we also discuss modeling of targeting performance and existing enhancements to ray pointing.

**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 [14] and Stuerzlinger and Oh [27] showed that laser pointing targeting performance is poor compared to the mouse (and around 1.4b/s or 3.0b/s respectively). Peck [23] parameterized the jitter of a laser pointer spot in terms of angle, and suggests that grip affects it. Myers and colleagues [15] studied the effect of different grips and postures, and found reduced jitter with a PDA-pointer held with two hands.

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 [12]).

**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 head-tracking is already present). 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 [1,3,4,13,32,33]. 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.

**Modeling of the Targeting Task.** Kondraske [10] suggested that rotational tasks (e.g. rotating a knob to a particular position) are better modeled by Fitts's law if the angular distances are taken into account. This modeling approach is relevant for ray pointing because the movement of the hand is often rotational, but also because, with large displays, linear measures do not represent faithfully *perceived distances and sizes* as seen by the user (e.g., targets of the same size on distant areas of the display appear narrower than targets close to the user). Kopper et al. [12] explored several models, including some based on angles, and others including a distance parameter. To date, however, linear models for ray pointing tasks are still predominant, since there is no strong evidence supporting a substantial modeling benefit of using angles.

**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 [6,30], its CD gain altered [11], alternative modalities blended into the action [29], and snapping mechanisms added [16,29]. We know that some of those mechanisms may improve pointing (e.g., fitering) but these also imply trade-offs (e.g., filtering implies delay [22], and *semantic snarfing* [16] 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.

## 4 Ray Pointing Performance Factors

There are many possible factors that might affect ray pointing performance (e.g., grip, number of hands and filtering [6,15]). In our study we concentrate only on *control type* and *parallax*, as described below. We chose these two factors based on previous empirical results from the literature (e.g., [18]) that have not been studied in relation to ray pointing.

Table 1. Technique classification according to the factors.

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

### 4.1 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 two *control types* 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).

### 4.2 Parallax

Our other factor of interest is visual parallax (Figure 4). We define visual parallax as the distance between the center of rotation used to specify the pointing direction (usually the center of 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.

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.

## 5 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).

### 5.1 Method

**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 5).

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 5). From this point of view, the display covered approximately 100º of the user's horizontal field of view and 68º vertically.

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 using 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). In all techniques a circular cursor was displayed on the screen.

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 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 height of the display (see Figure 5).

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 5 (Right) 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.

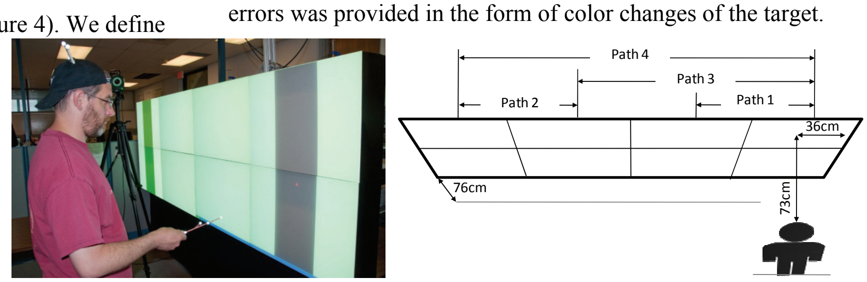


**Figure 5.** ***Left)*** The experimental setup during a horizontal targeting task. ***Right)*** Location of the participant and paths for the horizontal targeting task.

**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 basic demographic information, was tested for eye dominance (to determine the dominant eye for the image-plane technique), and received instruction in the four 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 [19].

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 1 hour.

**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 [26].

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

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

Fitts's law. 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 [12], 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 ($ID_{\text{Linear}}$), and one with a variant of the formula that substitutes D and W for the subtended angles of D and W ($\delta$ and $\omega$) from the location of the user (see Figure 6):

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

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

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

$x_1$ and $x_2$ correspond to the horizontal coordinates of the extreme points of the linear distance whose angle we are calculating (see Figure 7). In our experimental setup, $ID_{\text{Linear}}$ and $ID_{\text{Angular}}$ calculations proved substantially different from each other because of the large size of the display and the position of the user. Figure 8 plots the $ID_{\text{Linear}}$ of all tasks against their $ID_{\text{Angular}}$.

If, as we hypothesize, $ID_{\text{Angular}}$ predicts performance significantly better than $ID_{\text{Linear}}$ 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 [26]. 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.

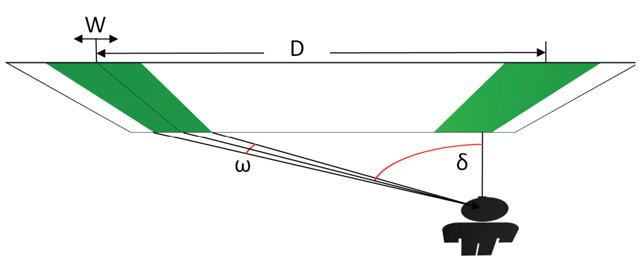


**Figure 6.** Geometrical relationships between D, W, δ and ω.

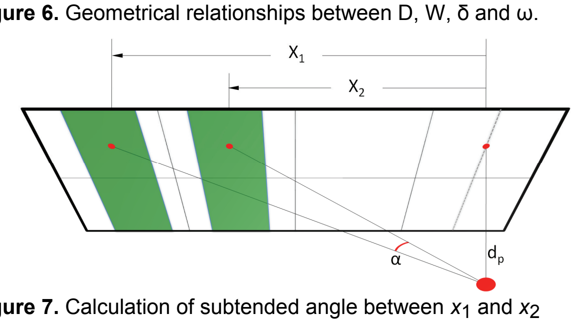


**Figure 7.** Calculation of subtended angle between $x_1$ and $x_2$

## 5.2 Results

We begin with our analysis of fit of the linear and angular models, follow with 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 $ID_{\text{Linear}}$ as a predictor variable shows an average $R^2 = 0.33$. For $ID_{\text{Angular}}$, 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 – see Figure 9). A repeated-measures ANOVA of throughput (which amalgamates accuracy and speed measures) with technique and task as factors

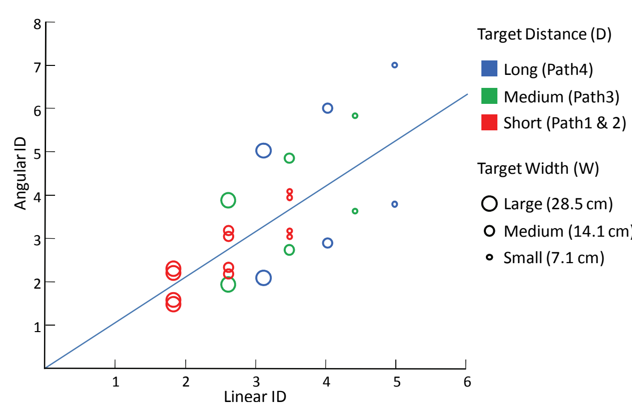


**Figure 8.** Relationships between linear IDs and angular IDs.

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 ($\mu_{laser} = 4.05$b/s, $\mu_{arrow} = 3.8$b/s, $\mu_{image\text{-}plane} = 3.6$b/s, $\mu_{fixed\text{-}origin} = 3.4$b/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 ($p < 0.022$).

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 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 10). Consistent with performance measures, participant preference rankings favored laser and arrow and placed fixed-origin pointing as the least-liked (see Table 2).

## 6 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,

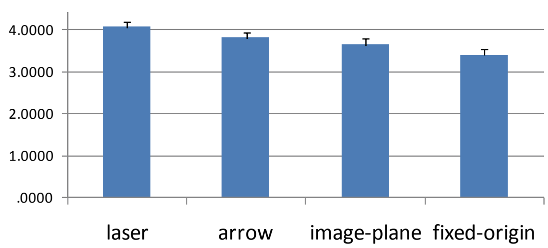


**Figure 9.** Throughput values (in bits/s) for the horizontal targeting task. Error bars represent standard error.

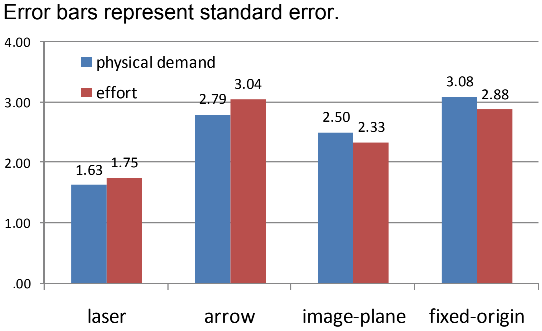


**Figure 10.** Average physical demand and effort responses (out of a 7-point Likert scale) in the horizontal targeting task (lower means less effort and less physical demand respectively).

**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 |

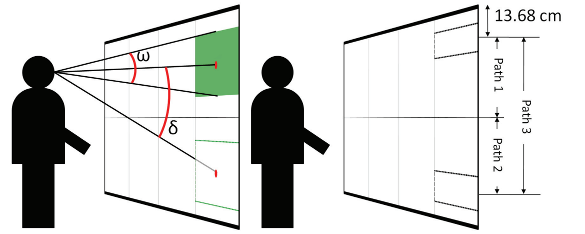


**Figure 11**. Vertical task angles and paths.

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.

### 6.1 Method

**Apparatus.** The apparatus and location of the participant with respect to the screen was identical to experiment 1. The task was performed on only two rightmost display modules (Figure 11).

**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.

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 11 (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. Experiments 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 user's eye position.

### 6.2 Results

**Analysis of fit.** We run regressions of the data with the angular and linear models and found the linear model to have a slightly (but statistically significant) better fit than the angular model

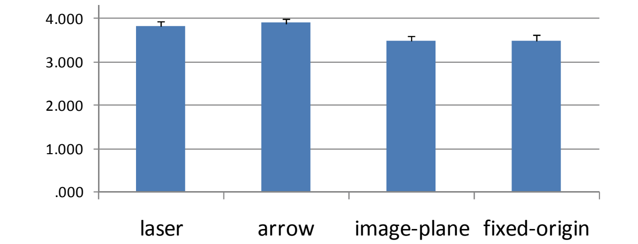


**Figure 12.** Throughput (in bits/s) for the vertical targeting task. Error bars represent standard error.


**Figure 13.** Tracing task tunnels and positions.

($R^2_{Linear} = 0.37 > R^2_{Angular} = 0.33$). Calculations of throughput were thus based on the linear model.

**Analysis of performance measures.** The averages of task completion times are in identical order to those found in the horizontal tasks ($\mu_{laser}$ = 391ms, $\mu_{arrow}$ = 421ms, $\mu_{image\text{-}plane}$ = 453ms, $\mu_{fixed\text{-}origin}$ = 453ms – see Figure 12). 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 ($\mu$ = 3.89b/s), followed by laser ($\mu$ = 3.82b/s), image-plane ($\mu$ = 3.48b/s), and fixed-origin pointing ($\mu$ = 3.47b/s). The throughput rankings are in slightly different order because throughput depends of both speed and errors.

The post-hoc tests show 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 ($p<0.009$ – the latter is only marginally significant with a 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%); image-plane was 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.

## 7 EXPERIMENT 3: TRACING

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

### 7.1 Method

**Apparatus and participants** were 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 13). 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.

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 13). 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 feedback, we changed the tunnel's outline color when the cursor was moved outside of it.

**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. *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).

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

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.

### 7.2 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 ($\mu$ = 107b/s), followed by arrow ($\mu$ = 83b/s). Laser and fixed-origin exhibited much lower performance ($\mu_{laser}$ = 56b/s, $\mu_{fixed\text{-}origin}$ = 65b/s – see Figure 14). 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 [26] greatly increases the effective angular *ID*s.

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 ($\mu$ = 18.4px) followed by the rest in a very tight group ($\mu_{arrow}$ = 21.7px, $\mu_{laser}$ = 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 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

**Table 3**. Preference ranks for experiment 2 and 3.

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

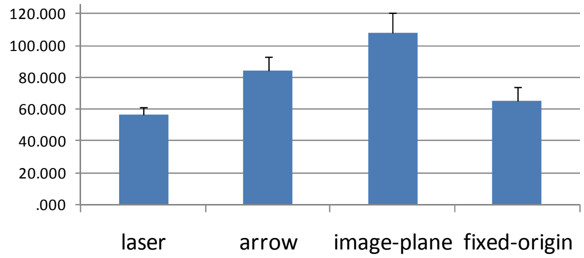


**Figure 14**. Throughput (in b/s) for the tracing task. Error bars indicate standard error.

($\chi^2(12) = 7.8$, $p < 0.024$). Interestingly, subjects judged the arrow as the least physically demanding technique, whereas the horizontal targeting task arrow was considered the second most demanding. The preference rankings (Table 3) also show very different results than experiment 1.

## 8 Discussion

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

### 8.1 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 seems to suggest that parallax is *not* important for targeting tasks, at least when cursor feedback is present (as in our experiment).

The advantage of rotational techniques over image-plane pointing is somewhat surprising because it contradicts some of the 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 depending on how the device is held 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 pointing is a good choice for activities that require both targeting and tracing; 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). Although participants found arrow pointing tiring for the horizontal targeting task, they did not for experiments 2 and 3. Longer term studies should test whether physical effort is really a serious issue; however, the current data suggests that, when comparing ray-pointing techniques, people's perception of effort is linked to their ability to perform well with that technique rather than to the differences on how the device is held.

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).

### 8.2 Models for Ray Pointing In Very Large Displays

We are not the first to suggest an angular adaptation of ID calculation for ray-pointing tasks. Kopper et al. [12] proposed a number of alternative models for ray pointing, some of which use angular measures. Although their data and our results both combine to support the use of alternative (angular) formulations of Fitts's law for ray pointing, their focus was specifically on comparing (slightly different) models for what we call laser pointing. Their fit calculations are performed after averaging all participants' trials for each ID, and result in very high fit values ($R^2 > .75$), even for models using linear distance and width. This approach is problematic because it eliminates most within-participant variance, and only highlights small differences in fit between models that do not strongly advocate for any particular model.

Instead, we decided to compare the linear formulation against the simplest angular model, without artificially eliminating variance. This resulted in a more accurate estimation of the real variance explained by the models and, most importantly, on very large differences in fit between the two models (up to 84% increase in fit) that strongly advocate for the use of angular models.

Additionally, the differences that we found in model fit were consistent across all techniques (not only ray pointing) and across tasks (not only targeting). This suggests that the better fit of the angular model is not only due to the rotational nature of input, but also due to the differences in perception of objects that are at different angles and are seen from different perspectives.

Having this strong evidence to support angular models is critical for research and design of large display interaction; using the linear model in large-display experiments (the current approach) will introduce a large amount of noise, which can dramatically reduce the power of statistical tests. Similarly, when using targeting and menu activation estimations for the design of large display interfaces, it is important to know that different locations in the display will be affected differently depending on the position of the user.

### 8.3 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 (see also the discussion in [12]).

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.

Because we found some differences between horizontal and vertical tasks, new experiments that investigate the weight of each for diagonal or circular tasks could also be useful.

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?

### 8.4 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

## 9 Conclusion

As very large displays become more common, it becomes increasingly important to develop input mechanisms appropriate for interaction at a distance: the class of pointing techniques that currently offers the most promise. As such, these techniques deserve very close scrutiny and refinement under various conditions. Specifically, 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 provide strong evidence that supports the use of angular indexes of difficulty for targeting and tracing tasks with any of the absolute ray pointing techniques.

## 10 Acknowledgements

This research was partially supported by NSERC, by iCORE, and by SMART Technologies, Inc. It has also received funding from the European Community's 7th Framework Programme (FP7/2007-2013) under grant agreement n° [217039]. Ricardo Jota was supported by the Portuguese Foundation for Science and Technology (SFRH/BD/17574/2004).

## References

[1] Argelaguet, F. and Andujar, C. Visual feedback techniques for virtual pointing on stereoscopic displays. *Proc. VRST'09*, 163-170.

[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] Kondraske, G.V. An Angular Motion Fitts' Law For Human Performance Modeling And Prediction. Proc. Engineering in Medicine and Biology Society 1994, 307.

[11] 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.

[12] Kopper, R., Bowman, D., and Silva, M. A Human Motor Behavior Model for Distant Pointing Tasks. Virginia Tech. Technical Report. (2008). http://eprints.cs.vt.edu/archive/00001056/01/chi2009-DistantPointing-Final.pdf

[13] Lee, S. Evaluation of pointing techniques for ray casting selection in virtual environments. *Proc. SPIE*'2003, 38.

[14] Mackenzie, I. and Jusoh, S. An Evaluation of Two Input Devices for Remote Pointing. In *Engineering for Human-Computer Interaction*. 2001, 235-250.

[15] 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.

[16] Myers, B., Peck, C., Nichols, J., Kong, D., and Miller, R. Interacting at a Distance Using Semantic Snarfing. *Proc. Ubicomp 2001*, 305-314.

[17] 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.

[18] Nacenta, M.A., Sallam, S., Champoux, B., Subramanian, S., and Gutwin, C. Perspective cursor: perspective-based interaction for multi-display environments. *Proc. CHI'06*, 289-298.

[19] NASA Human Performance Research Group. Task Load Index (NASA-TLX) 1987. NASA

[20] Olsen, Dan., and Nielsen, T. Laser pointer interaction. *Proc. CHI 2001*, 17-22.

[21] 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.

[22] Pavlovych, A. and Stuerzlinger, W. The tradeoff between spatial jitter and latency in pointing tasks. *Proc. EICS*'09, 187-196.

[23] Peck, C.H. Useful parameters for the design of laser pointer interaction techniques. *Ext.Abs. CHI '01*, 461-462.

[24] 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.

[25] Shoemaker, G., Tang, A., and Booth, K.S. Shadow reaching: a new perspective on interaction for large displays. *Proc. UIST'*07, 53-56.

[26] 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.

[27] Stuerzlinger, W and Oh, J. laser pointers as collaborative pointing devices. Proc. of Graphics interface 2002.

[28] Teather, R.J. and Stuerzlinger, W. Assessing the Effects of Orientation and Device on (Constrained) 3D Movement Techniques. *Proc. 3DUI*, 43-50.

[29] Tse, E., Hancock, M., and Greenberg, S. Speech-filtered bubble ray: improving target acquisition on display walls. *Proc Multimodal Interfaces 2007*, 307-314.

[30] Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large, high resolution displays. *Proc. UIST'2005*, 33-42.

[31] 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.

[32] 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.

[33] 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.