

An Evaluation of Techniques for Controlling Focus+Context Screens

Mark J. Flider

Brian P. Bailey

Department of Computer Science
University of Illinois at Urbana-Champaign
{mflider, bpbailey}@uiuc.edu

Abstract

We evaluated four techniques for controlling focus+context screens. We compared an egocentric versus exocentric View mixed with whether the display on the focus screen moves in the same (paper mapping) versus opposite (scroll mapping) direction as input force. Our results show that (i) View had little effect, (ii) users almost always allocated attention to the context screen when controlling the display, (iii) scroll mappings enabled a user to perform tasks faster, commit fewer errors, and be more satisfied with the system compared to paper mappings, and (iv) a user can better control focus+context screens when the frame of reference either *does* move or is *perceived* to move in the direction of input force. We discuss these results and recommend how to enable a user to better control focus+context screens.

Key words: Design, Focus+context, Sketching, Two-handed interaction.

1 Introduction

Designers need a larger electronic workspace effective for sketching [3]. To be effective for sketching, we believe that a larger electronic workspace should enable a designer to sketch on a high-resolution screen in an ergonomic and familiar manner, quickly and accurately navigate to different parts of a design, sketch details in context of the broader design, and instantaneously switch between a detailed portion of a design and its context.

To satisfy these requirements, we could use a visualization technique such as zooming [7, 8] or a distortion-based view [12, 17, 23, 25], use a single large screen [15], or augment a physical screen with a virtual screen [16]. While each of these techniques meets some of our requirements, none of them meet all of our requirements. For example, a zooming canvas does not enable a designer to sketch details in context of the broader design and the use of a single large screen provides poor ergonomics [15].

To create a larger electronic workspace that meets our requirements, we developed a focus+context screens system in which we tethered a high-resolution tablet to a large screen. As shown in Figure 1, the tablet provides



Figure 1: A user interacting with our focus+context system. She controls the frame of reference using her non-dominant hand while sketching with her dominant hand.

a logical frame of reference into the content shown on the large screen. In contrast to the focus+context screens system developed by Baudisch [4, 6], we physically separate the focus screen from the context screen to provide a designer with a wider visual angle to see more of the context. We position the large screen just above and behind the tablet to ensure that the tablet does not block the view of the context.

Since our focus screen is tied to, but physically separated from its logical frame of reference, we can scale the focus and context views independently, not possible with Baudisch's focus+context display [6]. Because of this, our system affords two perceptual views: an *exocentric* view, where a designer controls the frame of reference within the context screen, and an *egocentric* view, where the frame of reference remains fixed, and a designer controls the context about it [20]. For example, Photoshop's [2] navigator window uses an exocentric view: when the frame of reference moves in a direction on the context screen, the content on the focus screen moves in the *opposite direction*, causing a discrepancy of motion between the focus and context screens. In an egocentric view, however, when a user moves the content on the con-

text screen in a direction, the content on the focus screen moves in the *same* direction, eliminating the discrepancy of motion.

Consistent with bimanual theory [9, 18, 22], we enable a designer to control a focus+context screens system using a 6DOF input device in her non-dominant (ND) hand while sketching with her dominant hand. In the physical world, a person uses their ND hand to manipulate a frame of reference for their dominant hand, such as positioning a sheet of paper while sketching on it. In contrast, scroll-bars, the traditional method for interacting with large information spaces, scroll content in the *opposite* direction of physical movement. Thus, would a designer control a focus+context screens system most effectively with an egocentric or exocentric view, and a paper or scroll mapping? Furthermore, does the selection of a mapping depend upon the selection of a perceptual view?

While prior work has compared focus+context screens to different display techniques [4], investigated interactions for pan and zoom interfaces [8, 10, 19], and investigated how perceptual views affect navigation [11, 28] and object manipulation [24, 30] in virtual environments, our work is the first to empirically compare the mix of perceptual views with ND hand input mappings for controlling a focus+context screens system. Our results should lead to better techniques for controlling a focus+context screens system.

In our evaluation, we mixed View (ego- vs. exocentric) with Mapping (scrollbar vs. paper) and empirically compared the resulting four configurations. We measured task performance, goal and task errors, shifts of attention between screens, and user satisfaction among 14 users.

Our results show that (i) View had little effect, (ii) users almost always allocated attention to the context screen when controlling the display, (iii) scroll mappings enabled a user to perform tasks 25% faster, commit 70% fewer errors, and be much more satisfied with the system compared to paper mappings, and (iv) a user can better control focus+context screens when the frame of reference either *does* move or is *perceived* to move in the direction of input force. We discuss these results and recommend how to enable a user to better control focus+context screens.

2 Related Work

Our work differs from previous work in that we are using a focus+context screens system to provide a larger electronic workspace effective for sketching, using a ND hand input device for control, and comparing different perceptual views and input configurations for the system.

2.1 Larger Electronic Workspaces

To provide a virtually larger workspace, we could use a zooming canvas [7, 8], but this does not enable a designer to view details in context of the broader design [4, 6]. Distortion-based views [12, 17, 23] have been effective for visualizing large amounts of information, but the inherent spatial distortion would be awkward and inappropriate for visual design tasks such as sketching.

To provide a physically larger workspace, we could use a large digital desk. A large digital desk, however, typically has poor resolution, suffers from parallax effects, is not well-suited for users under 5'6", is not ergonomic, and users do not prefer it to a tablet for sketching [15].

While augmenting a physical screen with a virtual screen is useful for managing application windows [16], the interactions required to manage the screens would be overly awkward for electronic sketching. Previous systems which tied smaller displays to a large display [13, 26, 29] either did not support sketching or no evaluations of control techniques were performed.

Baudisch et al. surrounded a high-resolution small screen with a lower-resolution large screen [4, 5, 6]. Since the screens lie in the same plane, however, a designer would need to sit very close to the screens to sketch, overly limiting her visual angle to see design context. This work also did not evaluate techniques for controlling the display, which is a focus of our work.

2.2 Control Techniques for Focus+Context Screens

In virtual environments, researchers have investigated the use of ego- vs. exocentric views for navigation [11, 28] and object manipulation [24, 30]. Our work also investigates ego- vs. exocentric views, but from a perspective of how those views interact with mappings from a ND hand input device for controlling focus+context screens.

Photoshop's navigator [2] implements a focus+context screen (and in our terminology, an exocentric view with a scroll mapping) by providing an additional, smaller window. This competes for screen space, a precious commodity to designers [9]. Unlike our system, it does not allow the context view itself to be manipulated and there is ostensibly little empirical justification for this particular configuration of view and mapping. Most importantly, in our system, we enable a user to control the focus+context screens system using her ND hand.

Researchers have investigated the use of scroll-based pan and zoom interactions to enable a user to better navigate large information spaces [8, 10, 19]. In our work, we investigate how different mappings of ND-hand input for pan and zoom interfaces interact with perceptual views for controlling a focus+context screens system.

Buxton has shown that two-handed interaction can enable faster performance for drawing and other tasks

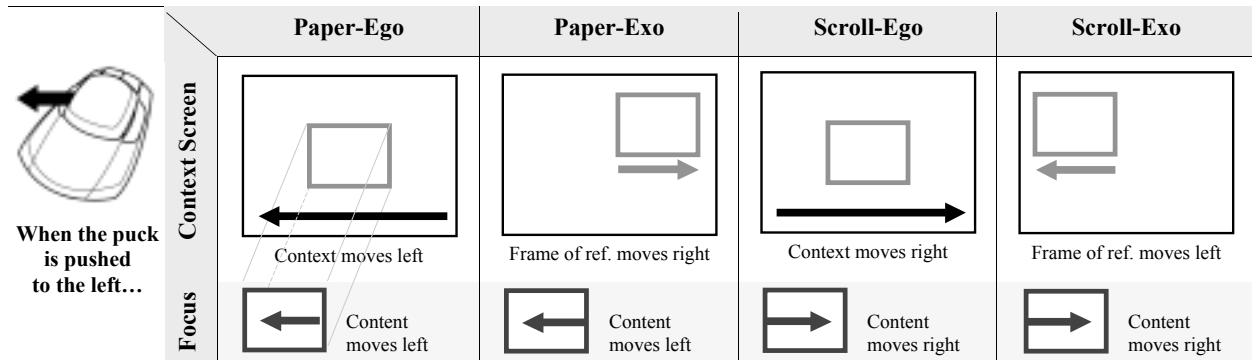


Figure 2: What a user sees on the focus and context screens when they push the puck left in each input configuration. On the context screen, the black arrows indicate the context (background) moving, while the lighter-colored arrows indicate the logical frame of reference moving. We encourage the reader to inspect this diagram carefully, and refer back to this diagram often as they read the paper.

[9, 22]. In our work, we investigate two-handed interaction for controlling a focus+context screens system.

3 Focus+Context Screens System

To provide focus+context, we tethered a high-resolution tablet to a lower-resolution large screen where the tablet, i.e., the focus screen, provides a frame of reference into the context screen (see Figure 1). With our focus+context screens system, a designer can use the tablet for sketching and can use the larger screen to see design context.

For the focus screen, we use a Wacom Cintiq 18SX graphics display tablet. For the context screen, we use a NEC LT260 high-lumen projector to project a large screen on the wall. We position the focus screen on a desk a few feet away from the projected screen and horizontally center the focus screen with respect to the projected screen.

Our approach is similar to Baudisch’s [6], except that we physically separate the two displays to afford a wider visual angle of the context screen, which is important for improving spatial performance [27]. To drive the focus and projected screen, we use a Macintosh dual processor machine with a dual head graphics card. We wrote the software for our focus+context screens system and evaluation using Objective-C and Cocoa under Mac OS X.

3.1 Frame of Reference

On the context screen, we draw a frame of reference to enable a designer to quickly identify their location in the context. The focus screen shows details of the design lying within the frame of reference. As shown in Figure 1, we draw the frame of reference using a thin red rectangle (which appears dark grey if printed in greyscale).

By controlling the frame of reference, a designer can control what information is shown on the focus screen.

When a designer controls the frame of reference, however, our system could either move the frame of reference across the context view or could move the context view around a fixed frame of reference. Thus, focus+context screens afford at least two different perceptual views:

- *An egocentric view.* The *frame of reference* remains fixed relative to the context screen and the context is panned and zoomed around it. For example, imagine viewing content on a piece of paper by moving the paper left and right and towards and away from you. Because it is fixed relative to the context screen, the frame of reference is always surrounded by a consistent amount of context space in each direction.
- *An exocentric view.* The *context* remains fixed while the frame of reference pans and zooms around it. For example, imagine viewing a large wall of information by walking left and right and towards and away from it. Since the frame of reference can be positioned anywhere within the context screen, the amount of surrounding context space can vary in each direction.

3.2 Input Device and Mappings

Consistent with theories of bimanual input [18], we enable a user to control the information shown on the focus screen using her non-dominant (ND) hand while sketching using a stylus in her dominant hand as shown in Figure 1. Buxton has shown that a two-handed interaction consistent with bimanual theory can provide better performance on drawing tasks than a single-handed interaction [9, 22]. His work, however, did not investigate controlling focus+context screens, which our work does.

For our input device, we selected 3D Connexion’s Cadman [1], a 6DOF isometric input device, which we refer

to as a “puck.” We chose this device because Jacob and Sibert recommend that an input device used for panning and zooming should enable a user to perform these tasks in parallel [21], which our puck does.

When a designer uses the puck to control the focus+context display, a significant design challenge is deciding how the display should respond. For example, when a designer pushes left on the puck, should content on the focus screen appear to move left to right, or right to left? We identify two methods to map user input to a change in the display:

- *Paper mapping.* When a user pushes the puck in one direction, the content on the focus screen moves in the *same* direction. For example, if a user pushes left on the puck, the content on the focus screen moves from right to left. For zooming, when a user pushes down – towards the floor – on the puck, the display on the focus screen zooms out, shrinking the content on the focus screen. We call this mapping ‘paper’ because it mirrors interactions with a piece of paper; pushing the paper left moves it left, and pushing it away shrinks it in your field of view.
- *Scroll mapping.* When a user pushes the puck in the direction of interest, the content on the focus screen moves in the *opposite* direction. For example, if a user pushes left on the puck, the content on the focus screen appears to move from left to right. For zooming, when a user pushes down on the puck, the display on the focus screen zooms in, enlarging the content on the focus screen. We call this mapping ‘scroll’ because most scrollbar behave in this way; sliding a scrollbar in a direction of interest moves a document in the opposite direction.

To understand how Mapping and View affect task performance, error rate, and satisfaction, we conducted a user evaluation that compared the four input configurations obtained by mixing Mapping with View: *paper-ego*, *paper-exo*, *scroll-ego*, and *scroll-exo*. In Figure 2, we show how pushing left on the puck would change the display that a user sees on the focus screen for each input configuration. Other directions and zooms are analogous.

4 User Evaluation

We designed the evaluation to answer five questions:

- How does View affect task performance, error rate, and satisfaction?
- How does Mapping affect task performance, error rate, and satisfaction?

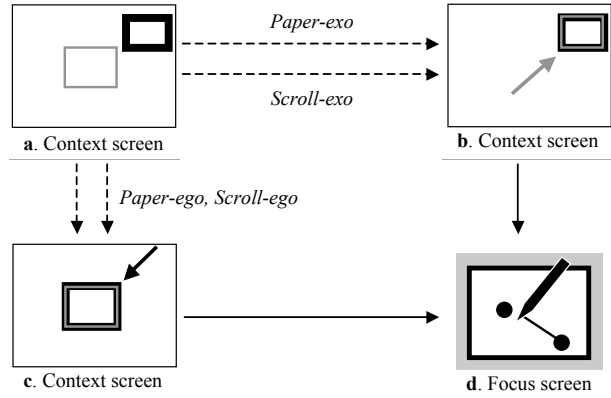


Figure 3: In a task (a), a user either moved the thin frame of reference to the thick, black, target rectangle (b), or moved the target rectangle to the frame of reference (c). A user performed each movement with paper and scroll mappings. Once the frame of reference was within the target’s border, two dots appeared on the focus screen (d).

- How do View and Mapping interact to affect task performance, error rate, and satisfaction?
- How do users divide their attention between the context and focus screens?
- How useful do users believe that our focus+context screens system would be for different types of design and work tasks?

4.1 Hardware and Software

We used the focus+context system described earlier. To set the sensitivity of the puck for the evaluation, we ran a small pilot study and set it to most users’ satisfaction.

4.2 Subjects

14 subjects (4 female) participated in the evaluation. Subjects were undergraduate or graduate students in computer science or psychology and were 18-30 years of age. Twelve users were right-handed and for the two users who were not, we moved the puck to the right side of the focus screen. Since the puck is symmetric, this had no impact on the tasks. We did not compensate participants.

4.3 Tasks

We wanted to design a task that would be appropriate for two-handed input, involve some amount of drawing with the dominant hand, and allow the user to switch attention between screens. Based on these goals, we developed a ‘connect the dots’ task, which was inspired by the drawing tasks used by Buxton [22].

At the beginning of each task, a user would see the frame of reference and a black rectangle (target) with a thick border on the context screen. The target rectangle

was a different size and in a different position relative to the frame of reference. Using the puck with the ND hand to pan and zoom, a user's task was to control the display and position the frame of reference within the border of the target rectangle. Depending on the configuration, the user would move either the frame of reference to the target rectangle (exocentric) or the target rectangle to the frame of reference (egocentric). Each movement was performed with both a paper and scroll mapping (see Figure 3). When a user controlled the display to where some of the black border of the target rectangle was visible along each edge of the focus screen, two dots would appear on the focus screen. The user would then use the stylus in her dominant hand to connect the dots, and click any button on the puck to end the task. We developed eight of these tasks where each task differed in the initial size and position of the target rectangle.

Because the border on the target rectangle was much wider than the frame of reference, a user would see its black border inside the focus screen when the frame of reference neared the target rectangle. Thus, a user could allocate her visual attention to either screen to perform the task. We expected users to look at the context screen while performing initial large movements, then switch attention to the focus screen to make finer adjustments.

4.4 Experimental Design

We used a 2 View (egocentric vs. exocentric) x 2 Mapping (paper vs. scroll) x 8 Task within-subjects design. We used a balanced Latin Square to define orderings of the configurations and randomly assigned users to them.

4.5 Procedure

When a subject arrived at the lab, we went through an informed consent process. We asked a user to fill out a questionnaire for demographic and handedness information. The experimenter then provided an overview of and demoed the focus+context system. The experimenter set up the first input configuration for the user and gave the user about five minutes of "free play" to become familiar with the configuration.

After a five minute warm-up, the experimenter described the 'connect the dots' task and gave the subject four practice tasks. During this time, the experimenter answered any questions that the subject had. The subject then performed each of the eight connect the dots tasks in a randomized order. The experimenter instructed the subject to perform the tasks as quickly as possible while maintaining accuracy on the tasks.

After completing each set of 8 tasks, the subject filled out a questionnaire while the experimenter arranged the next configuration. This process was then repeated three more times. The subject had five minutes of free play and

four practice trials for each input configuration.

After the last configuration, the subject filled out a post evaluation questionnaire. We video taped each session for later analysis. Each session lasted about 45 minutes.

4.6 Measurements

In the evaluation, we measured task performance, error rate, shifts of visual attention, and user satisfaction.

Task Performance

We logged performance in two parts: the time to control the display such that the frame of reference lies within the borders of the target rectangle (movement time) and the time from the first movement to the last pen event (total).

Error Rate

To measure error rate, we logged the following three data for each event generated by the puck:

- *Time*. When the event occurred.
- *Distance*. The distance between the centers of the frame of reference and target rectangle.
- *Zoom ratio*. The ratio of the height of the frame of reference to that of the target rectangle.

By logging both pan and zoom errors independently, we could analyze whether the input configuration affected these differently. We divided errors into two categories:

- *Goal error*. When a user's initial control action caused a change in the display opposite the user's intent. For example, a user pushed left on the puck to move the display left, but should have pushed right. We assume a user's goal was to always minimize distance to target.
- *Task error*. When a user began a control action in the right direction, but without a pause began moving away from the target rectangle, e.g., overshoot or undershot the target rectangle, or zoomed in or out too much.

By plotting *distance* and *zoom ratio* against *time*, we could inspect the shape of the resulting piecewise linear curve to identify the type and number of errors.

Shifts in Visual Attention

We measured shifts in visual attention by analyzing the video recordings and counting how often a user switched from the focus screen to the context screen and vice versa.

		Display Configuration Avg. (s.d.)				Analysis F=(1,13)		
		Scroll- ego	Scroll- exo	Paper- ego	Paper- exo	Map	View	View * Map
Movement Time		9.30 (3.08)	9.32 (2.10)	12.61 (3.20)	12.45 (2.86)	F=61.2 p<.001	F=.01 p<.93	F=.03 p<.87
Goal Error	Zoom	.46 (.38)	.38 (.25)	.71 (.39)	.70 (.21)	F=10.5 p<.01	F=.46 p<.51	F=.42 p<.53
	Move	.47 (.59)	.38 (.26)	1.46 (.69)	1.31 (.58)	F=36.0 p<.001	F=1.13 p<.31	F=.06 p<.80
Task Error	Zoom	.79 (.53)	.90 (.57)	.83 (.42)	.72 (.36)	F=.50 p<.49	F=.00 p<1.0	F<1.74 p<.21
	Move	1.6 (.90)	1.71 (.75)	2.05 (1.71)	1.78 (.79)	F=1.74 p<.21	F=.24 p<.63	F=.80 p<.39

Table 1: Task performance and error data.

User Satisfaction

After each configuration, we asked users to rate on a 5-point scale how much they disagreed or agreed with:

- You could perform the tasks easily.
- The input configuration was easy to use.
- The input configuration was easy to learn.
- The input configuration was well-suited to the tasks.
- Switching your attention between screens was easy.
- Controlling the display felt natural when looking at the context screen.
- Controlling the display felt natural when looking at the focus screen.

After the evaluation, we asked a user to rank the input configurations from best to worst and comment on why they gave the rankings. We also asked users to identify work tasks they felt our system would be most useful for.

5 Results

Because ‘Task’ did not produce meaningful effects, we collapsed the data across tasks and analyzed the results as a 2 (Map) x 2 (View) design.

5.1 Task Performance

Because movement time dominated total task time, we only report that information here.

As summarized in Table 1 and shown in Figure 4, the mapping used (scroll vs. paper) had a main effect on movement time ($F(1,13)=61.2$, $p<.001$). Post hoc analysis using paired t-tests showed that when a user was in either configuration with a scroll mapping, she performed better than when she was in either configuration with a paper mapping (scroll-ego vs. paper-ego $t(13)=4.56$, $p<.001$; scroll-ego vs. paper-exo $t(13)=4.35$, $p<.001$;

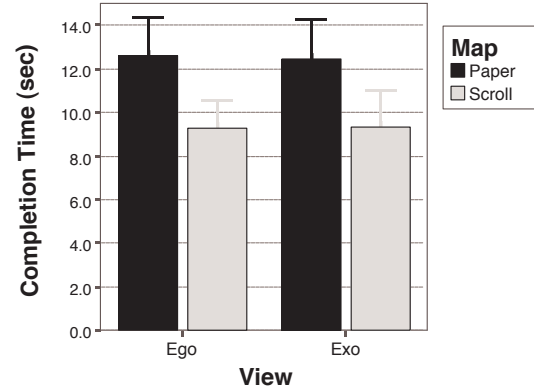


Figure 4: Task completion time for each configuration.

scroll-exo vs. paper-ego $t(13)=3.55$, $p<.004$; scroll-exo vs. paper-exo $t(13)=5.13$, $p<.001$). View did not have a main effect and there were no interactions.

On average, a user completed the tasks about 25% faster when using a scroll mapping than when using a paper mapping, regardless of View. Because our focus+context screens system is intended for daily use, the performance gain of using a scroll mapping is highly significant.

5.2 Error Rate

As summarized in Table 1 and shown in Figures 5 and 6, the mapping used (scroll vs. paper) had a main effect on goal errors for both zooming ($F(1,13)=10.5$, $p<.01$) and movement ($F(1,13)=36.0$, $p<.001$). For zooming goal errors, post hoc analysis using paired t-tests showed that when a user was in the scroll-exo configuration, the user committed fewer zooming goal errors than when in the paper-ego ($t(13)=2.63$, $p<.02$) or paper-exo ($t(13)=5.83$, $p<.001$) configurations. Although the scroll-ego bar is shorter than both bars for the paper mappings in Figure 5, the differences were not statistically significant.

For movement goal errors, post hoc analysis using paired t-tests showed that when a user was in either configuration with a scroll mapping, she made fewer errors than when in either configuration with a paper mapping (scroll-ego vs. paper-ego $t(13)=4.17$, $p<.001$; scroll-ego vs. paper-exo $t(13)=4.41$, $p<.001$; scroll-exo vs. paper-ego $t(13)=5.37$, $p<.001$; scroll-exo vs. paper-exo $t(13)=5.87$, $p<.001$).

The mapping used did not affect task errors, which we believe was because tasks errors depended more on the physical control of the puck and less on cognition. For both error categories, the perceptual view used did not affect error rate and there were no interactions. These results are consistent with the task performance results.

On average, a user made about 41% (zooming goal) and 70% (movement goal) fewer errors when using a

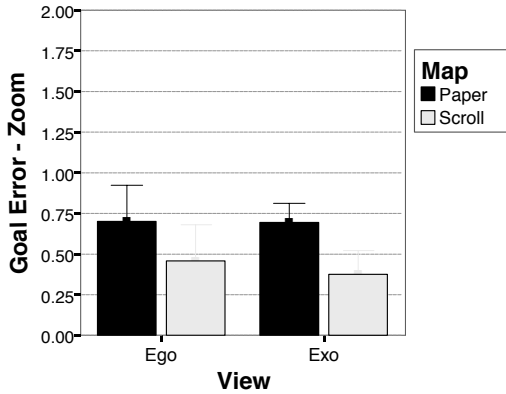


Figure 5: Zooming goal errors for each configuration.

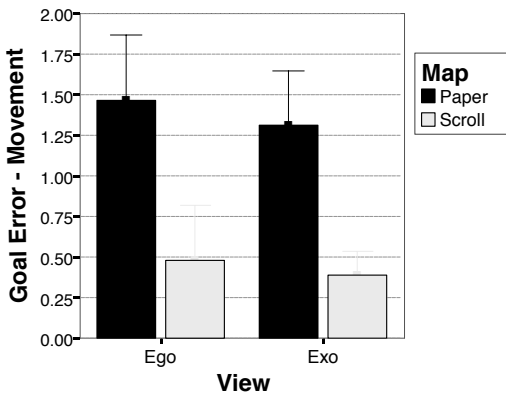


Figure 6: Moving goal errors for each configuration.

scroll mapping than when using a paper mapping, regardless of the perceptual view used. Although the absolute difference in error rate seems small (about one error), one should note that the tasks were very short in duration, each about twelve seconds. Thus, if a user frequently controlled the display, the scroll mapping would save a user about one error every 12 seconds of use.

5.3 Shifts of Attention

Although a few users shifted their attention to the focus screen to perform finer control of the display, most users almost always looked at the context screen when controlling the display during tasks. Only after the dots appeared on the focus screen did a user typically switch her attention back to that screen. This result suggests that when a designer interacts with a focus+context screens system, she will likely focus on the focus screen when performing local navigation tasks and shift her attention to the context to perform larger navigation tasks.

5.4 User Satisfaction

From the task questionnaires, summarized in Table 2, the mapping used had a main effect on a user’s response for all but the last two questions. Because users almost al-

The display configuration...	Display Configuration Avg, (s.d.)				Analysis F=(1,13)		
	Scroll-ego	Scroll-exo	Paper-ego	Paper-exo	Map	View	View * Map
Enabled you to perform the tasks easily	4.21 (.58)	3.64 (1.01)	2.86 (1.17)	2.5 (1.02)	F=14.7 p<.002	F=3.5 p<.08	F=.97 p<.53
Was easy to use	4.07 (.73)	3.79 (1.19)	2.71 (1.14)	2.71 (1.07)	F=13.9 p<.003	F=.20 p<.66	F=1.0 p<.34
Was easy to learn	4.43 (.65)	4.07 (1.07)	3.29 (.73)	2.86 (.86)	F=17.1 p<.001	F=3.1 p<.10	F=.04 p<.84
Was well-suited to tasks	4.14 (.66)	3.86 (.86)	2.93 (1.21)	2.50 (1.02)	F=16.4 p<.001	F=4.1 p<.07	F=.21 p<.66
Felt natural while looking at large screen	4.64 (.50)	4.43 (.65)	4.00 (.96)	3.71 (.99)	F=8.6 p<.01	F=2.6 p<.13	F=.04 p<.84
Felt natural while looking at tablet	3.15 (.90)	3.29 (.99)	2.85 (1.28)	2.93 (1.33)	F=1.4 p<.25	F=.81 p<.39	F=3.6 p<.08
Enabled you to switch between them	3.77 (.93)	3.79 (.89)	3.92 (.95)	3.71 (1.14)	F=.14 p<.71	F=.27 p<.61	F=.11 p<.75

Table 2: Responses from task questionnaires.

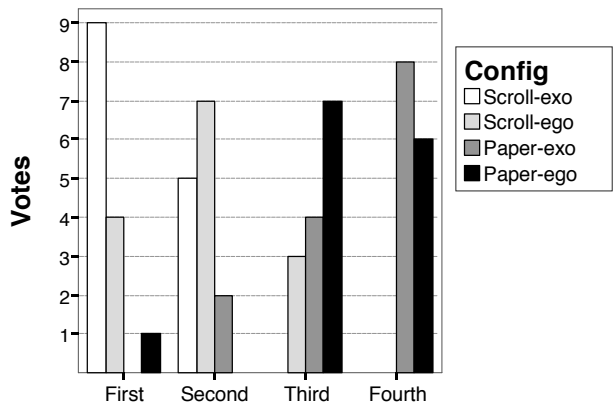


Figure 7: Qualitative ranking of input configurations.

ways focused their visual attention on the context display, the last two questions in Table 2 did not apply well, and thus, we did not expect or find any differences among the responses. The View used did not have a main effect and there were no interactions for any of the responses.

On average, users rated the use of the scroll mappings in our focus+context screens system a four or higher on a five-point scale. We believe this shows that users felt that our system was reasonably easy to use and learn in these configurations. Overall, the user satisfaction results are very consistent with the task performance and error data.

From the post evaluation questionnaire, as shown in Figure 7, users ranked scroll-ego and scroll-exo better than paper-ego and paper-exo. These results are consistent with the task performance and error measurements. As shown in Table 3, users identified graphic design,

Question	Avg (s.d.) - votes
Satisfied with this configuration	4.07 (0.83)
A similar setup would be useful in your work	3.71 (0.83)
Using both displays felt natural	3.71 (0.61)
Using both hands felt natural	3.86 (0.86)
Useful for graphic design	13 / 14
Useful for sketching	11 / 14
Useful for web design	10 / 14
Useful for video editing	6 / 14

Table 3: Responses from post-evaluation questionnaires. Users answered each question according to their most-preferred input configuration.

web design, video editing, and especially sketching as the tasks that our system would best support. These results are very encouraging, since these are the type of tasks for which we designed our system.

6 Lessons Learned and Recommendations

From the evaluation results, we learned that:

- *The perceptual view used (ego/exo) did not affect task performance, error rate, or satisfaction.* As shown in Figures 4-6 and Table 2, the perceptual view had little effect in either the paper or scroll mapping configurations. These figures and the table also show, however, that the mapping from a user’s direction of input force to a change in the display dominates the impact on task performance, error rate, and satisfaction.
- *Users performed tasks about 25% faster with the scroll mapping than with the paper mapping.* Because a force applied to the puck caused a change on the display consistent with a user’s expectation, the scroll mapping bridged the gulf of execution and enabled a user to perform the tasks faster. Although the absolute difference in performance was about four seconds, this performance difference is highly significant when extrapolated over many hours of use.
- *Using the scroll mapping, users committed fewer goal errors for both panning (70% fewer) and zooming (41% fewer) than with the paper mapping.* Although our results show that the absolute difference in goal errors was about one per task, each task was only about 12 seconds, so this small difference may be highly significant if extrapolated over many hours of use, for which our system is intended.

- *Users were more satisfied with the scroll mapping for performing tasks than with the paper mapping.* As shown in Table 2, users consistently preferred the scroll mapping, since the paper mapping often caused a change in display contrary to a user’s expectation. For example, in the questionnaires regarding a paper mapping, several users wrote “this configuration was very counter-intuitive.”
- *Users performed tasks faster, committed fewer errors, and were more satisfied controlling the system when the frame of reference **did** move or was **perceived** to move in the direction of input force.* In the scroll-ego configuration, when a user applied a directional force to the puck, the content on the context screen moved *opposite* that direction (see Figure 2). Although we initially believed that this configuration would cause worse performance and be rated poorly, we were surprised to discover that in this configuration users *perceived* that the frame of reference was actually moving in the *same* direction as input force. Thus, the red rectangle gave users a “perceptual handle” to make sense of the seemingly opposing motions. In fact, their perception was so strong that after being debriefed on the evaluation, several users exclaimed, “I could have sworn that the red rectangle was actually moving!” Our results show that *perceiving* the frame of reference as moving is almost as effective as *actually* moving it.

Regarding this last lesson, because bimanual input theory suggests a user would find a paper-ego configuration more natural than a scroll-ego configuration, we were surprised to find that users actually performed better and were more satisfied with the scroll-ego configuration. The reason is that most users reliably perceived the thin red rectangle (logical frame of reference) on the context screen as the element that actually moved, even though it was the context that was actually moving around the logical frame of reference. This effect is a perceptual phenomenon known as induced motion [14]. To test whether it was the bright red color of the rectangle that was causing the induced motion effect, we switched the color for half of our users from bright red to a much less luminescent dark blue. However, analysis of the data using color as a factor showed no performance or satisfaction differences.

Before identifying our final recommendations on how to enable a user to better control a focus+context screens system, let us review our results heretofore. Our results show that a user can control a focus+context screens system effectively in a scroll-ego configuration and slightly more effectively in a scroll-exo configuration. In the

scroll-ego configuration, a user controls the system effectively because he *perceives* the logical frame of reference moving in the direction of input force. In the scroll-exo configuration, a user controls the system effectively because the logical frame of reference actually *does* move in the direction of input force.

Although selecting the latter configuration for a focus+context screens system is tempting, a significant problem still remains: in a scroll-exo configuration, although the logical frame of reference moves in the direction of input force, the motion on the focus screen moves in the opposite direction (see Figure 2). An important contribution of our study lies in the discovery of induced motion in the scroll-ego configuration. This discovery enables us to recommend three alternative methods to enable a user to control the system effectively *and* resolve the discrepancy of motion between the two screens:

1. *Select a scroll-exo configuration uniformly and add a perceptual handle to the focus screen.* By placing a perceptual handle such as a cross-hair in the center of the focus screen, a user should latch onto the perceptual handle and perceive it as moving in the same direction as the logical frame of reference on the context screen. Thus, because we resolve the discrepancy of motion between the screens in a way consistent with our experimental results, a user should be able to control the system effectively when attending to either screen.
2. *Select a scroll-ego configuration uniformly and add a perceptual handle to the focus screen.* Again, by placing a perceptual handle such as a cross-hair on the focus screen, we can cause a user to perceive it moving opposite to the direction of content on the focus screen. In this configuration, a user would perceive the logical frame of reference on the context screen *and* the perceptual handle on the focus screen moving in the *same* direction as input force. Thus, in this configuration, we resolve the discrepancy in motion by *inducing the same visual effect on both screens*, allowing a user to control the system effectively when attending to either the context or the focus screen.
3. *Select a scroll mapping when a user attends to the context screen and select a paper mapping when the user attends to the focus screen.* If implemented using a system that detects visual gaze, when a user controls the display while attending to the context screen, the frame of reference would move in the direction of force, which our results show to be the best configuration. When the user shifts attention to the focus screen, however, the system should switch

to a paper mapping. With this mapping, the content on the focus screen appears to move in the same direction as the input force (see Figure 2). As our results show, this is required for a user to control the system effectively. The advantage of this approach over the previous two approaches is that it does not rely on induced motion through the addition of a perceptual handle to be effective.

For future work, we are designing another controlled study to determine how a user expects the logical frame of reference to move on the context screen, when the user makes fine control movements (nudges) while attending to the focus screen. This should enable us to select between our first and second recommendations above. In the study, we would also investigate the strength of the induced motion on the focus screen; if the effect is weakened – perhaps due to the smaller screen – we may implement and further investigate our third recommendation.

7 Conclusion

To develop a larger electronic workspace effective for sketching, we developed a focus+context screens system. To investigate effective techniques for controlling such a system, we compared an egocentric and exocentric view mixed with a paper and scroll mapping.

Our results show that (i) perceptual view had little effect, (ii) users almost always allocated attention to the context screen when controlling the display, (iii) scroll mappings enabled a user to perform tasks 25% faster, commit 70% fewer errors, and be much more satisfied with the system compared to paper mappings, and (iv) a user can better control focus+context screens when the frame of reference either *does* move or is *perceived* to move in the direction of input force.

Our results and recommendations should help systems designers develop more effective techniques for controlling focus+context screens systems.

References

- [1] 3Dconnexion CadMan.
<http://www.3dconnexion.com/cadman.htm>.
- [2] Adobe Photoshop.
<http://www.adobe.com/photoshop>.
- [3] B.P. Bailey and J.A. Konstan. Are informal tools better? Comparing DEMAIS, pencil and paper, and authorware for early multimedia design. In *CHI*, pages 313–320, 2003.
- [4] P. Baudisch, D. DeCarlo, A. Duchowski, and B. Geisler. Focusing on the essential: Considering attention in display design. *CACM*, 46(3):60–66, 2003.

- [5] P. Baudisch, N. Good, V. Bellotti, and P. Schraedley. Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. In *CHI 2002*, Minneapolis, Minnesota, April 2002.
- [6] P. Baudisch, N. Good, and P. Stewart. Focus plus context screens: Combining display technology with visualization techniques. In *Symposium on User Interface Software and Technology (UIST 2001)*, Orlando, Florida, November 2001.
- [7] B.B. Bederson, J.D. Hollan, K. Perlin, J. Meyer, D. Bacon, and G.W. Furnas. Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *Journal of Visual Languages and Computing*, 7:3–31, 1996.
- [8] F. Bourgeois and Y. Guiard. Multiscale pointing: Facilitating pan-zoom coordination. In *CHI*, pages 758–9, 2002.
- [9] W. Buxton and B. Myers. A study in two-handed input. In *CHI*, pages 321–326, 1986.
- [10] D. Byrd. A scrollbar-based visualization for document navigation. In *Proceedings of Digital Libraries*, pages 122–9, 1999.
- [11] R.P. Darken and J.L. Sibert. Navigating in large virtual worlds. *IJHCI*, 8(1):49–72, 1996.
- [12] D. Dee-Lucas and J. Larkin. Learning from electronic texts: Effects of interactive overview for information access. *Cognition and Instruction*, 13(3):431–468, 1995.
- [13] W. Donelson. Spatial management of information. In *Proceeding of SIGGRAPH*, pages 203–9, 1978.
- [14] K. Duncker. Uber induzierte bewegung: Ein beitrag zur theorie optisch wahrgenommener bewegung (Concerning induced movement). In W.D. Ellis, editor, *Source Book of Gestalt Psychology*, pages 180–259. Routledge & Kegan Paul, London, 1938.
- [15] A. Elliott and M.A. Hearst. A comparison of the affordances of a digital desk and tablet for architectural image tasks. *IJHCS*, 56(2):173–197, 2002.
- [16] S. Feiner and A. Shamash. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. In *UIST*, pages 9–17, 1991.
- [17] G.W. Furnas. Generalized fisheye views. In *CHI*, pages 16–23, 1986.
- [18] Y. Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Motor Behavior*, 19(4):486–517, 1987.
- [19] L. Holmquist. Focus+context visualization with flip zooming and the zoom browser. In *CHI Abstracts*, pages 263–264, 1997.
- [20] I.P. Howard. Spatial vision within egocentric and exocentric frames of references. In S.R. Ellis, editor, *Pictorial Communication in Virtual and Real Environments*. Taylor and Francis, London, 1991.
- [21] R. Jacob and L. E. Sibert. The perceptual structure of multidimensional input device selection. In *CHI*, pages 211–8, 1992.
- [22] P. Kabbash, W. Buxton, and A. Sellen. Two-handed input in a compound task. In *CHI*, pages 417–423, 1994.
- [23] J. D. Mackinlay, G. G. Robertson, and S. K. Card. The perspective wall: Detail and context smoothly integrated. In *Proc. of CHI-91*, pages 173–179, New Orleans, LA, 1991.
- [24] I. Poupyrev, T. Ichikawa, S. Weghorst, and M. Billinghamurst. Egocentric object manipulation in virtual environments: Empirical evaluation of interaction techniques. *Computer Graphics Forum*, 17(3):41–58, 1998.
- [25] G. Robertson and J.D. Mackinlay. The document lens. In *UIST*, pages 101–108, 1993.
- [26] N.A. Streitz, J. Geüüler, T. Holmer, S. Konomi, C. Müller-Tomfelde, W. Reischl, P. Rexroth, P. Seitz, and R. Steinmetz. i-land: An interactive landscape for creativity and innovation. In *CHI*, pages 120–127, 1999.
- [27] D.S. Tan, D. Gergle, P. Scupelli, and R. Pausch. With similar visual angles, larger displays improve spatial performance. In *CHI*, pages 217–24, 2003.
- [28] D.S. Tan, G.G. Robertson, and M. Czerwinski. Exploring 3d navigation: Combining speed-coupled flying with orbiting. In *CHI*, 2001.
- [29] M. Tani, M. Horita, K. Yamaashi, K. Tanikoshi, and M. Futukawa. Courtyard: Integrating shared overview on a large screen and per-user detail on individual screens. In *CHI*, pages 44–50, 1994.
- [30] S. Zhai and P. Milgram. Human performance evaluation of manipulation schemes in virtual environments. In *Proceedings of VRAIS*, pages 155–6, 1993.