The Personal Cockpit: A Spatial Interface for Effective Task Switching on Head-Worn Displays

Barrett Ens bens@cs.umanitoba.ca
Rory Finnegan umfinne2@cc.umanitoba.ca
Pourang Irani irani@cs.umanitoba.ca
Department of Computer Science, University of Manitoba, Winnipeg, MB, Canada

ABSTRACT
As wearable computing goes mainstream, we must improve the state of interface design to keep users productive with natural-feeling interactions. We present the Personal Cockpit, a solution for mobile multitasking on head-worn displays. We appropriate empty space around the user to situate virtual windows for use with direct input. Through a design-space exploration, we run a series of user studies to fine-tune our layout of the Personal Cockpit. In our final evaluation, we compare our design against two baseline interfaces for switching between everyday mobile applications. This comparison highlights the deficiencies of current view-fixed displays, as the Personal Cockpit provides a 40% improvement in application switching time. We demonstrate of several useful implementations and a discussion of important problems for future implementation of our design on current and near-future wearable devices.

Author Keywords
Head-worn display; head-mounted display; task switching; spatial input; spatial user interface; virtual window management; multi-display environment

ACM Classification Keywords
H.5.2 **[Information Interfaces and Presentation]**: User Interfaces - Interaction styles; Windowing systems

INTRODUCTION
The recent proliferation of lightweight, low-cost, transparent head-worn displays (HWDs) makes it possible for users to view and interact with information content at all times. Commercially available systems, are capable of giving users access to the same content as on mobile devices, through virtual 2D displays that appear to float in the user's field of view (e.g. Epson Moverio, Vuzix). However, the practical scope of these current interfaces is narrow due in part to limitations of display configuration. Content is fixed to a single view location, restricted to the periphery, or occluding the wearer's view of his surroundings (Figure 1b-c). These conditions will also inhibit task switching, as is the case with mobile devices, where nearly 30% of tasks involve multiple applications [7] and the costs of switching are severe [21].

B. Ens, R. Finnegan and P. Irani. The Personal Cockpit: A Spatial Interface for Effective Task Switching on Head-Worn Displays. In *CHI '14: Proceedings of the SIGCHI Conference on Human Factors and Computing Systems*, 10 pages, to appear, ACM, 2014.

Unlike their handheld counterparts, HWDs need not be limited by physical display constraints. Instead, designers can leverage the 3D capabilities of such devices to appropriate the abundance of space around the display wearer. Multiple virtual windows can appear to 'float' in surrounding space, remaining easily accessible, but without unwanted occlusion. As in real-world multi-monitor environments [14], we can use virtual windows to partition information by task and by users' current needs. For instance, an on-the-go multitasker might place a map below his line of sight, looking down to consult it only as directions are required. Later, while waiting for the bus, he may place his calendar and a real-time bus schedule side-by-side, viewing each at his leisure with a turn of the head. We explore the design space for such a mobile and user-configurable arrangement of multiple, floating displays we call the *Personal Cockpit* (Figure 1a).

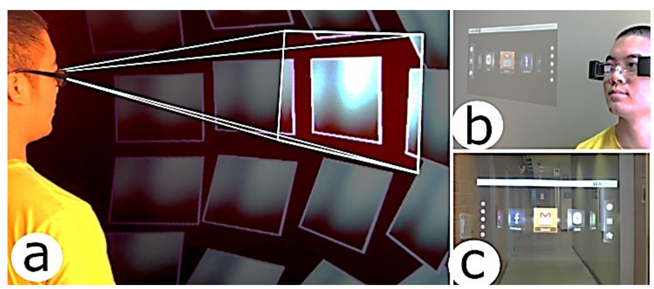

Figure 1. The Personal Cockpit (a) leverages an empirically-determined spatial layout of virtual windows. We investigate its design space, including field of view constraints of wearable displays. Our design is a shift from current interfaces (b, c), in which content remains fixed in the user's forward view.

Our work is inspired by past designs for HWDs, such as Feiner's implementation [12] of world- and body-fixed virtual windows. Subsequent studies [6, 9, 22, 34, 35] have indicated that leveraging users' perception of surrounding space may provide advantages over virtual navigation techniques. Our research builds on this prior work by exploring in depth a HWD interface that provides the benefits of proprioception and spatial memory. We craft a 2D layout customized to utilize head-motion with the constraints of a HWD's limited field of view (FoV). To fully exploit the potential of a spatial layout, we tune our design for use with direct input, akin to a personalized arrangement of floating 'touchscreens'. Our design process explores the relevant human factors, such as display size, distance, angular separation and spatial reference frames, in a high-fidelity, low-latency environment that minimizes confounds. Our study shows users can interact with the Personal Cockpit more effectively than with existing HWD navigation methods using view-fixed displays; when switching between a set of everyday applications, participants completed an analytic task 40% faster.

Our contributions are: 1) a design space outlining human factors related to effective information access on HWDs 2) an exploration of this design space that unveils relevant design issues and fine-tunes parameters for the Personal Cockpit; 3) an empirical demonstration, using high-fidelity systems, that our design is efficient for multi-task switching despite a limited FoV; and 4) a suite of interaction techniques for managing content in the Personal Cockpit. Our work is the first to apply such a rigorous design approach to body-centric HWD interfaces. Our thorough design space exploration provides an example for future research on interfaces for varying hardware configurations and lays several steps toward a user-configurable, multi-window management system.

RELATED WORK AND DESIGN FACTORS

We explore the design space for an interface ideally suited for multi-tasking on HWDs, the Personal Cockpit. Our work is inspired by a number of interfaces that leverage spatial memory to bridge the gap between real and digital worlds. Much of this work can be traced back to Fitzmaurice's information spaces [13], which map information to its associated physical locations in the real world. Feiner [12] later implemented a HWD interface with virtual windows mapped to world- and body-based reference frames. In Billinghurst's following work [5, 6], we see the potential of head-tracking for improving interaction with multiple displays. Many similar world- and body-centric concepts followed on other platforms such as spatially aware mobiles [30, 45] and projectors [8]. We build on these prior works by pinpointing relevant design issues that we use to guide our design process.

Field of View (FoV)

We are interested in how FoV limitations impact the Personal Cockpit. The human visual field spans about 200° horizontally and 130° vertically, however the detail-oriented foveal region of the eye spans only about 3° [28]. A wide FoV contributes to a user's sense of 'presence' in a virtual environment [23] and a limited FoV is known to hamper tasks relying on the user's peripheral view [28], such as navigation [40]. Due to limitations of technology, weight and cost, the display field of existing HWDs does not cover the entire human range. The available FoV on current low-cost HWDs varies between about 23º (e.g. Moverio) to 40º (e.g. Laster, Lumus). The impact of FoV on performance is gender dependent [11]. For tasks relying mainly on the foveal region, a 40º width may suffice [28].

Context Switching

Multiple displays have benefits over single displays for multitasking, particularly when display switching costs are minimized. Dual monitors can reduce workload and task time for frequent switching [19, 35] and much information can be safely relegated to a secondary display [35, 43]. Cauchard et al. [9] studied display separation in a mobile environment and found that context switching does not drastically impair performance of a visual search task, provided that head movement is minimized. Rashid et al. [29], however, found visual search to be slower when split between a mobile and large display than on either display alone. Spatial constancy in multi-window layouts can improve memorability and reduce switching time [31, 38].

Angular Separation

Given that the Personal Cockpit requires head movement, we consider the effects of angular separation between the multiple displays. The range of human neck motion for a normal adult is relatively large: about 85º for rotation to either side, 50º for vertical flexion (looking down) and 60-70º for vertical extension (looking up) [24]. However, the effective range for task switching is smaller. For example, Su and Baily [37] found that two displays on the same vertical plane can be displaced by up to 45º before negative effects on a docking task.

Display Size or Display Angular Width

Display size can influence task performance, although the effects are dependent on viewing distance. When viewing distance is held constant, we refer to display size as angular width. Ni et al. [26] found that large, high resolution displays improve performance on navigation, search and comparison tasks. Shupp et al. [34] found that large displays benefit certain tasks, such as search, but not others, such as route tracing. Ball and North [2] argue that the affordance of physical navigation has a greater effect on task performance than display size. Similarly, physical motion could prove advantageous for multitasking.

Window Distance

As our design of the Personal Cockpit includes direct user input, window distance is a primary design factor. For virtual displays, the impacts of depth are numerous. The minimum comfortable distance of binocular convergence is about 0.25 m [16], although ergonomics research recommends placing desktop monitors at a distance of at least one metre [1]. Tan and Czerwinski [39] found that performance is negatively impacted by mixed display distances. Thus our Personal Cockpit design should keep the working set of windows at a single depth. Estimation of depth is known to be impaired in virtual environments [10, 28, 40], due in part to FoV restrictions [43].

A well-understood phenomenon and cause of simulator sickness is vergence-accommodation mismatch. This effect occurs when the proprioceptive cues of focus and vergence become decoupled in stereoscopic environments [28, 33]. Until this issue is circumvented by technological advancements (e.g. [18]), HWD designers can reduce unpleasant effects by keeping the depth of virtual objects close to the surface of the virtual image plane [16, 33].

One further design consideration on HWDs with limited FoV is binocular overlap. As illustrated in Figure 2, the viewing frusta of both eyes typically overlap exactly at the distance of the display's virtual image plane. A device can be designed to allow a wider FoV by only partially overlapping the frusta. This choice comes at a trade-off in performance [28] due to monocular regions on the sides of the viewing region. Binocular overlap is also reduced when a large virtual object appears wider than the available viewing region. For example, the lower window in Figure 2 is cropped to a different region for each eye. One particular item of interest we explore is how the distance of a virtual display affects the interpretation of its contents.

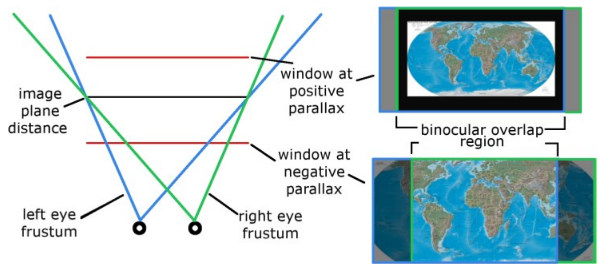

Figure 2. Binocular parallax creates an illusion of depth when objects appear in front of or behind the head-worn display's virtual image plane (left). If content appears wider than the available FoV (bottom right), binocular overlap is reduced.

Direct Input

Whereas the direct manipulation metaphor allows intuitive interaction with virtual objects [32], our Personal Cockpit design must take into account several issues inherent to 'touching the void': Depth perception of virtual objects is difficult and the depth estimation of a virtual surface is made more problematic by the lack of a tangible surface [10]. Furthermore, when distance is overestimated, the user's penetration of the surface can cause double vision, or diplopia [10, 42]. Also, interactive objects must remain within average maximum reach, about 50-60 cm to the front and 70-80 cm to the dominant side [25].

HWDs present additional challenges for direct input. In a wearable system, head-tracking and registration relies on body-fixed sensors. Thus, robust tracking and motion stabilization are required to create a convincing illusion of spatially situated objects. Also, since the display is located physically between the viewer and the locations of situated objects, a virtual display will occlude the user's hand as it reaches the surface. To make direct input feel natural, the system should detect the reaching hand and make it appear to occlude the virtual display. We circumvent these issues in our studies by emulating a HWD in a CAVE setting.

Spatial Reference Frames

A layout of multiple displays can be classified according to the spatial reference frame to which the displays are fixed (e.g. [5, 12]). For example, user elicitation study on organization of multi-display layouts [15] resulted in both environment-centric and user-centric layouts. We can similarly affix virtual displays to objects or location in the physical world (**world-fixed**) or to some part of the observer's body, such as the head (**view-fixed**), torso (**body-fixed**) or hand (**hand-fixed**). Because HWDs are easily portable, we explore the impact of different reference frames on direct input with the Personal Cockpit.

Display Layout Curvature

Shupp et al. [34] explored the differences between curved and flat layouts of multiple monitors when aligned to form a large display. Their study shows that performance with search and route tracing is 30% faster on the curved layout. This result may suggest that task switching is more efficient on a curved layout, which is well suited for reaching with an extended arm. Accordingly, we use a curved layout for the Personal Cockpit.

USER STUDIES

We refine our design of the Personal Cockpit as an advanced interface for multi-tasking on HWDs through four user studies. In the first 3 studies we fine-tune the design parameters (Figure 3) of display size (angular width), distance, reference frame and angular separation. In the last study we compare the Personal Cockpit against standard methods for task switching on view-fixed HWD interfaces.

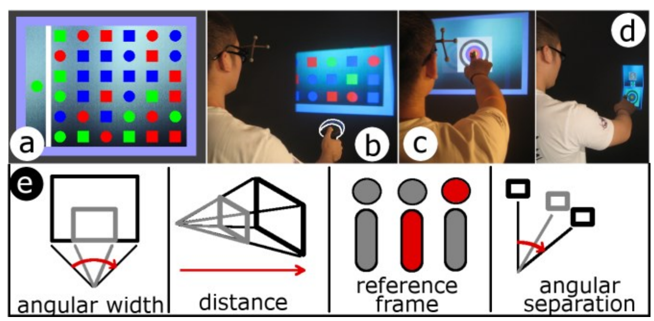

Figure 3. We used the results of our first 3 user studies (a-d) to tune the design parameters (e) of the personal cockpit.

Emulation Environment

As we focus on human factors limitations in our design, we run our studies in a projection-based CAVE environment. The low display latency and high-precision optical tracking enable us to explore designs not practical on today's hardware, to control for confounding background clutter and to examine previously untested design factors.

We emulate the restricted FoV of a HWD by clipping the viewing frusta of users to 40° × 30º (all windows in our studies have a 4:3 aspect ratio). We chose a 40° width because this angle is thought to be sufficient for detail-oriented tasks [28] and is within the range provided by currently available stereoscopic headsets (e.g. Lumus, Vuzix). As with actual see-through HWDs, the FoV restriction only affects virtual content; the real world view remains unobstructed.

To facilitate direct input, we explore visual output within reach of the user. However, many devices have a virtual image plane distance of 3 m or more (e.g. Moverio, Vuzix), which is impractical for use within reaching distance. Some devices have an adjustable image plane distance (e.g. Laster SmartVision), supporting objects in the near field. We emulate an image plane distance (Figure 2) of 1m, about the expected limit for use with direct input [16, 33]. As with FoV, this choice serves as a worst-case setting in which we evaluate the human-factors aspects of our design.

Our environment does not take all possible issues into account, for example vergence-accommodation mismatch (see Window Distance, above) or the problem of unwanted hand occlusion (see User Input, above and Transferring the Personal Cockpit to a HWD, below). However it allows us to examine issues related to FoV restriction such as the effects of binocular overlap and the efficiency of navigating to displays that are hidden out of view.

STUDY 1: SINGLE-DISPLAY VIEWING PARAMETERS

Our first study explores size and distance placement for a virtual display. These values depend on the FoV and distance limitations for direct input. Displays that appear wider than the FoV width are not fully visible from a single head position and may be difficult to interpret due to a reduction of the binocular overlap region (see Display Depth, above). We expect participants will be more efficient when the virtual display's angular width is equal to or less than the field of view. Participants may also prefer virtual displays that appear further away (i.e. with a lesser offset from the virtual image distance plane).

Participants, Task and Procedure

We recruited 10 university students (2 female, 21≤age≤40 years) from our local campus. We screened participants using a binocular depth test, which required them to differentiate between virtual displays placed at close (60 cm), intermediate (100 cm) and far (140 cm) distances. As a result of this test, we had to turn away 2 participants.

We implemented a visual search task to examine the effects of display width and distance. We use a conjunction search [41], in which the target and distracter objects share multiple properties. In our case, objects can share the same shape (square or circle) or colour (red, green or blue). The display is partitioned by a vertical line, with a target object appearing on left (Figure 3a-b). The participant must search an array of randomly generated objects on the right side of the line and count the number with the same shape and colour as the target object. Participants report their count by pressing one of four buttons on a handheld wand device. Virtual displays appear directly in front of the participant, centred at eye-height. Participants are asked to complete the task as quickly and as accurately as possible.

Design

We use a 5×4 within-participants design. The factors are *angular width* of the virtual display, relative to FoV (*50*, *75*, *100*, *125* or *150%)* and apparent *distance* of the window (*40*, *60*, *80* or *100 cm*). Conditions are presented to participants in random order to reduce learning effects. Within each condition, participants complete ten trials consecutively. To measure performance we record trial time and the number of incorrect selections. We collected 5 angular widths × 4 distances × 10 trials × 10 participants = 2000 data points. After each set of ten trials, participants provided perceived *effort* (on a 7-point scale) by answering the question "How hard did you have to work (mentally and physically) to accomplish your level of performance?"

Results

We analyzed data of recorded trial times and subjective scores of overall effort. In this study and those that follow, we remove outliers greater than 2SD from the mean.

Trial Time: We removed the first trial from each set (200 trials) to eliminate learning effects. We removed further 50 trials (2.78%) as outliers. The mean time of the remaining trials is 3.065 s (SD 1.157 s). We ran the univariate ANOVA for our analyses. Mean times for *angular width* and *distance* are shown in Figure 4. Results show a main effect of *angular width* ($F_{4,36.03}$=58.863, p<.001), but not distance ($F_{3,27.04}$=.106, p=.956). Post-hoc comparisons with Bonferroni corrections show significant differences between all *angular width* conditions (p<.001) except for *50* vs. *75* % (p=1.0).

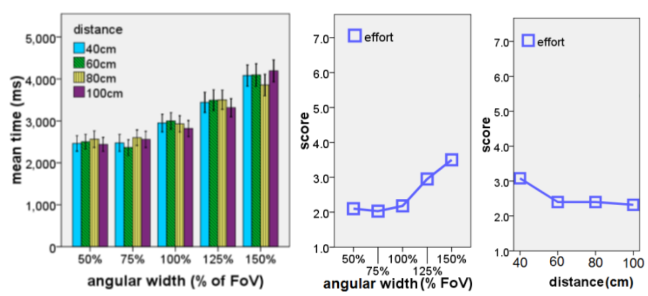

Figure 4. Mean trial times by *angular width* and *distance* (left). Mean *effort* for *width* and *distance* (right). Bars show ±2 SE

Effort: Participants provided scores after each condition for their perceived level of overall effort. We ran Friedman's ANOVA tests for each factor followed by post-hoc Wilcoxon tests. We found an effect of effort on both *angular width* ($\chi^2(4)$ = 63.44, p < .001) and *distance* ($\chi^2(3)$ = 22.15, p < .001). Mean scores are shown in Figure 4.

Discussion

We find that task time is directly influenced by the ratio of the display width to FoV. Task time is optimal when the virtual display is roughly ¾ the size of the FoV, likely due to reduced head motion. We see a small change from 100 to 75%, but no improvement with the smaller 50% ratio. Interestingly, perceived effort scores, in response to display width, follow an identical pattern to task time. We find that participants perceive increased discomfort at the nearest display distance (40 cm), but task performance is unaffected by distance. This result leaves open the possibility for direct input, as this latter factor is limited by the user's reach.

In the following study, we explore direct input for the average reach of 40-60 cm. As per our findings, in the remaining studies we restrict the window to be approximately ¾ width to fit completely within the FoV.

STUDY 2: SINGLE-DISPLAY INPUT PARAMETERS

Whereas Study 1 focused on visual output, Study 2 explores direct input. Our first goal is to determine which display distances best facilitate target selection. Our second goal is to see how the choice of spatial reference frame affects input that relies on reaching. In combination with Study 1, we can determine the ideal balance of design parameters to support both output and input. We expect that participants will benefit from proprioception with body-fixed or view-fixed windows, leading to lower targeting error.

Participants, Task and Procedure

We recruited 12 university students (2 female, 21≤age≤35 years). From a resting position, participants were asked to quickly and accurately 'touch' the centre of a 10 cm diameter bullseye target with their right hand (Figure 3c). The target is placed at one of 5 locations on a virtual window. Based on the outcome of the previous study, we chose a window width smaller than the FoV (70%). The target provided colour feedback to indicate correct or incorrect (the display is penetrated outside of the target region) selections. Participants began the next trial by returning their hand to the resting position. Input detection is provided by a Vicon tracking system.

Design

We used a 3×3 within-participants design. The factors are: spatial *reference frame* (*world-fixed*, *body-fixed* or *view-fixed*); *distance* of the display (*40*, *50* or *60* cm) and target *location* (*centre*, *top*, *bottom*, *left* or *right*). Body-fixed and view-fixed displays appeared at a set distance from the participant's body, as determined by the *distance* condition. World-fixed displays are initially set at the same distance, but are fixed to world-coordinates and do not move with the user. *Distance* and *reference frame* are presented in a random order to reduce learning effects. Within each condition, participants complete 5 blocks of trials. Within each block there is 1 trial at each *location*, presented in random order. To measure performance we record trial time and target selection error. Participants provide ratings of perceived *fatigue* for each combination *distance* and *reference frame* by answering the question "What was the level of fatigue from this task?" We collected 3 distances × 3 reference frames × 5 target locations × 5 trials × 12 participants = 2700 data points.

Results

We analyzed task completion *time*, pointing *error* and subjective ratings of *fatigue*. We found no effects of time.

Pointing Error: We define error as the distance between the detected selection and the target centre. For error analysis, we included all correctly completed trials. We compared error distances using a 3 × 3 × 5 univariate ANOVA. We found main effects of *distance* ($F_{2,22}$=4.443, p<.05), *reference frame* ($F_{2,22}$=13.759, p<.001) and *location* ($F_{4,44}$=4.780, p<.005) on pointing error. Post-hoc comparisons with Bonferroni corrections show significant differences between all pairs of *distance* ($p \leq .017$) and *reference frame* ($p \leq .003$). Mean pointing error distances are shown in Figure 5. There was also a significant interaction effect between *distance* and *location* ($F_{8,88}$=3.762, p=.001).

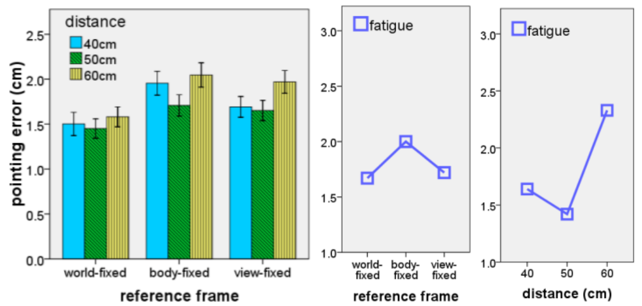

Figure 5. Mean pointing error by *reference frame* and *distance*. (left). Mean perceived *fatigue* (0-10) levels for *reference frame* (centre) and *distance* (right). Bars show ±2 SE

Fatigue: Participants rated *fatigue* on a 12-point Borg scale. As the Borg CR10 [27] scale was designed to be a linear mapping between perceived and actual intensity, we treat the resulting scores as scalar, using a univariate ANOVA. Results, as shown in Figure 5, reveal a significant effect of display *distance* ($F_{2,22}$ =13.162, p<.001). However, we did not find an effect of *reference frame* ($F_{2,22}$ =1.152, p=.334).

Discussion

We were surprised to find that target selection is clearly more precise in the world-fixed reference frame. Any benefits of proprioception in the other two reference frames were overshadowed by unintentional motion of the target window caused by the pointing motion. Although distance did not influence pointing speed, there was an unexpected effect of distance on pointing error. This effect was strongest in the body-fixed frame, i.e. when the window moves with the body, likely due to the unintentional window motion. Error was greatest at 60 cm, where participants' arms were near full extension. Precision was particularly bad in the top and left target locations, which required a slightly greater (right-handed) reach.

STUDY 3: MULTI WINDOW LAYOUTS

Having refined the distance parameter for direct input and visual output, we now investigate layouts of multiple windows, with target selection between two windows. Study 1 showed the best task performance when the window fits fully within view. Multiple tasks, however, are likely to occupy separate windows that span beyond the user's FoV. The ideal placement range is limited by human factors including the range of neck motion for a typical user and performance of direct input. As study 2 showed negative effects on pointing error from even subtle body motions, we use a world-fixed frame for optimal input. We choose a curved window layout for this study to keep targets within reach. However, a curved layout has a natural focal point. To determine if windows are best centred directly around the user, or offset to the side of the dominant pointing arm, we include *focal point* as a study factor. The *centre* focal point is symmetrical to the participant whereas the *right* focal point coincides with the right shoulder. All windows are placed at an equal distance (50 cm) from the current point of focus. Multiple windows are offset radially by a given *separation angle* (Figure 3e).

Participants, Task and Procedure

We recruited 8 university students (2 female, 1 left-handed, 21≤age≤35 years) from our local campus. Participants are presented with a two small windows (Figure 3d). One window contains a start button and is placed at shoulder height directly in front of the focal point (centre or right-offset). The second window contains a bullseye target, and is displaced either horizontally or vertically from the start window. The participant begins by 'touching' the start button, then moves quickly and accurately to the target.

Design

We use a 4×5×2 within-participants design. The factors are: *direction* of display displacement (*up*, *down*, *left* or *right*); displacement *angle* (*15°, 25°, 35°, 45°, 55°*) and point of *focus* (*centre* of body or *right* shoulder). For each *focus*, participants complete 10 consecutive blocks of trials, where 1 block contains all combinations of *direction* and *angle.* Trials in a block are presented in random order to prevent learning effects. The order of *focus* presentation is balanced between participants. We collected 4 directions × 5 displacement angles × 2 points of symmetry × 10 trials × 8 participants = 3200 total trials.

Results

Trial Time: Time is measured from the moment the start button is tapped until a selection is detected. For analysis of trial completion time, we included only correctly completed trials (i.e. the target selection falls within 5cm radius of the target centre.) We removed the first trial from each condition (320 trials) to reduce learning effects. We removed a further 88 trials (3.15%) as outliers. The mean time of the remaining trials is 0.81 s (SD 2.89 s). Mean trial times are shown in Figure 6. A univariate ANOVA reveals main effects of *direction* ($F_{3,21}$ =7.252, p<.005) and *angle* ($F_{4,28}$ =86.107, p<.001), but not for *focus*. Post-hoc tests with Bonferroni corrections showed significant differences between all pairs of *angles* and *directions* (all p<.001) except for *up* vs. *left* (p=1.0). There was also an interaction effect between *direction* × *angle* ($F_{12,84}$ =3.579, p<.001) as well as a 3-way interaction between *direction* × *angle* × *focus* ($F_{12,84}$ =2.678, p<.005).

Pointing Error: As in study 2, error is the distance between the detected selection and the target centre. For error analysis, we removed 138 (4.31%) outliers. Mean values are shown in Figure 6. A univariate ANOVA revealed main effects of *direction* ($F_{3,21.003}$=4.115, p<.05), *angle* ($F_{4,28.010}$=6.290, p<.001) and *focus* ($F_{1,7.002}$=21.204, p<.005).

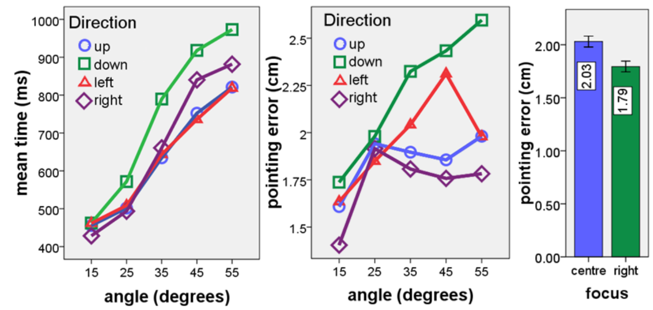

Figure 6. Mean trial times by *direction* and *angle* (left). Mean pointing error by *direction* and *angle* (middle). Mean pointing error by point of *focus* (right). Bars show ±2 SE

We also found a significant 3-way interaction for *direction* × *angle* × *focus* ($F_{12,84.156}$=2.816, p<.005).

Fatigue: Since this study requires both arm and head motion, we collected Borg ratings for both *arm* and *neck fatigue*. Due to the high number of conditions, we grouped separation angles into two *groups*, *low* (15º-35º) and *high* (45º-55º). We collected ratings for all combinations of *direction*, *group* and *focus*. We ran a 4×2×2 univariate ANOVA for each set of ratings. For *arm fatigue*, the test revealed significant effects of *direction* ($F_{3,21}$ =4.734, p<.05), *group* ($F_{1,7}$=15.465, p<.01) and *focus* ($F_{1,7}$ =5.984, p<.05). Neck fatigue showed only main effects of *direction* ($F_{3,21}$ =5.500, p<.01) and *group* ($F_{1,7}$ =13.213, p<.01). Results are shown in Figure 7.

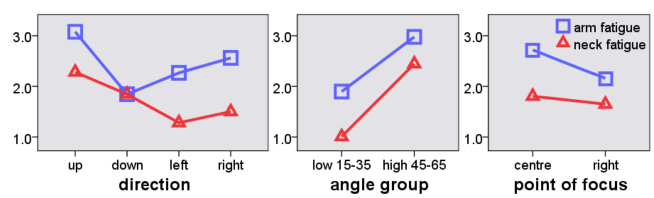

Figure 7. Mean perceived *arm fatigue* and *neck fatigue* for *direction*, *group* and *focus.*

Discussion

Time and error are both higher for targets in the *down* direction than for *up* (Figure 6). Despite this finding, several participants preferred the *down* direction to *up*, as it reduced arm fatigue (Figure 7). Pointing time generally increases with angle, as expected, due to increased travel distance. However, there is a steep increase in around the *35°* mark, when the start button and target both fit barely within view. Although *focus* doesn't affect pointing time, there is a significant reduction in error when the centre of curvature is shifted to align with the right shoulder. As a result of this finding, we explored various options for right-offset layouts before implementing the final study.

SUTDY 4: DISPLAY SWITCHING

The goal of this final study is to demonstrate that the Personal Cockpit, tailored based on the above set of results, facilitates effective task switching over current methods of application switching on HWDs. Whereas our first 3 studies

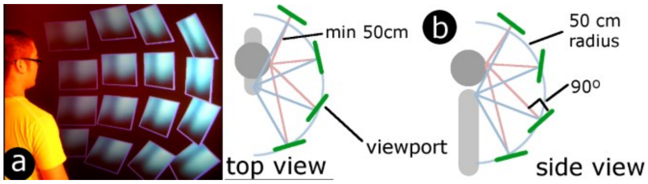

Figure 8. Our final design of the Personal Cockpit (a, b), based on findings from studies 1-3 and used in our final study.

explored subsets of the overall design space through abstract studies, we designed a more ecologically valid task for this 4th study.

Personal Cockpit Layout

We envision the Personal Cockpit as a versatile, configurable window manager that will be useful for many scenarios, including on-the-go multitasking. However, since study 2 showed that body-fixed windows are prone to target error, we use a world-fixed reference frame for our study. To keep windows within easy reach of the user, we chose a curved layout for the Personal Cockpit (Figure 8). Using the best input/output distance from studies 1 and 2, and the right-offset from study 3, we place each windows 50 cm from the user's right shoulder. To keep a 4×4 array within head range [24], we use a separation angle of 27.5º. To prevent window overlap, we reduce their width to 22 cm (60% of FoV at 50cm distance). Once the window position is determined, we set each window's orientation to face the user's point of view. Finally, based on results from study 1, we correct the window viewing distances. Since the right-shoulder focus causes some of the windows on the user's left to be placed uncomfortably close, we displace windows along the line of sight so each is a minimum of 50 cm viewing distance (Figure 8).

Participants, Task and Procedure

We recruited 12 university students (3 female, $21 \leq$ age ≤ 40) from a local campus. Participants are presented with a set of windows showing everyday applications, representing ones that might be used on a real HWD. The goal is to scan the windows for information needed to answer a question (Figure 9). The windows present all of the information required to select the correct answer, thus the participant must navigate *between* windows, but need not pan or scroll within the applications themselves.

An example task goes as follows: the participant begins a trial by pressing the button on the Start window, triggering the appearance of icons on the Question and Map windows. The participant navigates to the Question window to find out who he is meeting. Next, he finds the message next to that person in the Messages window. It looks like he is meeting for pizza, so he navigates to the Map window to locate the pizza icon marked with the letter 'a'. Finally, he returns to the Question screen to select the correct answer, 'a', ending the trial.

There are two question types, one with 4 applications (Start, Question, Messages, Map), as in the example, and a second

Figure 9. Example of the application windows presented to participants in study 4.

type that requires the participant to navigate 5 applications (Start, Question, My Contacts, Calendar and Map). The applications are randomly placed among empty desktop windows within an array of either 9 or 16 windows. The windows are laid out in space according to our Personal Cockpit design and the user switches applications by moving his head (Figure 10a).

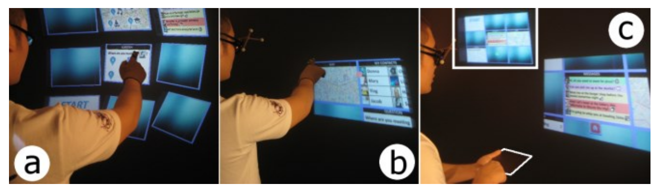

Figure10. Study 4 tested our design (a, shown without FoV constraint for demonstration) against techniques using direct (b) and indirect input (c) with view-fixed displays.

In addition to our Personal Cockpit design, participants must navigate using two baseline techniques with view-fixed displays: one with direct input and the other with indirect input (Figure 10b-c). In these techniques, the same application windows are arranged in a flat array, but the participant can only see those that fit within the 40° FoV. With the direct input technique, the user switches applications by panning the intangible, view-fixed surface (Figure 10b). This technique is analogous to panning on a typical smartphone. To assist direct input, we provide visual feedback to indicate whether the reaching finger is above, on, or behind the window surface. Based on previous work showing difficulties with depth judgement [10] and pilot testing, we provide a substantial 'surface' depth of 8cm.

The indirect technique uses a wireless trackpad, with which participants control a cursor on the view-fixed display (Figure 10c). To switch applications, the participant must select a home icon at the bottom of the display, which leads to an overview of the entire array (c, inset). From the overview, he can select any window in the array, which brings the corresponding window into full view. This technique is similar to the menu interface on some existing HWDs (i.e. Moverio). For consistency, all application windows are sized to 22 cm width and placed at 50cm viewing distance for both view-fixed techniques.

Design

We use a 3 × 2 × 2 within-participants design: *technique* (*PC* – Personal Cockpit with direct input*;* *VD* – view-fixed with direct input; or *VI* – view-fixed with indirect input); *complexity* (*3×3* or *4×4* array of virtual windows) *and question type* (type *I* or *II*). Within each technique, participants completed 4 sets of questions, 1 for each combination of *complexity* and *question type*. For each new set, applications were moved to new random window locations, but with a minimum of one application for each row and column in the layout array. Each set of 4 questions was completed using the same window layout. *Techniques* and *complexities* were fully balanced between participants. Type I questions always preceded type II.

Results

Trial time was measured as the duration between the task start and the selection of the correct answer. We collected 3 techniques×2 complexities×2 question types×4 questions × 12 participants = 576 data points. Of these we removed 24 outlier trials (4.17%). The mean time was 19.91s. Conditional means are shown in Figure 11.

Participants completed the trials significantly faster ($F_{2,22}$=94.845, p<.001) using *PC* (mean 13.57 s) than either of the view-fixed techniques (23.73 s for *VD* and 23.45 s for *VI*, Figure 11). Post-hoc tests with Bonferroni corrections showed significant differences between techniques *PC* vs. *VD* and *PC* vs. *VI* (both pairs p<.001), but not between *VD* vs. *VI* (p=.547).

The univariate ANOVA revealed significant effects of *complexity* ($F_{1,11.187}$ =39.937, p<.001) and *question type* ($F_{1,11.051}$ =11.143, p<.01). The simpler *3×3* complexity had a mean time of 18.33s while the *4×4* trials averaged 22.15s. Question type *I* was also faster than type *II* (18.65 vs. 21.82s). We also found interaction effects (Figure 11) between *technique* × *complexity* ($F_{2,22}$ =5.976, p<.01) and *technique* × *question type* ($F_{2,22}$=3.747, p<.05).

Fatigue: We collected subjective ratings of *arm fatigue* and *neck fatigue* for each combination of *technique* and *complexity*. Means are shown in Figure 12. We ran 3×2 univariate ANOVAs for both arm and neck fatigue. For *arm fatigue*, the test revealed significant effects of *technique* ($F_{2,22}$ =22.045, p<.001) and *complexity* ($F_{1,11.090}$ =7.510, p<.05). Post-hoc tests with Bonferroni corrections show differences between PC vs. VD and VD vs. VI (p<.05 for both pairs). There was also a significant interaction between *technique* and *complexity* ($F_{2,20}$ =2.761, p<.05). For *neck fatigue*, there was a main effect of *complexity* ($F_{1,11.168}$ =8.822, p<.05) but not technique ($F_{2,22.405}$ =7.334, p=.055).

DISCUSSION

Our final study shows potential for the Personal Cockpit as an alternative to interaction with view-fixed displays. Our technique is more efficient than both tested view-fixed techniques and less tiresome than direct input on a view-fixed display. Of the 12 participants in our study, 10 chose

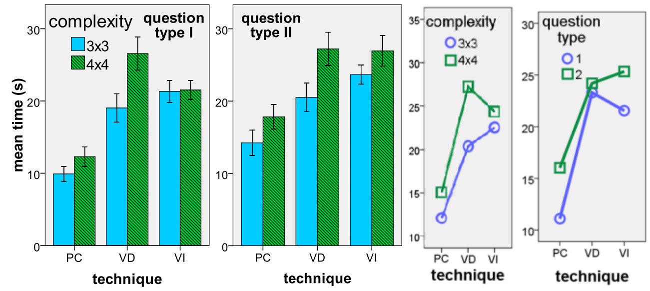

Figure 11. Mean time for study 4 by *technique*, *complexity* and *question type*. Bars show ±2 SE (left). Interaction effects for *technique* × *complexity* and *technique* × *question type* (right).

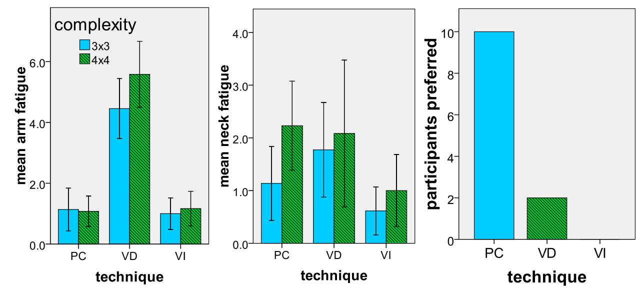

Figure 12. Mean Borg scale ratings for perceived *arm fatigue* (left) and *neck fatigue* (middle). 10 of 12 participants (83.3%) preferred the PC technique (right). Bars show ±2 SE

the Personal Cockpit as their preferred technique in a post-study questionnaire (Figure 12). Several participants commented that the Personal Cockpit was "*easy to navigate*". One participant said, "*I liked the speed of navigation - I was able to move around quickly and in such way it reduced the amount of work.*" Others mentioned that it was "*productive*" and "*the most natural*".

The Personal Cockpit is also scalable. Whereas the panning technique (*VD*) shows a large increase in time with a greater number of application windows (Figure 11), the Personal Cockpit shows only a small increase, as with the indirect method (VI). Despite the use of direct input and necessity of head motion, participants rated the Personal Cockpit on par with the indirect interaction technique (*VI*).

Our results are positive but come with some limitations. We tested only 2 baseline techniques. Although faster untested navigation techniques may exist (e.g. joystick or trackball), these may not have all of the advantages of the Personal Cockpit (i.e. unoccluded forward view; facilitates both navigation between windows *and* interaction with window contents). Further study with additional tasks is required for generalization, however our results are in line with those of prior research [6, 9, 22]. Further studies with actual HWD hardware are required for ecological validity.

Personal Cockpit Interactions

To further illustrate the utility of the Personal Cockpit, we created several interactive demonstration concepts. We implemented these in our CAVE emulation to show how the Personal Cockpit might be useful in real-life scenarios.

Window Overview: Although the Personal Cockpit user can access many applications quickly and easily, there may be times when an overview (also known as a 'World in Miniature' in VR [36]) of all open windows is useful. With a command gesture, the user can shrink the Cockpit layout into a palm-sized sphere (Figure 13a), which fits easily into view. Attached to the non-dominant hand, the user can manipulate the sphere for convenient viewing.

Changing Frames of Reference: The Personal Cockpit is as mobile as the HWD device and can be designed to follow the user on the go with a body-fixed reference frame. When at work or at home, the Cockpit windows can be fixed to a wall or other available space. In this demo, a user can switch between a flat, world-fixed layout and a curved, body-fixed layout with a tap on the HWD (Figure 13b).

Manual Arrangement: Our Cockpit design in Study 4 demonstrates a customized automatic layout. Depending on the situation, the user may want to rearrange the windows manually. In this demonstration, the user can grab, move and resize windows at his leisure using in-air pinching gestures. To open a new application window, the user grabs an icon from a task-launcher window and places it in an empty location, where a new window springs into view (Figure 13c). A hand-fixed reference frame is convenient for bimanual interaction techniques.

Window Intercommunication: In multitasking situations, two or more windows may be tied to the same task. For instance, many tasks can have peripheral information or tool palettes in a secondary display. We demonstrate a colour-picker tool, in which the user can select a colour to tint a photo in a separate application window (Figure 13d).

Translating the Personal Cockpit to a HWD
Our next step is to demonstrate that Personal Cockpit's design advantages transfer to a real-world HWD. As display and tracking technologies advance, systems will be able to support fully-mobile implementations. We outline some important challenges for this realization.

Body-Fixed Stabilization: In Study 2, we found that naïvely fixing windows to body-fixed coordinates leads to selection inaccuracies with slight body motions. Based on this finding, we envision a hybrid between world- and body-fixed reference frames for mobile use. When the user is standing still, the layout becomes purely world-fixed. When he begins moving again, the system detects this and brings the Cockpit along. Other approaches include using a low-pass filer to smooth and stabilize motion.

Pseudo-Occlusion: An important problem we discussed earlier (Direct Input, in Design Factors section) is that a HWD lies between the viewer and the input space. This causes the display image to occlude any outside objects, including the user's hands. We propose the concept of *pseudo-occlusion* to solve this. The system would accurately track the hands' position in nearby space. When the hand is placed between the HWD and a virtual window, the system subtracts the interfering region from the rendered window, making it appear occluded by the hand.

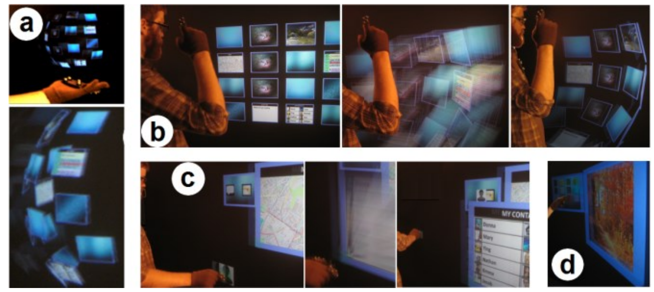

Figure 13. Personal Cockpit interactions scenarios: Changing from world-fixed to body-fixed layout (a); opening a new application window (b); window intercommunication (c); and a shrinking the Cockpit to a palm-sized overview (d).

Transfer to Wearable Technology: Our emulation of the FoV limitation in a CAVE environment provided us with several advantages in implementation and tracking. Further research is required to discover the limitations of applying a functional Personal Cockpit interface on current and near-future hardware with variations such as different image plane distances. We also must answer questions about the effectiveness of transparent displays in real situations, such as with objects moving in the background or while walking.

Lessons Learned
We take away the following lessons from our investigation: 1) The spatial multi-window layout of our design allows fast task switching, requiring only 60% of the time of the 2 tested view-fixed interaction techniques. 2) Virtual windows are compatible with direct input, even with a limited FoV. Windows can be placed as close as 50cm, even with a 1m distant virtual image plane. 3) Body-fixed reference frames are subject to higher targeting error than world-fixed windows, due to unintentional perturbations caused by reaching motion. 4) A curved layout is subject to lower error and arm fatigue when offset to align with the dominant limb. 5) The Personal Cockpit is scalable within reasonable limits. Greater window offset angles, and thus greater window numbers, lead to increased head motion. This can lead to longer task switching times.

CONCLUSION AND FUTURE WORK
We have explored the design space of the Personal Cockpit, a design concept for fast task switching on head-worn displays. We refined our design based on the outcomes of 3 user studies and tested our final design against techniques using direct and indirect interaction with view-fixed displays. We work towards a window management system for HMDs by demonstrating a set of interactions using the Personal Cockpit. We lay out several challenges that must be addressed in translating our design to a wearable device. In future, we plan to move forward by implementing a wearable prototype system. We will also continue exploring techniques for multi-window management [3, 4] as well as features for enabling direct input, by better understanding limitations of perception [10] and user fatigue [17].

ACKNOWLEDGEMENTS

This project was supported in part by NSERC.

REFERENCES

1. Ankrum, D.R. Visual ergonomics in the office – Guidelines. *Occupational Health & Safety, 68, 7* (1999), 64-74.
2. Ball, R. and North, C. The effects of peripheral vision and physical navigation on large scale visualization. *Proc. GI '08*, (2008), 9-16.
3. Bell, B.A. and Feiner, S.K. Dynamic space management for user interfaces. *Proc. UIST '00*, ACM (2000), 239-248.
4. Bell, B., Feiner, S. and Höllerer, T. View management for virtual and augmented reality. *Proc. UIST '01*, ACM (2001), 101-110.
5. Billinghurst, M., Bowskill, J., Jessop, M. and Morphett, J. A wearable spatial conferencing space. *Proc. ISWC '98*, IEEE (1998), 76-83.
6. Billinghurst, M., Bowskill, J., Dyer, N. and Morphett, J. An evaluation of wearable information spaces. *Proc. VRAIS '98*, IEEE (1998), 20-27.
7. Böhmer, M., Hecht, B., Schöning, J., Krüger, A. and Bauer, G. Falling asleep with Angry Birds, Facebook and Kindle – A large scale study on mobile application usage. *Proc. MobileHCI '11,* (2011), 47-56.
8. Cao, X. and Balakrishnan, R. Interacting with dynamically defined information spaces using a handheld projector and a pen. *Proc. UIST '06*. ACM (2006), 225-234.
9. Cauchard, J., Löchtfeld, M., Fraser, M., Krüger, A. and Subramanian, S. m+pSpaces: Virtual workspaces in the spatially-aware mobile environment. *Proc. MobileHCI '12*, ACM (2012), 171-180.
10. Chan, L.W., Kao, H.S., Chen, M.Y., Lee, M.S., Hsu, J. and Hung, Y.P. Touching the void: Direct-touch interaction for intangible displays. *Proc. CHI '10*, ACM (2010), 2625-2634.
11. Czerwinski, M., Tan, D.S., and Robertson, G.G. Women take a wider view. *Proc.CHI '02*, ACM (2002), 195-202.
12. Feiner, S. MacIntyre, B., Haupt, M. and Solomon, E. Windows on the world: 2D windows for 3D augmented reality. *Proc. UIST '93*. ACM (1993), 145-155.
13. Fitzmaurice, G.W. Situated information spaces and spatially aware computers. *Communications of the ACM 36, 7*, ACM (1993), 39-49.
14. Grudin, J. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. *Proc. CHI '01*, ACM (2001), 458-465.
15. Ha, V., Ziola, R., Wallace, J. and Inkpen, K. My MDE: Configuring virtual workspaces in multi-display environments. *CHI '06 Extended Abstracts*. ACM (2006), 1481-1486.
16. Hezel, P.J. and Veron, H. Head mounted displays for virtual reality. *Proc. SPIE, 2218* (1994), 41-50.
17. Hincapié-Ramos, J.D., Guo, X., Moghadasian, P. and Irani, P. Consumed Endurance: A metric to quantify arm fatigue of mid-air interactions. *Proc. CHI '14*, ACM (2014), 10 pages.
18. Hoffman, D.M., Girshick, A.R., Akeley, K. and Banks, M.S. Vergence-accomodation conflicts hinder visual performance and cause visual fatigue. *Vision, 8, 3, Art. 33* (2008), 30 pages.
19. Kang, Y. and Stasko, J. Lightweight task/application performance using single versus multiple monitors: A comparative study. *Proc. GI* '08, CIPS (2008), 17-24.
20. Laster SmartVision. www.laster.fr
21. Levia, L.A., Böhmer, M., Gehring, S. and Krüger, A. Back to the app: The costs of mobile application interrruptions. *Proc. MobileHCI '12*, ACM (2012), 291-294.
22. Li, F.C.Y, Dearman, D. and Truong, K.N. Virtual shelves: Interactions with orientation aware devices. *Proc.UIST '09*, ACM (2009), 125-128.
23. Lin, J., Duh, H. Parker, D., Abi-Rached, H and Furness, T. Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment. *Proc. VI '02*, *EEEI* (2002), 164-171.
24. Mayer, T., Brady, S., Bovasso, E., Pope, P. and Gatchel, R.J. Noninvasive measurement of cervical tri-planar motion in normal subjects. *Spine*, 18, 15 (1993), 2191-2195.
25. NASA-STD-3000 Man-Systems Integration Standards. Rev B, 1995.
26. Ni, T., Bowman, D.A, and Chen, J. Increased display size and resolution improve task performance in information-rich virtual environments. *Proc. GI '06*, CIPS (2006), 139-146.
27. Noble, B.J., Gunnar, A.V.B., Jacobs, I., Ceci, R. and Kaiser, P. A category-ratio perceived exertion scale: Relationship to blood and muscle lactates and heart rate. *Medicine and Science in Sports and Exercise, 15, 5* (1983), 523-528.
28. Patterson, R., Winterbottom, M. and Pierce, B. Perceptual issues in the use of head-mounted visual displays. *Human Factors, 48, 3*, 555-573.
29. Rashid, U., Nacenta, M.A. and Quigley, A. The cost of display switching: A comparison of mobile, large display and hybrid UI configurations. *Proc.AVI '12*, ACM (2012), 99-106.
30. Rekimoto, J. and Nagao, K. The world through the computer: Computer augmented interaction with real world environments. *Proc. UIST '95*, ACM (1995), 29-36.
31. Robertson, G. van Dantzich, M, Robbins, D., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D. and Gorokhovsky, V. The task gallery: A 3D window manager. *Proc. CHI '00*, (2000), 494-501.
32. Shneiderman, B. Direct manipulation: a step beyond programming languages. *IEEE Computer, 16, 8* (1983), 57-69.
33. Shibita, T., Kim, J., Hoffman, D.M. and Banks, M.S. The zone of comfort: Predicting visual discomfort with stereo displays. *Vision, 11, 8, Art. 11* (2011), 29 pages.
34. Shupp, L., Ball, R., Yost, B., Booker, J. and North, C. Evaluation of viewport size and curvature of large, high-resolution displays. *Proc. GI '06*. CIPS (2006), 123-130.
35. St. John, M., Harris, W. and Osga, G.A. Designing for multi-tasking environments: Multiple monitors versus multiple windows. *Proc. HFES*, SAGE (1997), 1313-1317.
36. Stoakley, R., Conway, M.J. and Pausch, R. Virtual reality on a WIM: Interactive worlds in miniature. *Proc. CHI '95*, ACM (2005), 265-272.
37. Su, R.E. and Bailey, B.P. Put them where? Towards guidelines for positioning large displays in interactive workspaces. *Proc. INTERACT '05*, Springer (2005), 337-349.
38. Tak, S., Cockburn, A., Humm, K., Ahlström, D., Gutwin, G. and Scarr, J. Improving window switching interfaces. *Proc. INTERACT '09*, Springer (2009), 187-200.
39. Tan, D. and Czerwinski, M. Effects of visual separation and physical discontinuities when distributing information across multiple displays. *Proc. INTERACT'03*, (2003) 252-260.
40. Toet, A., Jansen, S.E.M., and Delleman, N.J. Effects of field-of-view restrictions on speed and accuracy of manoeuvring. *Perceptual and Motor Skills, 105* (2008), 1245-1256.
41. Treisman, A. M. and Gelade, G. A feature-integration theory of attention. *Cognitive psychology, 12* (1980), 97–136.
42. Valkov, D., Steinicke, F., Bruder, G., Schöning, J., Daiber, F. and Krüger, A. Touching floating objects in projection-based virtual reality environments. *Proc. EGVE-JVRC '10*, Eurographics (2010), 17-24.
43. Willemsen, P., Colton, M.B., Creem-Regehr, C. and Thompson, W.B. The effects of head-mounted display mechanical properties and field of view on distance judgments in virtual environments. *Applied Perception, 6, 2*, *Art. 8* ACM (2009), 14 pages.
44. Yang, X.D., Mak, E., McCallum, D., Irani, P. Cao, X. and Izadi, S. LensMouse: Augmenting the mouse with an interactive touch display. *Proc. CHI '10*, ACM (2010), 2431-2440.
45. Yee, K. Peephole Displays: Pen interaction on spatially aware handheld computers. *Proc. CHI '03*, ACM (2003), 1-8.