

five: Enhancing 3D Wall Displays with a 2D High-Resolution Overlay

Daniel Steffen¹, Achim Ebert¹, Matthias Deller¹, and Peter Dannenmann²

¹ Deutsches Forschungszentrum für Künstliche Intelligenz GmbH,
Trippstadter Str. 122, 67663 Kaiserslautern, Germany
{Daniel.Steffen,Achim.Ebert,Matthias.Deller}@dfki.de

² Fachhochschule Wiesbaden, Fachbereich Ingenieurwissenschaften,
Am Brückweg 26, 65428 Rüsselsheim, Germany
dannenmann@mndu.fh-wiesbaden.de

Abstract. Projection-based stereoscopic wall displays allow users to immerse themselves into virtual scenes, such as architecture simulations or games. However, the usually low resolution (dpi) of such displays and slight alignment offsets between the two projectors result in a loss of detail and bad readability of textual information.

We propose addressing the problem by overlaying a third projector, so that its image is visible to both eyes. We eliminate offset artifacts by extracting 2D contents from the scene and rendering it using this dedicated “2D projector”. In addition, we locally increase resolution by focusing the 2D projector onto a smaller region. This allows us to reduce the size of overlaid 2D annotations, thereby reducing interference with the 3D scene. In our paper, we describe the design of our display system called *five*. Following, we present two detailed user studies that compare *five* with an overview-and-detail and a pan-and-zoom interface.

Keywords: Focus-plus-context screens, stereoscopy, peripheral vision, experimental evaluation.

1 Introduction

Projection-based stereoscopic wall displays allow users to immerse themselves into virtual scenes. Application scenarios include construction in mechanical or civil engineering, process control, and interactive games. In remote surgery, stereoscopic presentation allows exact hand-eye coordination. In architecture, 3D visualization helps convey complex spatial relationships without requiring experience with the interpretation of blueprints. Displaying contents stereoscopically can also improve users’ perception.

Many 3D scenarios involve context information. Providing this information to the users in the form of annotations, such as text overlays, allows users to perform visual sensemaking tasks [18]. Unfortunately, stereo projection systems can only display detail and especially text in a limited way, because the wide viewing angle required for immersion comes at the expense of detail. In addition, subtle alignment offsets between the two projectors result in double images, further impacting on the




Fig. 1. The five display system complements wide-angle stereoscopic projection with a more focused, 2D hi-res projector. This allows overlaying readable annotations, such as text.

readability of 2D contents, such as text (Fig. 3a). As a result, displaying a 3D scene on a stereoscopic projection wall limits the number of annotations that can be displayed at reasonable readability. At the same time, the annotations have to be comparably large, a fact which can lead to substantial interference with the 3D contents it is overlaid onto.

2 The *five* Display System

We propose solving the problem by using a hybrid display system that combines stereoscopic projection with an additional third projector dedicated to rendering 2D annotations. Fig. 1 illustrates a model of the display system we built. Fig. 2 shows a schematic view of it. Since the system combines 2D and 3D, we named it *five*. It is inspired by *focus-plus-context screens* [4], but extends the concept into *mixed dimensionalities*, i.e., the context area is not only of a different resolution, but also of a different dimensionality. The resulting system offers benefits for tasks that require stereoscopic presentation and high resolution. Various application scenarios can be found, in which the data to be visualized is inherently 3D as well as 2D. Construction datasets (like in car manufacturing, architectural design, urban planning, etc.) are challenging examples that contain three-dimensional geometries and also user manuals, measurements, etc. Thus, only an interweaved visualization can properly transport the needed contextual information.

As shown in Fig. 2, the first two projectors create a wide-angle stereoscopic image, i.e. one projector renders the scene from the perspective of the user's left eye, the other one from the right eye. Filters over each projector polarize the light emerging from them; users wear matching filter glasses; this assures that each of the user's eyes sees only the image intended for it. We are using a standard PowerWall™ system for that purpose.

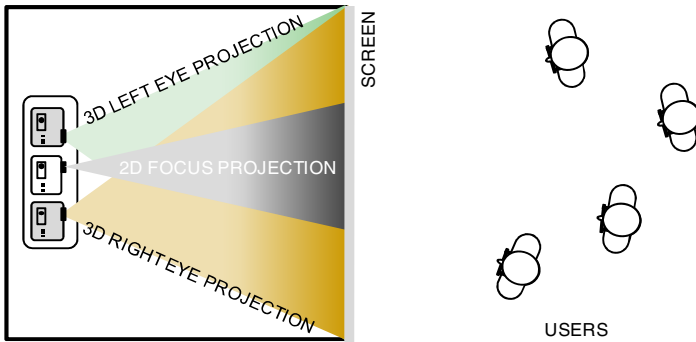


Fig. 2. Top view of the *five* prototype setup

The third projector does not use a polarization filter, which makes its projection visible to both eyes, and thus 2D. We focus this projection to a smaller region, resulting in an area of increased resolution. We typically point this *focus* projector to the display center, as shown in Fig. 2. This allows displaying information in high detail in the area of interest and does not affect three-dimensional context space by reason of its small size. By focusing the projector, light intensity per surface increases. We use a software-based filter to adapt the brightness of the focus projector to the same level as the one of the *context* projectors.

Compared to an all-high-resolution stereoscopic display, the resulting display has two limitations. Firstly, hi-res annotations can only be displayed in that area since the focus-projector covers only a sub-area of the wall display. Secondly, since both eyes see the same focus image, overlays offer no stereo disparity. Hi-res annotations are therefore flat and are always perceived as being located in the plane of the projection screen.

Before we display a scene on *five*, we partially alleviate these limitations by pre-processing the scene in the following five steps. (1) We translate the 3D scene such that as many annotations as possible fall onto the focus area. In Fig. 1, for example, we center a vanishing point over the focus area, as it holds the majority of objects in the scene. We then translate the scene in z -direction in order to position annotations as close to the projection plane as possible. This minimizes local z offsets, when we switch the annotation layer to 2D. We illustrate this process at the example of a document environment in the next section. (2) We extract the 2D annotations from the 3D scene. Since projection is based on the additive color model, we mask the space that held the annotation in the 3D model with a black area. Projecting onto the black mask with the focus projector will then result in the original brightness. (3) Where beneficial, we move additional annotations from the context area into the focus area. We link these annotations with the objects they describe, typically using lines turning them into *external labels*, as described in [7]. In cases where adding the line would introduce more clutter than can be saved by having it in the hi-res area, we do not translate that annotation but keep it low-res and stereo-projected. (4) We scale annotations down to reduce interfere with the 3D scene (Fig. 3c) or to be able to show additional annotations or additional details of an annotation (Fig. 3d).




Fig. 3. (a) On a stereo projection wall: low resolution and alignment offsets impact readability of 2D text (b) Overlaying the text using a separate dedicated “2D” projector eliminates these artifacts. The underlying 3D scene is still rendered using the stereo projectors. (c) The increased resolution allows reducing the size of the 2D annotation, here revealing a leak in the fuel container. This minimizes interference with the underlying 3D scene or (d) can be used to overlay additional, here potentially vital 2D contents.

Alternatively, we render annotations at their original size to maximize readability of the overlay (Fig. 3b) or we mix all three objectives. (5) We render the scene.

In summary, delegating 2D annotations to a dedicated projector leads to a display that is more readable and/or more informative. The additional detail can be used to (1) reduce artifacts substantially, resulting in more readable overlays, as shown in Fig. 3b, (2) minimize the occlusion of 3D contents, as shown in Fig. 3c, (3) display additional annotations or to display annotations in additional detail, as shown in Fig. 3d, or a combination of these.

Even though high-res annotations are limited to a certain display area and z-depth, adding the third projector significantly increases the amount of information that can be displayed and that can be directly associated with 3D on-screen objects. When regular stereoscopic projection runs out of annotation space, additional annotations can only be made available by flipping back and forth between multiple views or by switching back and forth visually between the display wall and an external display. As we found in the two studies mentioned earlier, retrieving such external annotations costs time and imposes cognitive load on users. Overlaying a focus projector, in contrast, increases the information density in situ, thereby minimizing switching effort and cognitive load.

3 Demo Application

To demonstrate our prototype and to illustrate the pre-processing process of 3D scenes, we implemented a visual document management application. It allows users to retrieve documents based on spatial memory. By mounting the focus projector in portrait mode we maximize the benefit for document viewing in this application.

Fig. 4 shows a walkthrough. In clockwise order: (a) In the default view, each document is shown as a thumbnail. Documents are visually grouped by shifting them horizontally and vertically. Proximity to the viewer, i.e. promotion in z , indicates relevance either absolute or with respect to a query entered by the user. The user selects a thumbnail to see that document in full, readable detail. (b) Resolution and




Fig. 4. (a) To view a document, (b) the document space is distorted to make space and the selected thumbnail moves to the center, (c-d) the thumbnail grows, and (e-f) is finally replaced by a readable high-resolution version of the page

sharpness are insufficient for displaying the selected stereoscopic projection in situ, but the embedded 2D focus projector offers sufficient resolution. The context space distorts to free up the focus area and the selected thumbnail is *copied* into the focus region, (c) where it is smoothly scaled (d) until it fills the entire focus area. At the same time, the document moves towards the user until it reaches the projection plane (stereo-disparity zero). At this point, the entire display is still projected stereoscopically. (e) The focus part of the 3D projects is faded to black and (f) the 2D focus area is faded in, turning the document high-resolution.

To visually clarify which thumbnail was moved, the system can alternatively link the thumbnail copy to its original using a *rubber band*, as shown in Fig. 4f (see *drag-and-pop* [5]).

4 The Five Prototype

Our prototype is based on a PowerWall™, which serves as the 3D context display. Two JVC DLA-SX21S projectors with a 1400 x 1050 resolution rear-project the stereoscopic image pairs onto a 2.9m x 2.3m semi-translucent screen through circular polarizing filters of opposite handedness. A Windows PC controls the stereoscopic visualization. It features two NVIDIA GeForce 8800 GTX graphics cards, 2 GByte main memory as well as two Intel Core 2 Duo CPUs with 2.4 GHz each.

The focus display is implemented using an Epson EMP-TW 1000 projector, located below the two JVC projectors. It offers a resolution of 1920 x 1080 pixels. We focused it on an area of 1.03m x 0.58m. As a result, a context pixel corresponds in size to approximately 5x5 focus pixels (Fig. 3).




Fig. 5. Calculating the position of the 2D focus area within the focal plane of the 3D context

five runs on top of our OpenGL and C++ based framework called *AnyScreen* on MS Windows. It controls all three projectors. Moreover, it offers four stereoscopic (side-by-side, horizontally and vertically interlaced, and anaglyph stereoscopy) and one non-stereoscopic rendering mode for each render window. In the traditional stereo configuration, we set up the two stereo projectors in side-by-side mode. We add the focus projector as a third screen located in the center of the stereoscopic projection (Fig. 2).

In order to integrate the 2D focus projection seamlessly into the 3D virtual environment, the position of the focus region is configured using the following calibration process: Using the nomenclature from Fig. 5, the calibration process determines the coordinates (x_{3D}, y_{3D}, z_{3D}) in the 3D context that correspond to the coordinates (x_{2D}, y_{2D}) in the 2D focus region. Using this correspondence, we compute the 3D context coordinates for every object in the 2D focus region.

Since the 2D focus region is located on the projection plane, its z value is predefined. We have set up our 3D coordinate system in a way that the viewing direction is always along the negative z -Axis and the projection plane is orthogonal to the viewing direction. Thus, in this system the projection plane is described by $z = const$. Therefore, the z_{3D} coordinate is easily determined for every object that is lying in the 2D focus region.

For calculating the x_{3D} and y_{3D} coordinates, the 3D coordinates $(x_{l,3D}, y_{l,3D})$ and $(x_{r,3D}, y_{b,3D})$ that correspond to the upper left-hand and lower right-hand corners of the 2D focus region have to be determined. Once this information is available, the corresponding x and y 3D-coordinates for any pair (x_{2D}, y_{2D}) of 2D-coordinates can be computed as:

$$\begin{pmatrix} x_{3D} \\ y_{3D} \end{pmatrix} = \begin{pmatrix} x_{l,3D} \\ y_{l,3D} \end{pmatrix} + \begin{pmatrix} x_{2D} w_{2D}^{-1} (x_{r,3D} - x_{l,3D}) \\ y_{2D} h_{2D}^{-1} (y_{b,3D} - y_{l,3D}) \end{pmatrix}$$

To determine the coordinates $(x_{l,3D}, y_{l,3D})$ and $(x_{r,3D}, y_{b,3D})$ of the corners of the 2D focus area calibration, a calibration dialog is displayed on the focus and the context projectors.

5 Related Work

The purpose of the *five* system is to display additional information in limited space. It therefore addresses the same problem as corresponding software techniques (zooming, overviews, and fisheyes) and in particular focus-plus-context screens.

5.1 Zooming, Overviews, and Fisheyes

Zoomable interfaces [15] allow users to navigate information sequentially by panning and zooming. Semantic zooming makes it possible to assign different appearances to an object at different zoom factors. While zooming into a map, for example, additional detail may be revealed. Other variants include, *constant density zooming* [22] and *non-linear* panning and zooming [11].

Overview-and-detail interfaces consist of two or more display windows, one of which typically shows the entire information space, while one or more detail views provide a close-up view [12, 16]. These interfaces have been applied to a variety of applications, including program code editors [7] and image collections [16]. Positioning the detail view from within the overview allows efficient navigation [6]. In addition, the overview window helps users to keep track of their current position in the information space [16] so that they have the feeling that they are in control [21]. On the flipside, overviews require additional screen space and users are forced to switch between views to reorient themselves, which costs time [3].

Focus-plus-context techniques, in contrast, avoid switching between windows by combining detail area and periphery into a single view (e.g., the *bifocal display* by Apperley and Spence [1]). Examples for focus-plus-context techniques are *fish-eye views* [9, 10], the *perspective wall* [14], and *document lens* [17]. The main drawback of focus-plus-context approaches is that they introduce visual distortion, which can interfere with the user's task [4, 13].

5.2 Focus-Plus-Context Screens

This shortcoming is addressed by *focus-plus-context screens* [4], large low-resolution screens with an embedded high-resolution screen. Information is displayed across both display units, so that scale is preserved, while resolution varies. As a result, distortion is avoided. The original screen device was developed by Baudisch et al [4]. In a quantitative evaluation [3], participants performed map-navigation and chip-design tasks faster when using a focus-plus-context screen than when using a zooming or overview interface with the same number of pixels.

Several researchers have adapted and enhanced the idea of focus-plus-context screens. Ashdown created a more flexible system by using projectors for both focus and context [2]. Sanneblad and Holmquist made the focus area moveable by combining a high-resolution tablet with a large projected screen [19]. Flider and Bailey physically separated the focus screen from the context screen [8]. However, none of the current research has addressed the usage of the focus-plus-context screen in stereoscopic environments.

In order to verify our design, we conducted two user studies. The first one evaluates *five* in the context of a static document, the second in the context of a dynamic, real-time task.

6 User Study 1: Visual Search in a Static Scene

In this study, participants performed a visual search task on a static document: participants located five “buildings” of a given height within a map of simple, abstract buildings of different heights. Participants used a pan-and-zoom interface, an overview-and-detail interface, and an interface based on *five*.

Our main hypothesis was that participants would complete the task faster when using the *five* interface than when using the pan-and-zoom interface (which required manual navigation) and the overview-and-detail interface (which required users to reorient themselves when visually switching between views). We expected these results to be reflected in terms of subjective satisfaction as well. Given that participants could verify findings before committing to them, we expected negligible error rates across interfaces.

Participants. Twelve participants volunteered. Participants were students of the local university or technical/non-technical staff of our department. Participants ranged in age from 23 to 40 years, average 30.92. All participants had normal or corrected-to-normal vision.

Task. During each trial, participants were presented with a birds-eye view of an abstract 75x50 cell “map” containing 90 tall “buildings”, as shown in Fig. 6a. Participants were also given a height value. Their task was to locate the five buildings that were of the specified height on the 90-building map.

Participants could *estimate* building heights, e.g. based on the 3D cues provided by the stereo projection system. However, to be certain, participants had to look up the




Fig. 6. (a) The *five* interface condition allowed participants to use 3D depth cues to locate candidate “buildings” of the sought height. (b) The exact height and the building’s number were annotated onto the building’s base (here shown in focus view)

exact height, which was written onto the building's base as shown in Fig. 6b. Whenever a building of the sought height was found, participants would call out the building number shown just above the building height.

Height and number annotations were scaled to fit the width of the building. We thereby simulated the space limitations of a real city where all cells would be filled, preventing text annotations to grow into the space of neighboring cells. Both annotations were readable while in the 2D focus area, but not readable while located in the 3D context.

Interfaces. There were three interface conditions: *pan-and-zoom*, *overview-and-detail*, and *five*. All interface conditions were implemented using the hardware setup described earlier. Participants controlled the system using a mouse located on a 1.1m high desk about 2.5 meters right of the center of the projection as shown in Fig. 6a.

In the *pan-and-zoom* condition, the skyscraper scene including buildings and annotations was rendered using the stereo projectors. In order to read annotations, participants had to zoom in by a factor of 2.8 or more by rolling the mouse wheel. As a result, participant alternated between zooming out to locate a candidate building and zooming in to examine the exact height and determine the building number. Alternatively, participants could pan the display by moving the mouse with the left button pressed.

Also the *overview-and-detail* used only the stereo projectors to render the scene. Fig. 7 shows how this interface condition was set up. The stereoscopic display wall always showed a zoomed out perspective on the 3D scene. The additional *detail* projection located right of the display wall was zoomed in deep enough to show annotations. A red marker frame in the 3D representation indicated which region of the scene was shown in the detail view. When the 3D scene was panned, the detail in the additional projector changed accordingly.

The *five* interface displayed the scene also using the stereo projectors. Annotations, however, were rendered using a focus projector directly onto the base of the respective buildings whenever that building was located inside the focus area; otherwise no

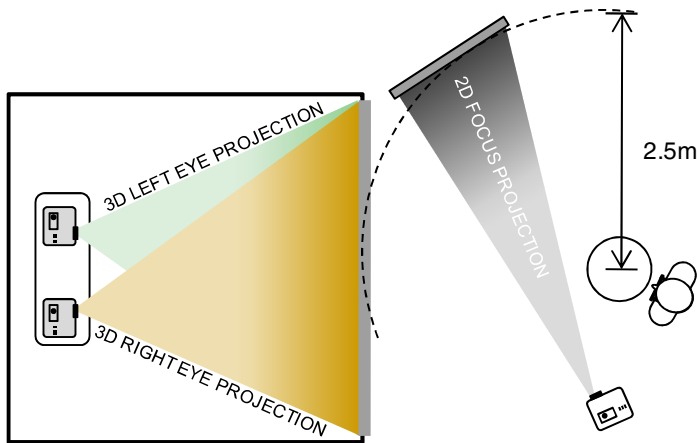


Fig. 7. Top view of the overview-and-detail interface

annotation was shown. As in the other conditions, participants panned using the mouse. The *five* interface did not support zooming.

Procedure. We used a within-subject experimental design, e.g. each participant carried out each task on all three interfaces. Interface order was counterbalanced following a Latin square of order 3. Maps were pre-computed and their order randomized. For each trial, we recorded task time and error rate (calling out the wrong building number).

Participants received a verbal explanation and training for each interface condition. Participants filled in a questionnaire after completing each interface condition. It contained questions on the participants' general impression of this interface (how comfortable / how easy to use / how fatigue-proof is it), its usability, and its suitability to the task, provided on a 7-item Likert scale. The study took less than 30 minutes per participant including training.

Results and Discussion. As expected, participants did not make any errors, across interface conditions. Fig. 8 shows overall task times.

We performed a one-way repeated measure ANOVA evaluating the within-subjects effect of display type (*five*, overview-and-detail and pan-and-zoom). The sphericity assumption was not met so the Huynh-Feldt correction was applied. The main effect of display type was significant, $F(1.476, 16.236) = 16.997$, $p < 0.005$.

Post-hoc comparisons were performed using the Bonferroni adjustment for multiple comparisons. The comparison revealed that the *five* interface yielded significant faster completion times than both the overview-and-detail ($p < 0.001$) and the pan-and-zoom ($p < 0.001$). However, the mean task completion time of overview-and-detail and pan-and-zoom was not significantly different ($p = 0.182$).

As hypothesized, participants performed the visual search task faster when using the *five* interface than when using the two competing interface conditions. Subjective ratings reflect these findings. We noted that the participants definitely preferred the *five* interface over both other ones. Summarizing, all participants rated the *five* interface most applicable to the given task (Question 5). Only a slight difference occurred between the ratings of the overview-and-detail and pan-and-zoom interface.

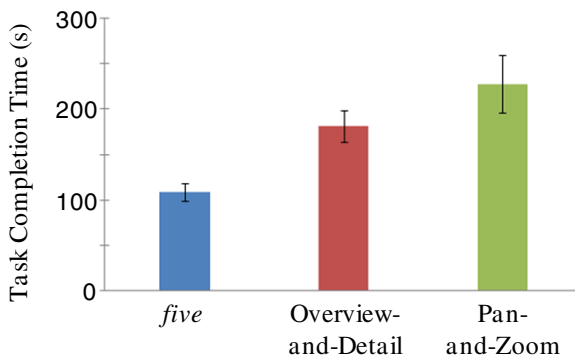


Fig. 8. Average task completion time in seconds (+/- standard error of the mean)

A repeated-measures ANOVA revealed a significant main effect of display ($F(2,22)=17.868$, $p<0.001$). Post-hoc comparisons using Bonferroni adjustment revealed that the *five* interface was more suitable for the task than the overview-and-detail one ($p<.001$) and the pan-and-zoom one ($p<.01$).

The performance benefit of the *five* interface over the pan-and-zoom condition was expected, as it was predictable that panning buildings into the focus area would be faster than zooming in and out.

The performance benefit compared to the overview-and-detail interface, however, does suggest that displaying information in situ leads to performance benefits. These findings are in accordance with [4], who found that requiring users to visually switch between overview-and-detail view in a 2D display system influences task performance.

The absence of time constraints allowed participants to complete the given task without error on all interface conditions. To investigate performance under time constraints, we conducted a second study. For this second study, we also used a more realistic scene geometry: in the first study, all annotations were located in the same plane; this prevented the introduction of depth error when flattening the annotation space to the 2D view. In the second study, we therefore used a scene with annotations along substantial depth.

7 User Study 2: Real-Time Interaction, Dynamic Scene

In the second user study, we investigated whether the *five* interface provides an advantage over overview-and-detail interface in the context of a dynamic scene. Participants performed a simplified real-time flight simulation that required them to simultaneously pay attention to both a 3D task (avoiding asteroids) and a 2D detail task (pick-up the right types of power-ups labeled in 2D located along the path).

We expected participants to obtain higher scores (more fuel, fewer trashcans and asteroids) when using the *five* condition. The reason was that the overview-and-detail condition would require participants to visually switch between views, introducing brief moments of reduced orientation. Again, we expected subjective satisfaction ratings to mirror this.

Participants. The same twelve participants from User Study 1 took part in User Study 2. We accepted that participants had additional experience using the different display setups, because potential training effects would have applied equally to both interface conditions.

Task. Participants were flying a simulated spacecraft, as shown in Fig. 9. The participants' task was to pick up fuel canisters (by simply flying through them) while avoiding trash canisters and asteroids. In order to do so, participants switched between four "lanes" arranged in a 2x2 grid (top left, top right, bottom left, and bottom right) using the arrow keys on a keyboard.

Fuel canisters and trashcans were floating in the middle of one of the lanes. Both were represented by similar looking blocks, labeled "fuel" or "trash" respectively. Fuel and trash canisters appeared nearly indistinguishable, then turned distinguished in the focus/detail view, and finally turned distinguishable also in the




Fig. 9. Participant's task was to pick up fuel canisters while avoiding trashcans and asteroids in real-time (the four white labels and lines are part of the figure and were not shown on screen)

context/overview, approximately one second before flying through them. Asteroids appeared in pairs from different sides of the lane and tumbled across lanes.

The participants were informed that fuel canisters were worth one point, trash cans a minus point, and asteroid hits five minus points. For each interface condition, participants played a 90-second sequence containing 40 fuel canisters, 40 trashcans, and 50 asteroids (25 pairs). Sequences were generated randomly for every user and trial.

Interfaces. There were two interface conditions: The overview-and-detail interface corresponding to the overview-and-detail interface from the first user study. However, the only interaction possible was lane switching. Participants switched lanes using the arrow keys on the keyboard. The *five* condition corresponded to the *five* condition in the first user study. As in the overview-and-detail, participants navigated using arrow keys. We did not include the pan-and-zoom interface, because the speed of the dynamic nature of the task left no time for zooming.

Procedure. As in the first study, participants received instructions and training, completed the task on all interface conditions in counter balanced order, and filled in the same questionnaire. For each trial we recorded the number of picked-up fuel canisters as well as the collisions with asteroids and trash canisters. The study took between 10 and 15 minutes per participant.

Results and Discussion. Fig. 10 summarizes error rates. A repeated-measure ANOVA evaluating the within-subjects effects (*five*, O+D) found a highly significant mean effect of interface, for all three performance indicators. When using the *five* interface, participants collected more fuel packages (45% more, $F(1,11)=39.105$, $p<.001$) while colliding with fewer trashcans (47%, $F(1,11)=113.626$, $p<.001$) and asteroids (51%, $F(1,11)=55.579$, $p<.001$) than when using the *overview-and-detail* interface.

As expected, subjective preferences reflected these findings. Again, we could note that the participants definitely preferred the *five* interface over the other one. Summarizing, all participants rated the *five* interface most applicable to the given task. A




Fig. 10. Mean numbers of fuel packages picked-up (higher is better) and obstacles collided-with (lower is better) during flight task (+/- standard error of the mean)

repeated-measures ANOVA showed a significant effect for interface type $F(1,11)=17.963, p<.001$.

Unlike our first study, the real-time constraints caused the performance difference to impact error rate rather than task time (which was controlled). Our findings support our main hypothesis: Displaying 2D focus and 3D context in a single integrated view helps users keep track and stay in control in interactive applications. The overview interface, in contrast required users to switch visual attention back and forth, which interfered with the double attention task. These findings broaden the results given in [3] by adding dimensionality as an additional parameter.

8 Conclusion

In this paper, we have presented a hybrid display system. By complementing wide-angle stereoscopic projection with a 2D “focus projector” the resulting display setup allows to display detailed annotations *collocated with* the objects they are referring to, thus minimizing clutter and cognitive load. The findings of our two user studies suggest that a display like *five* can offer benefits when working with static as well as dynamic real-time applications.

Encouraged by our findings, we have started to adapt additional applications to the *five* display system. A commercial first person shooter (Unreal Tournament 2004) is now running on the *five* display. The additional 2D projector makes the game’s heads-up display legible; alternatively, we could have used the increased clarity of the 2D projector to reduce the overlay size, thus reducing occlusion of 3D game contents.

As future work, we are planning to design visualization and interaction techniques that allow users to interact with objects and annotation on the *five* display system more easily.

Acknowledgments. We thank Patrick Baudisch (Microsoft Research) for the close cooperation, fruitful discussions, and constructive comments. We also thank all our study participants. This research has been funded by BMBF.

References

1. Apperley, M.D., Spence, R.: A bifocal display technique for data presentation. In: Proc. Eurographics 1982, pp. 27–43 (1982)
2. Ashdown, M., Robinson, P.: *Escritoire: A Personal Projected Display*. IEEE Multimedia 12(1), 34–42 (2005)
3. Baudisch, P., Good, N., Bellotti, V., Schraedley, P.: Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In: Proc. CHI 2002, pp. 259–266 (2002)
4. Baudisch, P., Good, N., Stewart, P.: Focus plus context screens: combining display technology with visualization techniques. In: Proc. UIST 2001, pp. 31–40 (2001)
5. Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P., Bederson, B., Zierlinger, A.: Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems. In: Proc. Interact 2003, pp. 57–64 (2003)
6. Beard, D., Walker, J.: Navigational Techniques to Improve the Display of Large Two-Dimensional Spaces. Behaviour and Information Technology 9(6), 451–466 (1990)
7. Bell, B., Feiner, S., Höllerer, T.: View management for virtual and augmented reality. In: Proc. UIST 2001, pp. 101–110 (2001)
8. Flider, M.J., Bailey, B.P.: An evaluation of techniques for controlling focus+context screens. In: Proc. GI 2004, pp. 135–144 (2004)
9. Furnas, G.W.: Generalized Fisheye Views. In: Proc. CHI 1986, pp. 16–23 (1986)
10. Furnas, G.W.: The Fisheye View: A New Look at Structured Files. In: Card, Mackinlay, Shneiderman (eds.) *Information Visualization: Using Vision to Think*, pp. 145–152. Morgan Kaufmann Publishers, San Francisco (1999)
11. Furnas, G.W., Zhang, X.: MuSE: a multiscale editor. In: Proc. UIST 1998, pp. 107–116 (1998)
12. Hornbaek, K., Frokjaer, E.: Reading of electronic documents: the usability of linear, fish-eye, and overview+detail interfaces. In: Proc. CHI 2001, pp. 293–300 (2001)
13. Leung, Y.K., Apperley, M.D.: A review and taxonomy of distortion-oriented presentation techniques. ACM TOCHI 1(2), 126–160 (1994)
14. Mackinlay, J.D., Robertson, G.G., Card, S.K.: The perspective wall: detail and context smoothly integrated. In: Proc. CHI 1991, pp. 173–176 (1991)
15. Perlin, K., Fox, D.: Pad: An Alternative Approach to the Computer Interface. In: Proc. SIGGRAPH 1993, pp. 57–64 (1993)
16. Plaisant, C., Carr, D., Shneiderman, B.: Image-browser taxonomy and guidelines for designers. IEEE Software 12(2), 21–32 (1995)
17. Robertson, G.G., Mackinlay, J.D.: The Document Lens. In: Card, Mackinlay, Shneiderman (eds.) *Information Visualization: Using Vision to Think*, pp. 562–569. Morgan Kaufmann Publishers, San Francisco (1999)
18. Russel, D.M., Stefik, M.J., Pirolli, P., Card, S.K.: The cost structure of sensemaking. In: Proc. CHI 1993, pp. 269–276 (1993)
19. Sanneblad, J., Holmquist, L.E.: Ubiquitous Graphics: Combining Hand-held and Wall-size Displays to Interact with Large Images. In: Proc. AVI 2006, pp. 373–377 (2006)
20. Shneiderman, B.: The Eyes Have It: A Task by Data Type Taxonomy for Information Visualization. In: Proc. IEEE Symposium on Visual Languages (VL 1996), pp. 336–343 (1996)
21. Shneiderman, B., Plaisant, C.: *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, 4th edn. Addison-Wesley, Reading (2005)
22. Woodruff, A., Landay, J., Stonebreaker, M.: Constant Information Density in Zoomable Interfaces. In: Proc. AVI 1998, pp. 110–119 (1998)