

High-Precision Magnification Lenses

Caroline Appert^{1,2}
appert@lri.fr

Olivier Chapuis^{1,2}
chapuis@lri.fr

Emmanuel Pietriga^{2,1}
emmanuel.pietriga@inria.fr

¹LRI - Univ. Paris-Sud & CNRS
Orsay, France

²INRIA
Orsay, France

ABSTRACT

Focus+context interfaces provide in-place magnification of a region of the display, smoothly integrating the focus of attention into its surroundings. Two representations of the data exist simultaneously at two different scales, providing an alternative to classical pan & zoom for navigating multi-scale interfaces. For many practical applications however, the magnification range of focus+context techniques is too limited. This paper addresses this limitation by exploring the *quantization* problem: the mismatch between visual and motor precision in the magnified region. We introduce three new interaction techniques that solve this problem by integrating fast navigation and high-precision interaction in the magnified region. *Speed* couples precision to navigation speed. *Key* and *Ring* use a discrete switch between precision levels, the former using a keyboard modifier, the latter by decoupling the cursor from the lens' center. We report on three experiments showing that our techniques make interacting with lenses easier while increasing the range of practical magnification factors, and that performance can be further improved by integrating speed-dependent visual behaviors.

Author Keywords

Focus+Context, Lenses, Quantization, Navigation, Selection

ACM Classification Keywords

H. Information Systems H.5 Information Interfaces and Presentation H.5.2 User Interfaces (H.1.2, I.3.6)

INTRODUCTION

Although display technologies continue to increase in size and resolution, datasets are increasing even faster. Scientific data, e.g., telescope images and microscope views of the brain, and generated data, e.g., network visualizations, geographical information systems and digital libraries, are too big to be displayed in their entirety, even on very large wall-sized displays. In Google Maps, the ratio between extreme scales is about 250,000. Vast gigapixel images, such as the 400,000-pixel wide image of the inner-part of our galaxy from the Spitzer telescope also require huge scale

factors between a full overview and the most detailed zoom. Users do not necessarily need to navigate through the entire scale range at one given time, but still, they need interaction techniques that will allow them to fluidly navigate between focused and contextual views of large datasets. Such techniques are typically based on the following interface schemes [8]: overview + detail, zooming, focus + context; none of which offers an ideal solution. The task determines which technique is most appropriate, taking scale range, the nature of the representation, input device, available screen real-estate, and of course, the user's preferences, into account.

This paper introduces techniques designed to improve lens-based focus+context interfaces. Our goals are to extend the range of practical magnification factors, which is currently very limited, and to make low-level interactions easier. For the sake of clarity, we illustrate all of our techniques with one common type of lens: constrained magnification lenses [4, 18, 19]. However, our improvements are generic and apply to all types of lenses. They can also be adapted to other focus+context interfaces, including hyperbolic trees [16] and stretchable rubber sheets [20].

QUANTIZATION IN FOCUS+CONTEXT INTERFACES

Constrained lenses provide in-place magnification of a bounded region of the representation (Figure 1-a). The *focus* is integrated in the *context*, leaving a significant part of the latter unchanged. Typical examples of such lenses include magnifying glasses and many distortion-oriented techniques

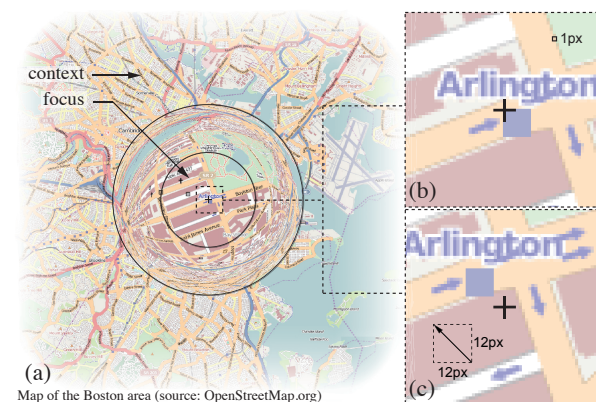


Figure 1. (a) In-place magnification by a factor of 12; (b) center of magnified region with cursor in the middle (detail); (c) same region after moving the lens by one pixel both South and East.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10 – 15, 2010, Atlanta, Georgia, USA

Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.




Figure 2. Space-scale diagram of possible locations for lens center (each ray corresponds to one pixel in context space).

such as the so-called graphical fisheyes. Early implementations of magnification techniques only magnified the pixels of the context by duplicating them without adding more detail, thus severely limiting the range of useful magnification factors (up to 4x). Newer implementations [4, 18] do provide more detail as magnification increases. Theoretically, this means that any magnification factor can be applied, if relevant data is available. In practice, this is not the case as another problem arises that gets worse as magnification increases: *quantization*.

Lenses are most often coupled with the cursor and centered on it. The cursor, and thus the lens, are operated at context scale. This allows for fast repositioning of the lens in the information space, since moving the input device by one unit makes the lens move by one pixel at context scale. However, this also means that when moving the input device by one unit (dot), the representation in the magnified region is offset by MM pixels, where MM is the focus’ magnification factor. This means that only one pixel every MM pixels can fall below the cursor in the magnified region. In other words some pixels are unreachable, as visual space has been enlarged in the focus region but motor space has not.

This problem is illustrated in Figure 1: between (b) and (c), the lens has moved by 1 unit of the input device, corresponding to 1 pixel in the context, but the magnified region is offset by 12 pixels. Objects can thus be difficult or even impossible to select; even if their visual size is above what is usually considered a small target (less than 5 pixels). The square representing Arlington station in Figure 1 is 9-pixel wide, yet its motor size is only 1 pixel.

Figure 2 illustrates the problem with a space-scale diagram [11]: the center of the lens can only be located on a pixel in the focus window that is *aligned* – on the same ray in the space-scale diagram – with a pixel in the context window. Since the focus window is MM^2 larger than the context window, and since the cursor is located at the lens’ center, only one out of MM^2 pixels can be selected. Figure 2 shows that as MM increases, more pixels become unreachable.

Beyond the general problem of pixel-precise selection in the magnified region, quantization also hinders focus targeting, i.e., the action that consists in positioning the lens on the object of interest [12, 18]. This action gets harder as the magnification factor increases, even becoming impossible at extreme magnification factors.

This quantization problem has limited the range of magnification factors that can be used in practice; the upper limit reported in the literature rarely exceeds 8x, a value relatively low compared to the ranges of scale encountered in the information spaces mentioned earlier.

In this paper, we introduce techniques that make it possible to perform both fast navigation for focus targeting and high-precision selection in the focus region in a seamless manner, enabling higher magnification factors than those allowed by conventional techniques. After an overview of related work, we introduce our techniques. *Speed* continuously adapts motor precision to navigation speed. *Key* and *Ring* use a discrete switch between two levels of precision (focus and context), the former using an additional input channel, the latter by decoupling the cursor from the lens’ center. We then report the results of two controlled experiments that evaluate focus targeting and object selection performance. Finally, we iterate our designs by integrating speed-dependent visual behaviors from the Sigma Lens framework [18]. The resulting hybrid lenses further improve performance, as shown in a third controlled experiment.

RELATED WORK

Most techniques for navigating multi-scale information spaces are based on either overview + detail, zooming or focus + context (see Cockburn et al. [8] for a very thorough survey). Zooming interfaces, e.g., [21, 14] display a single level of scale and therefore require a temporal separation to transition between “focus” and “context” views. They usually do not suffer from quantization effects, but both views cannot be observed simultaneously. Overview+detail interfaces [13, 22] show both views simultaneously using spatial separation, still requiring some mental effort to integrate the two views. They usually allow pixel-precise selections in the detail region, but focus targeting is also subject to quantization problems in conventional bird’s eye views.

Focus+context techniques “aim to decrease the short term memory load associated with assimilating distinct views of a system” [8] by integrating the focus region inside the context. This integration, however, limits the range of magnification factors of practical use. Basic magnifying glasses occlude the surroundings of the magnified region [12]. To address this issue, distortion oriented techniques provide a smooth transition between the focus and context views. Distortion, however, causes problems for focus targeting and understanding of the visual scene. Carpendale et al. [4] describe elaborate transitions that enhance the rendering of the distorted area and make higher magnifications comprehensible from a visual perspective. Gutwin’s *Speed-coupled flattening lens* [12] cancels distortion when the lens is repositioned by the user, thus removing a major hindrance to focus targeting. The Sigma Lens framework [18] generalizes the idea of speed-coupling to a larger set of lens parameters. For example, the *Speed-coupled blending lens* makes focus targeting easier from a motor perspective by increasing the focus region’s size for the same overall lens size, using a dynamically varying translucence level to smoothly transition between focus and context.

Although their primary goal is different, focus+context interfaces share issues with techniques designed to facilitate pointing on the desktop. The decoupling of visual and motor spaces plays a central role in techniques designed to facilitate the selection of small targets, e.g., [6, 7, 17] – see [2] for a detailed survey. Not designed for exploratory multi-scale navigation, but closer to our problem are pointing lenses [19], which punctually enlarge both visual and motor space to facilitate small target selection through stylus input. However, visual space is enlarged by duplicating the pixels of the original representation. The popup vernier [1] enables users to make precise, sub-pixel adjustments to the position of objects by transitioning from coarse to fine-grain dragging mode through an explicit mode switch. The technique provides visual feedback based on the metaphor of vernier calipers to make precise adjustments between both scales.

LENSES WITH HIGH-PRECISION MOTOR CONTROL

The quantization effect is due to the mismatch between visual and motor space precision in the focus region. This mismatch, in turn, is caused by the following two properties of conventional lenses:

- (P1) the cursor is located at the center of the lens, and
- (P2) the cursor location is controlled in context space.

These properties cause problems with the two low-level actions performed by users: *focus targeting*, and *object selection* within the magnified region. In this section we introduce three techniques that address these problems by breaking the above properties.

For all our techniques, lens displacements of less than MM focus pixels, corresponding to displacements of less than 1 context pixel, are achieved by slightly moving the representation in the focus region while keeping the cursor stationary (see discussion of Experiment 2’s results for more detail).

Precision through Mode Switching: the *Key* technique

The first approach to address the problem is to provide a way of controlling the location of the lens in focus space (as opposed to context space). We immediately discard the solution that consists in solely interacting in focus space because of obvious performance issues to navigate moderate to large distances (all distances are multiplied by MM in focus space). The simplest technique uses two control modes: a *context speed* mode and a *focus speed* mode. This requires an additional input channel to perform the mode switch, for instance using a modifier key such as SHIFT. Users can then navigate large distances at *context speed*, where one input device unit is mapped to one context pixel, i.e., MM focus pixels, and perform precise adjustments at *focus speed*, where one input device unit corresponds to one focus pixel.

Figure 3 illustrates this technique, called *Key*: the first case (No modifier) is represented by the topmost grey line; the second case (Shift pressed) by the bottommost grey line. When SHIFT is pressed, (P2) is broken. A similar “precision mode” is already available in, e.g., Microsoft Office to freely position objects away from the intersections formed by the underlying virtual grid using a modifier key.




Figure 3. Displacement in focus space (in pixels) for one input device unit move in function of the input device speed ($MM = 4$).

The *Key* technique represents a simple solution. However, as the selection tools based on Magic Lenses [3], an additional channel is required to make the explicit mode switch. Bi-manual input techniques are still uncommon. Modifier keys tend to be used for other purposes by applications, and their use often results in a “*slightly less than seamless interaction style*” [2]. The next two techniques we propose do not require any additional input channel.

Speed-dependent Motor Precision: the *Speed* technique

Following recent works that successfully used speed-dependent properties to facilitate pointing [5] and multi-scale navigation [12, 14, 18], our first idea was to map the precision of the lens control to the input device’s speed with a *continuous* function, relying on the assumption that a high speed is used to navigate large distances while a low speed is more characteristic of a precise adjustment (as observed for classical pointing [2]).

The black line (Speed) in Figure 3 illustrates the behavior of our speed-dependent precision lens. Cursor instant speed s is computed as the mean speed over the last four move events. It is mapped to the lens’ speed so as to break (P2) as follows:

- (i) if $s < MIN_SPEED$ then the lens moves at *focus speed* ;
- (ii) if $MIN_SPEED \leq s \leq MAX_SPEED$ then the lens moves by x focus-pixels for 1 input device unit, where x is $1 + (1 - \frac{MAX_SPEED - s}{MAX_SPEED - MIN_SPEED}) \times (MM - 1)$;
- (iii) if $s > MAX_SPEED$ then the lens moves at *context speed* like a conventional lens.

Cursor-in-flat-top Motor Precision: the *Ring* technique

The last technique is inspired by Tracking menus [10]. Consider a large rigid ring (e.g., a bracelet) on a flat surface (e.g., a desk). The ring can be moved by putting a finger inside it and then moving that finger while keeping it in contact with the surface to pull the ring. This is the basic metaphor used to interact with the *Ring* lens: the ring is the lens’ focus region (called the *flat-top*) and the cursor is the finger.

The *Ring* lens breaks property (P1): it decouples the cursor from the lens center; the cursor can freely move *within* the flat-top at *focus scale*, thus enabling pixel-precise pointing in the magnified region (bottommost grey line (Inside ring) in Figure 3). When the cursor comes into contact with the flat-top’s border, it pulls the lens at *context speed*, enabling fast repositioning of the lens in the information space (topmost

grey line (Pushing ring) in Figure 3). Figure 5 illustrates the lens behavior when the cursor comes into contact with the ring: the segment joining the lens center (g) to the contact point (p) is progressively aligned with the cursor's direction.

Decoupling the cursor's location from the lens' center has a drawback when changing direction: because the user has to move the cursor to the other end of the flat-top before she can pull the lens in the opposite direction. We tried to address this issue by pushing the physical metaphor: we introduced friction in the model to make the ring slide when the cursor stops, with the effect of repositioning the lens' center so as to match the cursor's position. We were not able however to get a satisfying result, and abandoned the idea.

EXPERIMENTS

We conducted two experiments to compare the performance and limits of the three lenses described above. Participants were asked to perform a simple task: selecting an object in the magnified area. The targets were laid out in a circular manner and the order of appearance forced participants to perform the task in every direction, following the recommendations of the ISO 9241-9 standard [9]. Only one target was visible at a time so that participants could not take advantage of the layout to facilitate the task: as soon as the participant clicked on one target, the next target appeared. The recorded movement time is the interval between the appearance of the target and a click on it. The target is presented as a yellow circle on a gray background, and is always surrounded by a 10-pixel red square clearly visible in the context view. The background is also decorated by a grid to help participants understand the transition between context and focus view, and to minimize desert fog effects [15] that can occur with scenes that are too uniform.

Analysis of the Task

A *pointing task* with a lens is typically divided in two main phases: (i) *focus targeting*, which consists in putting a given target inside the flat-top of the lens (Figure 4-(a) and (b)) and (ii) *cursor pointing* to precisely position the cursor over the target (Figure 4-(b) and (c)).

The *focus targeting* task has an index of difficulty of about:

$$ID_{FT} = \log_2 \left(1 + \frac{D_c}{(W_{FT_c} - W_c)} \right)$$

where W_{FT_c} and W_c are the respective sizes of the flat-top and the target in context pixels, and D_c is the distance to the target in context pixels as well¹. This formula clearly shows that difficulty increases as distance increases, as the size of the flat-top decreases, and as the size of the target decreases. The size of the flat-top in context pixels is directly related to the magnification factor of the lens, MM . Indeed, the size of the flat-top is fixed in terms of focus pixels, so the higher MM , the smaller the size of the magnified area in context pixels (see [18] for an analysis of the difficulty of a focus targeting task).

¹ ID_{FT} is the exact index of difficulty when the target must be fully contained in the flat-top. Here the task is slightly easier because the target just has to intersect the flat-top.




Figure 5. Bottom: behavior of a *Ring* lens when the cursor comes into contact with the flat-top's border at the bottom of the ring and then moves to the right. Top: Computation of the ring's location.

The final *cursor pointing* task mainly depends on the area of the target in focus space that intersects the flat-top after the focus targeting task. The larger this area, the easier the cursor pointing task. We can at least consider the best case, i.e., when the target is fully contained in the flat-top. In this case, the difficulty of the cursor pointing task can be assessed by the ratio $\frac{D_f}{W_f}$ where D_f is the distance between the cursor and the target, and W_f is the motor size of the target when magnified in the flat-top. The distance D_f is small, i.e., smaller than the flat-top's diameter, so we assume that the difficulty of the cursor pointing task is mainly caused by the value of W_f . Note that for regular lenses, the value of W_f is actually the size of the target at context scale because the target is only visually magnified. With our lenses however, since pixel-precise selections are possible, W_f is the magnified size of the target (at focus scale). We provide additional details about the division between the two subtasks in the following sections.

The first experiment tests pointing tasks with an average level of difficulty, while the second one tests pointing tasks with a very high level of difficulty, involving targets smaller-than-a-pixel wide at context scale. Our experimental design involves the three factors that determine the pointing task difficulty introduced above: the distance to the target (D_c), its width (W_c), and the lens' magnification factor MM .

Experiments: Apparatus

We conducted the experiments on a desktop computer running Java 1.5 using the open-source ZVTM toolkit. The display was a 21" LCD monitor with a resolution of 1600 x 1200 (≈ 100 dpi). The mouse was a regular optical desktop mouse at 400 dpi with the default acceleration function.

Experiment 1: Design

The goal of the first experiment is to test whether any of the three techniques we introduced in the previous section degrade performance when compared with regular lenses (*Reg*). We expect them to improve overall performance because the overall task difficulty is theoretically lower. On the one hand, the focus targeting task should not be harder: since we test small targets with lenses having the same flat-top size, the distance in context space is the main factor contributing to difficulty. All our lenses are able to navigate large distances like a regular lens, i.e., move at *context speed* (Key:

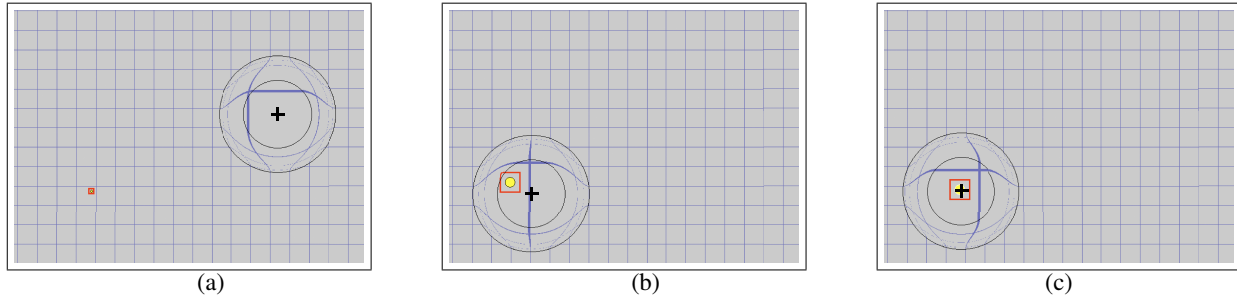


Figure 4. Screenshots of our experimental task: *focus targeting* from (a) to (b) and, *cursor pointing* from (b) to (c). Screenshots have been cropped to show details, and cursors have been made thicker to improve readability.

when SHIFT is released; *Ring*: when the cursor pulls the lens; *Speed*: when the lens moves fast enough). On the other hand, cursor pointing should be easier since the difficulty of this second phase mainly depends on the target’s motor width in focus space. Since all of our lenses allow to navigate at *focus speed*, they can take benefit of the magnified target size whereas this is not the case with a regular lens: even though it is magnified, the target size in motor space is the same as if it were not magnified.

Sixteen unpaid volunteers (14 male, 2 female), age 20 to 35 year-old (average 26.8, median 26), all with normal or corrected to normal vision, served in Experiment 1. Experiment 1 was a $4 \times 2 \times 2 \times 3$ within-subject design with the following factors:

- Technique: $TECH \in \{Speed, Key, Ring, Reg\}$
- Magnification: $MM \in \{4, 8\}$
- Distance between targets (context pixels): $DC \in \{400, 800\}$
- Target width (context pixels): $WC \in \{1, 3, 5\}$

We grouped trials into four blocks, one per technique ($TECH$), so as not to disturb participants with too many changes between lenses. The presentation order was counterbalanced across participants using a Latin square. Within a $TECH$ block, each participant saw two sub-blocks, one per value of magnification factor (MM). The presentation order of the two values of MM was also counterbalanced across techniques and participants. For each $TECH \times MM$ condition, participants experienced a series of 12 trials per $DC \times WC$ condition, i.e., 12 targets laid out in a circular pattern as described earlier. We used a random order to present these $2 \times 3 = 6$ series within a sub-block. We removed the first trial of each series from our analyses as the cursor location is not controlled when a series begins. To summarize, we collected $4 \text{ TECH} \times 2 \text{ MM} \times 2 \text{ DC} \times 3 \text{ WC} \times (12-1) \text{ replications} \times 16 \text{ participants} = 8448$ trials for analysis. Before each $TECH$ condition, the experimenter took 2-3 minutes to explain the technique to be used next. Participants were told each time the value of MM was about to change, and had to complete 4 series of practice trials for each new $TECH \times MM$ condition.

Experiment 1: Results and Discussion

Our analysis is based on the full factorial model:

$$TECH \times MM \times WC \times DC \times \text{Random}(\text{PARTICIPANT})$$

with the following measures:

- *FTT*, the focus targeting time;

- *CPT*, the cursor pointing time;
- $MT = FTT + CPT$, the time interval between the appearance of the target and a successful mouse press on it (this measure includes penalties caused by errors); and
- *ER*, the error rate (an error is a press outside the target).

Analysis of variance reveals an effect of $TECH$ on MT ($F_{3,45} = 15.2, p < 0.0001$). A Tukey post-hoc test shows that *Reg* is the significantly slowest technique and that *Key* is significantly faster than *Ring*. Note that there is no significant difference between *Ring* and *Speed*, nor between *Speed* and *Key*. Participants also made more errors with *Reg* than with our techniques. We expected *Reg* to perform worse since, as we already mentioned, the target’s motor size is in context pixels for *Reg* whereas it is in focus pixels for *Key*, *Speed* and *Ring*. The target is thus much harder to acquire in the *CPT* phase. Analysis of variance shows a significant effect of $TECH$ ($F_{3,45} = 18.5, p < 0.0001$) on *ER*. Figures 6-(a) and (b) respectively show the time MT and error rate *ER* for each $TECH \times WC$ condition.

We find a significant effect of DC ($F_{1,15} = 121.9, p < 0.0001$) on movement time MT . It is consistent with our expectations: DC has a significant effect on FTT ($F_{1,15} = 165, p < 0.0001$) while it does not on *CPT* ($p=0.4$). The higher the value of DC , the harder the focus targeting phase. Our techniques do not seem to be at a disadvantage in this phase compared to *Reg* since the effect of $DC \times TECH$ on FTT is not significant ($p=0.9$).

MM also has a significant effect on MT ($F_{1,15} = 249.6, p < 0.0001$), the effect being distributed across both FTT ($F_{1,15} = 515, p < 0.0001$) and *CPT* ($F_{1,15} = 79, p < 0.0001$). Figure 6-(c) shows the three measures per $TECH \times MM$: a bar represents MT per condition while the line shows the repartition between FTT (lower part of the bar) and *CPT* (upper part)². This clearly shows that a high MM leads to high FTT since the flat-top size in context pixels directly depends on MM , as explained in the previous section. A higher MM also means a larger target width in focus pixels. This can explain the effect of MM on *CPT*: *CPT* decreases as MM increases.

The target width in focus pixels is of course also related to WC , which is consistent with our observations: WC has an effect on both (i) FTT ($F_{2,30} = 45, p < 0.0001$) and (ii) *CPT* ($F_{2,30} = 1110, p < 0.0001$), and also on MT ($F_{2,30} =$

²Error bars in the figures represent the 95% confidence limits of the sample mean ($mean \pm StdErr \times 1.96$).




Figure 6. Movement time (a) and error rate (b) per TECH \times WC. (c) Movement time per TECH \times MM. For (a) and (c), the lower part of each bar represents focus targeting time, the upper part cursor pointing time.

623.8, $p < 0.0001$, Figure 6-(a)). Indeed, as we expected, the smaller wc, the higher the focus targeting time (i). Also, the larger wc, the larger the target in focus pixels to improve focus pointing time (ii). Regarding error rate, wc ($F_{2,30} = 17.5$, $p < 0.0001$) and MM ($F_{1,15} = 16.8$, $p = 0.0009$) have a significant effect on ER: participants made more errors when the target size was small. This is a simple interpretation that explains the difference in means that we observe; but we have to refine it to reflect the more complex phenomenon that actually takes place.

Coming back to the effect of TECH, we also observe two significant interaction effects that involve TECH on MT.

First interaction effect: TECH \times MM ($F_{3,45} = 4.7$, $p = 0.0063$) which can be observed on Figure 6-(c). A Tukey post-hoc test shows that for MM = 4, Speed, Key and Ring are significantly faster than Reg but this test also shows that for MM = 8, only Key and Speed are significantly faster than Reg (Ring no longer is). A closer look at the focus targeting phase explains why Ring seems to suffer from high magnification factors. We know that FTT increases as MM increases. We can observe on Figures 6-(c) and (a) that Ring is actually slower than the other techniques for this FTT phase. This is probably due to the cost of repairing overshoot errors during this phase: changes in direction are costly with Ring since the user first has to move the cursor to the opposite side of the flat-top before being able to pull the lens in the opposite direction.

Second interaction effect: TECH \times WC ($F_{6,90} = 55.1$, $p < 0.0001$) which can be observed on Figure 6-(a). A Tukey post-hoc test shows a significant difference in mean for wc=1 between Reg and the other techniques, while this difference is not significant for wc=3 and wc=5. To better assess the interpretation of such a result, we consider finer analyses on CPT. Figure 7 shows CPT for each TECH \times MM \times WC condition. Analyses reveal significant effects of TECH, MM and WC and significant interactions TECH \times MM and TECH \times WC (all $p < 0.0001$) on CPT. Tukey post-hoc tests show that Key, Speed and Ring are globally faster than Reg for cursor pointing. This is not surprising since the motor size of the target is smaller for Reg than for the others, as we said earlier. However, this significant difference holds only for wc=1 and wc=3, not for wc=5. In the latter case, only Speed is significantly faster than Reg. Moreover Ring is faster than Key for wc= 1, while Speed is not. These results suggest that Ring is particularly efficient for very small targets and that Speed is more appropriate for larger ones.




Figure 7. Cursor pointing time per TECH \times MM \times WC condition.

The latter observations suggest that modeling the movement time MT as the sum of FTT and CPT ($MT = FTT + CPT$) may be too naive to explain the subtle differences between techniques. For instance, this model does not explain the differences between Ring and Speed that depend on wc. In the same spirit, we observe that the difference between Reg and other lenses for wc=5 is very small considering that the target's motor size is 5 for Reg and 20 (MM=4) or 40 (MM=8) for Key, Speed and Ring. The additive model based also fails to explain the following observation: Speed features significantly higher FTT values than Key and Reg for MM=8 only. We tentatively explain this by the increased difficulty of controlling a lens with speed-dependent precision when the slope of the mapping function is too steep (linear function from MIN_SPEED to MAX_SPEED, i.e., focus speed to context speed on Figure 3). We tried several variations that, e.g., depend on the difference between these two speeds, without success. Using a gentler slope is frustrating because of the stickiness caused by the large movements required to reach the MAX_SPEED threshold. The more subtle differences we reported in the second part of this section may be explained by the fact that a transition phase between the focus targeting phase (FTT) and the cursor pointing phase (CPT) actually exists for our lenses: pressing a key for Key, stop pulling the flat-top for Ring, performing speed adjustments with Speed.

At the end of the experiment, participant were asked to rank the lenses (with ex-aequo allowed) using two criteria: perceived usability and performance. These two rankings were almost the same for all participants. All but one ranked Reg as their least preferred technique (one participant ranked it third with Speed fourth). There was no significant difference among other lenses. For instance, 8 participants ranked

Speed first, 3 ranked it second; 6 participants ranked *Key* first, 5 ranked it second, and 5 participants ranked *Ring* first, 7 ranked it second. We also asked participants to comment on the techniques. The main reason for the bad ranking of *Reg* is the great difficulty to acquire small targets, related to the *cursor jumping* effect due to quantization. Regarding *Speed*, most participants found the technique “*natural*”; some found the speed “*difficult to control*”. The participants who ranked *Key* high justified it by a “*transparent control*”; other participants complained about the need to use two hands. Regarding *Ring*, the cursor pointing phase was found easier because the lens is stationary, but participants also raised the overshooting problem discussed earlier.

To summarize, in comparison with regular lenses, precision lenses increase pointing accuracy. They also increase selection speed for small targets and are as fast for larger ones.

Experiment 2: Design

This second experiment evaluates our techniques on extreme tasks: very small target sizes and high magnification factors. We discard the *Reg* technique as it is not capable of achieving sub-pixel pointing tasks, i.e., involving targets that are smaller-than-a-pixel wide in context space. Another difference with Experiment 1 is that we use W_F as a factor instead of W_C . This allows us to isolate the effects of W_F and MM . Indeed, since $W_F = W_C \times MM$, two values of MM correspond to two different values of W_F for the same W_C value.

Twelve participants from Experiment 1 (10 male, 2 female), age 20 to 35 year-old (average 27.25, median 26.5), also served in Experiment 2. Experiment 2 was a $3 \times 2 \times 2 \times 3$ within-subject design with the following factors:

- $TECH \in \{Speed, Key, Ring\}$
- $MM \in \{8, 12\}$
- $DC \in \{400, 800\}$
- $W_F \in \{3, 5, 7\}$

As in Experiment 1, trials were blocked by technique, with presentation order counterbalanced across participants using a Latin square. The experimenter explained the technique to be used during 2-3 minutes before each $TECH$ condition. For each $TECH$, participants saw the two values of MM , grouped into two sub-blocks (sub-block presentation order were counterbalanced across techniques and participants). Each sub-block contained 6 series of 8 trials, 1 series per $DC \times MM$ condition, presented in a random order. To summarize, we collected $3 TECH \times 2 MM \times 2 DC \times 3 W_C \times (8-1) replications \times 12 participants = 3024$ trials for analysis. As in Experiment 1, participants were alerted by a message each time the MM value changed and had to complete 4 practice series for each $TECH \times MM$ condition.

Experiment 2: Results and Discussion

Our analysis is based on the full factorial model:

$$TECH \times MM \times W_F \times DC \times Random(PARTICIPANT)$$

We consider the same measures as in Experiment 1: task completion time MT , focus targeting time FTT , cursor pointing time CPT and error rate ER .




Figure 8. Movement time per $TECH \times MM$. The lower part of each bar represents focus targeting time, the upper part cursor pointing time.

Analysis of variance reveals simple effects of W_F ($F_{2,22} = 68$), MM ($F_{1,11} = 393$) and DC ($F_{1,11} = 65$) on MT (all $p < 0.0001$). As expected, MT increases as W_F decreases, as MM increases and as DC increases. Participants also make significantly more errors when W_F decreases (3.67% for $W_F = 7$, 5.36% for $W_F = 5$ and 8.82% for $W_F = 3$).

The differences in movement time MT among techniques is significant ($F_{2,22} = 21.6, p < 0.0001$) while the difference in error rate is not (6.15% for *Speed*, 6.05% for *Key* and 5.65% for *Ring*).

There is an interaction effect $TECH \times MM$ on MT ($F_{2,22} = 24.8, p < 0.0001$): Tukey post-hoc tests show that *Ring* and *Key* are significantly faster than *Speed* but only for $MM=12$ while these differences are not significant for $MM=8$. Figure 8 shows that this large difference at $MM=12$ is due to a sharp increase of focus targeting time (FTT) for *Speed*. Comments from participants confirm that the speed dependent control of motor precision is too hard when the difference between context scale and focus scale is too high, resulting in abrupt transitions. With *Speed*, participants did not succeed in controlling their speed: either they overshoot the target (targeting speed too high) or spent a lot of time putting the target in focus (speed too low). Therefore, *Speed* does not seem to be a suitable lens for pointing with a very high magnification factor: at $MM=12$, the linear function linking focus speed to context speed is too steep to be usable.

Figure 8 shows that focus targeting performance of *Ring* degrades as MM increases. However, good cursor pointing performance compensates for it, resulting in good overall task completion time. Figure 9 shows CPT for each $TECH \times MM \times W_C$ condition. Analysis of variance reveals a significant effect of W_F ($F_{2,22} = 230, p < 0.0001$) on CPT . As mentioned earlier, the larger W_F , the easier the cursor pointing task. However, the effects of MM ($F_{1,11} = 154, p < 0.0001$) and $TECH$ ($F_{2,22} = 64, p < 0.0001$) on CPT are less straightforward to interpret. CPT is higher when $MM=12$ than when $MM=8$, *Ring* is faster than *Key* and *Speed*, and the difference between *Ring* and both *Key* and *Speed* is larger when $MM=12$ than when $MM=8$ (the $TECH \times MM$ interaction is indeed significant on CPT , $F_{2,22} = 9.8, p = 0.0009$).

A plausible explanation for these effects lies in the differences in terms of Control-Display (C-D) gain among tech-




Figure 9. Cursor pointing time per TECH \times MM \times Wf condition.

niques in the cursor pointing phase³. Figure 10 illustrates the difference in terms of control-display gain among lenses, all in high-precision mode. During the cursor pointing phase, *Ring* is stationary; only the cursor moves inside a static flat-top. This is not the case for *Key* and *Speed* for which high-precision cursor pointing is achieved through a combination of cursor movement and flat-top offset. In Figure 10, to achieve a mouse displacement of 15 units, the cursor has moved by 1 context pixel (= 8 focus pixels) and the representation has moved by 7 focus pixels to achieve an overall displacement of 15 focus pixels. As a result, the control-display gain is divided by MM for *Key* and *Speed*. This might be the cause for the observed performance degradation. This interpretation is consistent with the stronger degradation for *Key* and *Speed* than for *Ring* from MM=8 to MM=12. Note, however, that there is still a small degradation of CPT from MM=8 to MM=12 for *Ring*, that we tentatively explain by a harder focus targeting phase when MM=12 that influences the transition from focus targeting to cursor pointing.

To summarize, when pushed to extreme conditions, the *Speed* lens becomes significantly slower than the other precision lenses while *Ring* remains as fast as *Key* without requiring an additional input channel for mode switching.

MOTOR CONTROL COMBINED WITH VISUAL FEEDBACK

Previous experiments show that techniques with advanced *motor* behaviors enable higher-precision focus targeting and object selection while increasing the upper limit of usable magnification factors. The Sigma Lens framework [18] takes a different approach at solving the same general problem by proposing advanced *visual* behaviors. We now explore how to combine these two orthogonal approaches to create hybrid lenses that further improve performance.

Sigma Lenses with High-Precision Motor Control

The two Sigma lens visual designs reported as the most efficient ones in [18] can be directly combined with our motor designs. The first one is the *Speed-coupled blending* (abbreviated *Blend*): it behaves as a simple magnifying glass whose translucence varies depending on lens speed. Smooth transition between focus and context is achieved through dynamic alpha blending instead of distortion. This enables a larger flat-top for the same overall lens size, reducing the

³The ratio between the distances traveled by the *cursor* and the *input device*, both expressed in metric units.




Figure 10. Difference in control-display gain between *Ring* and *Speed/Key* lenses (MM=8). In *italic*: cursor location on screen.

focus targeting task's index of difficulty. The other design (abbreviated *Flat*) is a variation on Gutwin's original *Speed-coupled flattening* [12]. The lens flattens itself into the context as its speed increases so as to eliminate the problems caused by distortion. Figure 11 illustrates both behaviors.

We designed four new techniques that result from the combination of one of the above two visual behaviors with either speed-dependent motor precision (*Speed*) or cursor-in-flat-top motor precision (*Ring*). *Key* was discarded because it proved awkward to combine explicit mode switching with speed-dependent visual properties.

Speed + Flat: this lens behaves like the original *Speed* design, except that the magnification factor decreases toward 1 as speed increases (Figure 11-a). The main advantage is that distortion no longer hinders focus targeting. Additionally, flattening provides indirect visual feedback about the lens' precision in motor space: it operates in context space when flattened, in focus space when not flattened.

Ring + Flat: This lens behaves like the original *Ring* design, with the magnification factor varying as above. As a consequence, the flat-top shrinks to a much smaller size (time stamp *t3* on Figure 11-a), thus making course corrections during focus targeting easier since the cursor is still restricted to that area. As above, distortion is canceled during focus targeting.

Ring + Blend: This distortion-free lens behaves like the original *Ring* design, except that the restricted area in which the cursor can evolve (the flat-top) is larger (time stamps *t1* and *t5* in Figure 11-b). As speed increases, the flat-top fades out, thus revealing the context during the focus targeting phase (time stamps *t2* to *t4*). An inner circle fades in, representing the region that will actually be magnified in the flat-top if the lens stops moving. The cursor is restricted to that smaller area, making course corrections less costly.

Speed + Blend: This lens behaves like the original *Speed* design without any distortion. As above, the flat-top fades out as speed increases and fades back in as speed decreases. Again, the larger flat-top reduces the focus targeting task's index of difficulty. In a way similar to *Speed + Flat*, blending provides indirect visual feedback about the lens' precision in motor space: it operates in context space when transparent, in focus space when opaque.

Experiment 3: Design

Our goal is to evaluate the potential benefits of combining techniques that enable higher motor precision with visual behaviors based on speed-coupling. We use *Static* versions,

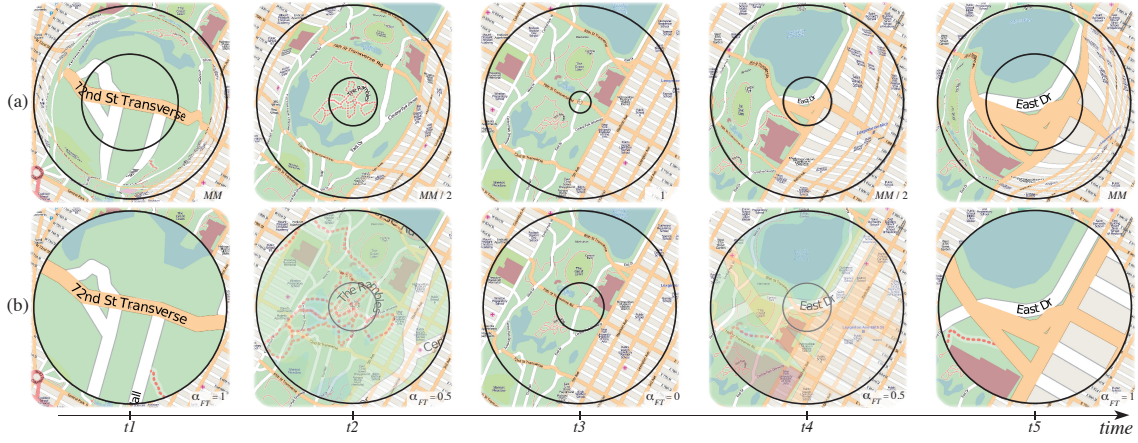


Figure 11. Behavior of two Sigma lenses during a focus targeting task ending on *East Drive* in Central Park. (a) As speed increases, the *speed-coupled flattening lens* smoothly flattens itself into the context (from $t1$ to $t3$), and gradually reverts to its original magnification factor when the target has been reached ($t4$ and $t5$). The inner circle delimits the region magnified in the flat-top. (b) As speed increases, the *speed-coupled blending lens* smoothly fades into the context (from $t1$ to $t3$), and gradually fades back in when the target has been reached ($t4$ and $t5$). The inner circle fades in as the lens fades out; it delimits which region of the context gets magnified in the lens. The magnification factor remains constant.

i.e., without any dynamic visual behavior, of our *Ring* and *Speed* techniques as baselines. Experiment 2 revealed that problems arise for the *difficult* tasks. We thus consider here difficult conditions in terms of magnification and target size. To reduce the length of the experiment, we discarded the Dc factor (distance between targets) as it did not raise any particular issue for any of the techniques.

Twelve participants from the previous experiments served in Experiment 3. Experiment 3 was a $2 \times 3 \times 2 \times 3$ within-subject design with the following factors:

- Motor precision technique: $TECH \in \{Speed, Ring\}$
- Visual behavior: $VB \in \{Blend, Flat, Static\}$
- Magnification: $MM \in \{8, 12\}$
- Target width in focus pixels: $WF \in \{3, 7, 15\}$

Trials were grouped into two main blocks, one per technique ($TECH$). These blocks were divided into three secondary blocks, one per visual behavior. The presentation order of $TECH$ main blocks and VB secondary blocks was counterbalanced across participants using a Latin square. Within a $TECH \times VB$ block, each participant saw two sub-blocks, one per magnification factor (MM); presentation order was counterbalanced as well. For each $TECH \times VB \times MM$ condition, participants experienced 3 series of 8 trials, one per value of WF , presented in a random order. We collected $2 \times 3 \times 2 \times 3 \times 3 \times (8-1)$ replications $\times 12$ participants = 3024 trials for analysis. As with the other two experiments, participants received a short explanation before each $TECH \times VB$ condition and performed 3 practice trial series per $TECH \times VB \times MM$ condition.

Experiment 3: Results and Discussion

As in Experiments 1 and 2, we perform analyses of variances with the full factorial model $VB \times TECH \times MM \times WF \times Random(PARTICIPANT)$ for MT , FTT , CPT and ER . Tukey post-hoc tests are used for pairwise comparisons.

As expected, we find a simple effect of VB on MT ($F_{2,22} =$




Figure 12. Movement time (MT) per VB by $TECH \times MM$ condition. The lower part of each bar represents focus targeting time (FTT), the upper part cursor pointing time (CPT).

$67, p < 0.0001$) revealing that visual behaviors significantly improve overall performance. Even if CPT is significantly degraded, the gain in FTT is strong enough (significantly) to decrease MT (see Figure 12). The degraded cursor pointing performance observed here is not surprising. It can be explained by the time it takes for a speed-coupled blending lens to become opaque enough or for a speed-coupled flattening lens to revert to its actual magnification factor. The performance gain measured for the focus targeting phase is consistent with previous experimental results [12, 18]. Overall, the gain in the focus targeting phase is strong enough to improve overall task performance.

The effects of WF and MM on MT are consistent with the previous two experiments: MT increases as WF decreases and as MM increases. *Ring* is still significantly faster than *Speed* ($TECH$ has a significant effect on MT : $F_{1,11} = 153, p < 0.0001$). Even if visual speed-coupling improves the performance of *Speed* more than that of *Ring* (significant interaction effect of $TECH \times VB$ on MT : $F_{1,11} = 11, p = 0.0005$), *Ring* remains faster than *Speed* for each MM . However, the advantage of *Ring* over *Speed* is significant only for $MM=12$ when we consider only the two speed-coupling techniques ($TECH \times MM$ on MT is significant, $F_{1,11} = 227, p < 0.0001$, as well as $VB \times TECH \times MM$, $F_{2,22} = 21, p < 0.0001$).

Note that we do not observe a significant advantage of *Blend* over *Flat* as reported in [18]. The main difference is that our targets are much smaller than those tested with Sigma lenses (0.25 to 1.9 context pixels in our experiment vs. 8 context pixels in [18]). Small targets probably cause more overshoot errors that are more expensive to repair with *Blend* than with *Flat*: if the larger flat-top of *Blend* is supposed to make focus targeting easier under an error-free hypothesis, it also causes an area of occlusion that is a significant drawback when trying to correct overshoots. Our participants actually reported that observation; in case of an overshoot they often left the target zone completely to perform a new focus targeting task. However this interpretation should be taken carefully since we did not record the number of overshoot errors. We only measured *ER*, the percentage of clicks outside the target (5.15% for *Blend*, 5.55% for *Flat* and 4.36% for *Static*). As in Experiment 2, the only factor that has an effect on error rate is target width w_T .

SUMMARY AND FUTURE WORK

Large differences in scale between focus and context views cause a quantization problem that makes it difficult to precisely position lenses and to acquire small targets. Quantization severely limits the range of magnification factors that can be used in practice. We have introduced three high-precision techniques that address this problem, making focus targeting and object selection more efficient while allowing for higher magnification factors than regular lenses. This is confirmed by the results of our evaluations, which also reveal that some lenses are more robust than others for extreme conditions, with the *Ring* technique performing the best. Our high-precision techniques can be made even more efficient by combining them with speed-dependent visual behaviors drawn from the Sigma lens framework, as shown in the last experiment.

We analyzed our observations based on a model for target acquisition that sums the *focus targeting* and *cursor pointing* time to get the overall task time. Our results suggest that this model is too simple as it ignores the transition period between the two subtasks. This is especially true for lenses with a speed-dependent behavior, because of the delay to revert back to their stationary configuration. As future work we plan to refine the additive model to better account for these transitions. We also plan to adapt our techniques to other focus-context interfaces and investigate non-circular focus shapes.

REFERENCES

1. Y. Ayatsuka, J. Rekimoto, and S. Matsuoka. Popup vernier: a tool for sub-pixel-pitch dragging with smooth mode transition. In *Proc. UIST '98*, 39–48. ACM, 1998.
2. R. Balakrishnan. "Beating" Fitts' law: virtual enhancements for pointing facilitation. *IJHCS*, 61(6):857–874, 2004.
3. E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Toolglass and magic lenses: the see-through interface. In *Proc. SIGGRAPH '93*, 73–80. ACM, 1993.
4. S. Carpendale, J. Ligh, and E. Pattison. Achieving higher magnification in context. In *Proc. UIST '04*, 71–80. ACM, 2004.
5. O. Chapuis, J. Labrune, and E. Pietriga. Dynaspot: Speed-dependent area cursor. In *Proc. CHI '09*, 1391–1400. ACM, 2009.
6. A. Cockburn and P. Brock. Human on-line response to visual and motor target expansion. In *Proc. GI '06*, 81–87, 2006.
7. A. Cockburn and A. Firth. Improving the acquisition of small targets. In *Proc. BCS-HCI '03*, 181–196, 2003.
8. A. Cockburn, A. Karlson, and B. B. Bederson. A review of overview+detail, zooming, and focus+context interfaces. *ACM CSUR*, 41(1):1–31, 2008.
9. S. A. Douglas, A. E. Kirkpatrick, and I. S. MacKenzie. Testing pointing device performance and user assessment with the ISO 9241, part 9 standard. In *Proc. CHI '99*, 215–222. ACM, 1999.
10. G. Fitzmaurice, A. Khan, R. Pieké, B. Buxton, and G. Kurtz. Tracking menus. In *Proc. UIST '03*, 71–79. ACM, 2003.
11. G. W. Furnas and B. B. Bederson. Space-scale diagrams: understanding multiscale interfaces. In *Proc. CHI '95*, 234–241. ACM & Addison-Wesley, 1995.
12. C. Gutwin. Improving focus targeting in interactive fish-eye views. In *Proc. CHI '02*, 267–274. ACM, 2002.
13. K. Hornbæk, B. B. Bederson, and C. Plaisant. Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM ToCHI*, 9(4):362–389, 2002.
14. T. Igarashi and K. Hinckley. Speed-dependent automatic zooming for browsing large documents. In *Proc. UIST '00*, 139–148. ACM, 2000.
15. S. Jul and G. W. Furnas. Critical zones in desert fog: aids to multiscale navigation. In *Proc. UIST '98*, 97–106. ACM, 1998.
16. J. Lamping and R. Rao. Laying out and visualizing large trees using a hyperbolic space. In *Proc. UIST '94*, 13–14. ACM, 1994.
17. M. J. McGuffin and R. Balakrishnan. Fitts' law and expanding targets: Experimental studies and designs for user interfaces. *ACM ToCHI*, 12(4):388–422, 2005.
18. E. Pietriga and C. Appert. Sigma lenses: focus-context transitions combining space, time and translucence. In *Proc. CHI '08*, 1343–1352. ACM, 2008.
19. G. Ramos, A. Cockburn, R. Balakrishnan, and M. Beaudouin-Lafon. Pointing lenses: facilitating stylus input through visual-and motor-space magnification. In *Proc. CHI '07*, 757–766. ACM, 2007.
20. M. Sarkar, S. S. Snibbe, O. J. Tversky, and S. P. Reiss. Stretching the rubber sheet: a metaphor for viewing large layouts on small screens. In *Proc. UIST '93*, 81–91. ACM, 1993.
21. J. J. van Wijk and W. A. Nuij. A model for smooth viewing and navigation of large 2d information spaces. *IEEE TVCG*, 10(4):447–458, 2004.
22. C. Ware and M. Lewis. The DragMag image magnifier. In *Proc. CHI '95*, 407–408. ACM, 1995.