Ethereal Planes: A Design Framework for 2D Information Spaces in 3D Mixed Reality Environments

Barrett Ens
University of Manitoba
Winnipeg, Canada
bens@cs.umanitoba.ca

Juan David Hincapié-Ramos
University of Manitoba
Winnipeg, Canada
jdhr@cs.umanitoba.ca

Pourang Irani
University of Manitoba
Winnipeg, Canada
irani@cs.umanitoba.ca

ABSTRACT
Information spaces are virtual workspaces that help us manage information by mapping it to the physical environment. This widely influential concept has been interpreted in a variety of forms, often in conjunction with mixed reality. We present Ethereal Planes, a design framework that ties together many existing variations of 2D information spaces. Ethereal Planes is aimed at assisting the design of user interfaces for next-generation technologies such as head-worn displays. From an extensive literature review, we encapsulated the common attributes of existing novel designs in seven design dimensions. Mapping the reviewed designs to the framework dimensions reveals a set of common usage patterns. We discuss how the Ethereal Planes framework can be methodically applied to help inspire new designs. We provide a concrete example of the framework's utility during the design of the Personal Cockpit, a window management system for head-worn displays.

Author Keywords
Information spaces; mixed reality; design framework; head-worn displays; spatial user interfaces

ACM Classification Keywords
H.5.2 **Information Interfaces and Presentation]**: User Interfaces – Theory and methods

INTRODUCTION
The recent proliferation of low-cost yet robust display and sensing technologies is opening the door to new paradigms for everyday computing. Displays and sensors are quickly becoming small and lightweight enough for wearable applications while approaching benchmarks in latency and fidelity that make them practical. Similar to the shift from mouse and keyboard toward the more intuitive paradigm of direct touchscreen manipulation, we now foresee the widespread adoption of spatial interaction and mixed reality for everyday information management in platforms such as head-worn displays (Figure 1). Yet these platforms are still in their relative infancy and there is a lack of methodological tools to support the design of everyday applications.

B. Ens, J.D. Hincapié-Ramos and P. Irani. Ethereal Planes: A Design Framework for 2D Information Spaces in 3D Mixed Reality Environments. In *SUI '14: Proceedings of the 2nd symposium on Spatial user interactions*, 11 pages, to appear, ACM, 2014.

© ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version will be published in *SUI 2014, October 4–5, 2014, Honolulu, USA*.

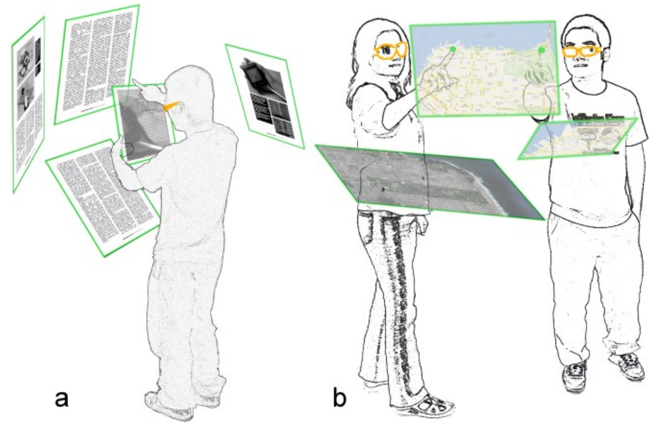

Figure 1. Our design framework, Ethereal Planes, facilitates the classification and comparison of designs that use 2D information spaces in 3D mixed reality environments. Analysis techniques can inspire the construction of new designs. Informed decision-making is an important step toward advanced productivity features for multitasking (a), analytic reasoning and co-located collaboration (b).

In this paper we aim to assist the design process by collecting and organizing concepts introduced and explored in previous research endeavors. Based on a systematic literature review, we present a design framework we call Ethereal Planes. Ethereal Planes describes the design space of planar (2D) interfaces in 3D mixed reality environments. We focus on 2D designs because they are familiar [30,36], intuitive [23], and have advantages in efficiency, speed, precision and reduction of clutter [15,16,52]. While there are many instances where 3D interfaces will prove useful, 2D interfaces are currently ubiquitous both within and beyond the realm of computing interfaces and will remain suitable for a wide range of uses, particularly those involving information simplification or abstraction (e.g. text, floor plans, control panels).

Ethereal Planes employs the concept of information spaces [24] in assisting the design of advanced and productive interfaces. Information spaces support intuitive computing interaction by mapping information to real world space, allowing us to look beyond the boundaries of the computing device and perceive information where it belongs – in the surrounding environment. Information spaces have been implemented in diverse platforms including spatially-aware handheld devices, personal projectors [12,67] , tabletops [59] and digital paper [58]. Ethereal Planes is primarily aimed at supporting interface design on head-worn displays (HWDs) [6,22], which due to their wearable nature are always-available and hands-free, in a way not possible with previous technologies. Ethereal Planes is intended for interaction designers of mixed-reality HWDs applications.

Ethereal Planes was derived from a systematic literature review of information spaces with 2D instantiations. We encapsulate the recurring design themes into seven design dimensions. By analyzing common design choices from existing implementations we identified common design patterns. Further, we discuss several analysis techniques (e.g. tweaking, combining) that can help inspire new designs, and discuss our own use of the framework in the design of a system called the Personal Cockpit [3].

BACKGROUND

Our goal in defining Ethereal Planes is to support the design of user interfaces for emerging HWD technologies. However, we look beyond the individual technical challenges of these novel technologies towards a framework to encourage the development of everyday user interfaces for everyday applications. We encourage new and useful designs by providing a unifying foundation for the description and categorization of tools needed for manipulating spatially distributed information. In this section we introduce the concepts of design frameworks and mixed-reality technologies.

Design Frameworks

Design frameworks are conceptual tools created to help designers conceptualize the nuances of particular technologies and formalize the creative process. Design frameworks have an established history in interface design, and have shown their value in providing terminology to categorize ideas [50] and organize complex concepts into logical hierarchies [46]. Design frameworks often accompany either the introduction of a previously unexplored concept (e.g. Graspable User Interface [25]) or the exploration of existing work in a new light (e.g. Ambient Information Systems [49], Availability Sharing Systems [35], and Ephemeral User Interfaces [20]).

Several frameworks related to spatial and mixed reality interactions have previously been developed for immersive virtual environments. For example, Bowman and Hodges 8 describe a framework outlining techniques for virtual navigation. Poupyrev et al. [48] present a taxonomy of virtual object manipulation techniques. Mine et al. [44] introduce a framework to leverage proprioception to assist interaction with virtual objects. Also, a well-known survey by Hinckley et al. [36] discusses many general issues relevant to spatial user interaction. In contrast to these previous frameworks, Ethereal Planes specifically addresses interface design for 2D, mixed reality information spaces and draws from work developed for a wide variety of mixed reality platforms.

In creating Ethereal Planes we used techniques also applied to HWD interface design by Robinett [54] and similar to those formalized in Zwicky's General Morphological Analysis [53]. This method treats a set of defined taxonomical terms as a set of orthogonal *dimensions* in a geometric *design space*. The resulting theoretical matrix provides a structure for objective classification and comparison. The methodical filling-in of this structure helps to categorize existing concepts, differentiate ideas, and identify unexplored terrain. In summary, there are three basic steps in the development and usage of our design framework, which we follow through the course of this paper:

1. Review of existing designs to distill a set of characteristic dimensions
2. Categorization of existing designs among these dimensions to identify both gaps and common usages
3. Generation of new designs through an analytic process of combining and altering design choices

Along these steps, our Ethereal Planes framework fulfills several purposes: The distillation from existing literature of a set of general but widely encompassing design dimensions provides a taxonomy for designers, researchers, teachers and students to express their creations. The dimensional organization also helps the understanding of existing designs by providing a means to categorize them; by contrasting and comparing these, designers gain insight into general patterns and identify gaps in the dimensional framework where designs do not yet exist. Designers can then use this information to assist with the creation of new designs, either by applying the strengths of existing patterns to the correct contexts or thorough experimentation, by altering one or more dimension and then imagining the resulting implications.

Mixed Reality Technologies

Mixed reality, the combination of real and virtual objects, has its roots in the see-through HWD technology introduced by Sutherland [60]. Buxton and Fitzmaurice [11] identified three potential platforms for realizing information spaces: Caves, HWDs and handheld devices. These technologies, and more recently, projection, have since have since become staples of mixed reality. These methods cover the breadth of visual output platforms that surface in our literature review.

Each of these technologies has its advantages and limitations. Caves can produce high-fidelity immersive environments, but size and cost restricts them from common use. HWDs are recently available in lightweight form factors, both monocular [27] and stereoscopic [9,63]. The latter hold promise for mixed reality due to their capability for producing convincing 3D effects similar to those available in a Cave environment. Moreover, HWDs possess an advantage over Caves in their capability to produce different perspectives of the same object for multiple viewers 1. Handheld devices are now ubiquitous, making them a popular target platform, but only serve as a small window to virtual content (e.g. [68]). Projectors are also becoming popular with the advent of compact portable versions (e.g. [12,40]). Projectors are spatially less restrictive than handhelds, but require an external surface for projection.

We created the Ethereal Planes framework primarily for the design of next-generation HWD interfaces. The potential versatility and affordance for mobility of HWDs, along with support of integrated sensors [47,56] for sophisticated user input (e.g. mid-air gestures), makes these devices a promising future ubiquitous mixed-reality platform.

ETHEREAL PLANES FRAMEWORK

The foundation of our Ethereal Planes design framework is an organizational taxonomy for classifying designs that incorporate virtual 2D workspaces.

Research Method

The taxonomy was the product of an extensive review of literature related to information spaces, and spatial interaction. Within this body of work, we found a subset of designs that embody the concept of Ethereal Planes. We began with a thorough archive search for papers exploring spatial user interfaces that occupy real world space, extending or existing fully beyond the limits of a conventional display screen. We focused on designs involving planar information spaces thus excluded designs that do not explicitly discuss 2D workspaces, for example those that involve navigating 3D workspaces through a 2D display. We also excluded papers that do not introduce distinct differences from previous designs, for example the use of an existing design in a new context or focus on the technology for implementing a known design. To begin, we manually sifted through the previous 5 years' proceedings of CHI, UIST, ISWC and VRST. We also conducted a tree search of references and citations of the initial papers we identified and of seminal papers on spatial interaction frameworks (e.g. [8,36,44,48]). The final list, containing 34 papers, is not intended to be exhaustive, however represents a diverse selection of designs from which we draw. (A complete list of all 34 designs in our survey, along with their dimensional classifications, may be found on our project page: http://hci.cs.umanitoba.ca/projects-and-research/details/personal-cockpit-spatial-user-interface)

From the papers in our literature review, we distilled a set of design dimension using a bottom up approach resembling open coding. We began with [18] candidate dimensions that fit the concepts found in the reviewed literature, then iteratively reduced these into a set small enough to manage in a concise framework, yet containing enough dimensions to make it useful. We eliminated dimensions, for example, that expressed concepts that we deemed relatively insubstantial (e.g. fidelity), that were later incorporated into other dimensions (e.g. spatial reference frame) or that were substantial enough that treatment in our current framework would be superficial (e.g. co-located collaboration). Several important concepts that deserve further consideration are listed in a later section (Framework Extensions). This process resulted in seven design dimensions, listed in Table 1. We further organized the dimensions into three groups based on the strongest dependencies between them. This grouping is used to organize several resulting design recommendations.

Group	Dimension	Values		
Reference Frame	*Perspective*	egocentric	exocentric	
	Movability	movable	fixed	
Spatial Manipulation	*Proximity*	far	near	on-body
	Input mode	direct	indirect	
	Tangibility	tangible	intangible	
Spatial Composition	*Visibility*	high	intermediate	low
	Discretization	continuous	discrete	

Table 1. Seven dimension of our design framework, their three groups and their potential values.

Design Space Dimensions

Perspective denotes the conceptual viewpoint of the observer. To delineate this dimension, we borrow the terminology of *egocentric* and *exocentric* reference frames, used in early virtual reality literature [65] and later included in a taxonomy for virtual object manipulation by Poupyrev et al. [48]. The exocentric perspective the viewer is an outside observer, whereas the egocentric perspective is immersive. These terms correspond to the sub-divisions of world- and body-based coordinate systems used in other taxonomies, such as that of Cockburn et al. [16]. Feiner et al. [22] expanded these to three possible reference frames for virtual windows, view-fixed, surround-fixed or object-fixed. Billinghurst [6] similarly refers to head-, body- or world-stabilized information displays. Hinckley et al. [36] use the terms relative and absolute gesture to denote motions in body- and world-centric space, respectively. In our framework, *egocentric* reference frames denote 'first person' (body-centric) reference points, such as the head or body, whereas *Exocentric* frames are set relative to any object or other real-world (world-centric) reference point.

Movability denotes whether workspaces are *movable* or *fixed* with respect to a given frame of reference. Fixed workspaces are indefinitely locked in place to their respective coordinate systems. Movable ones can be relocated in relation to their egocentric or exocentric reference point. In most contexts, we consider a hand-fixed information space as *movable* because it can be moved to different coordinate points within the reference fame, whether body- or world-centric. A mobile device display, for example, can be often relocated with respect to the user's head or body, thus does not usually qualify as *fixed*.

Proximity describes the distance relationship between an information space and its user. We use a set of regions drawn from neuropsychology [21,34] also used by Chen et al. [14]: *on-body* (coincides with pericutaneous space, on the body surface), *near* (peripersonal space, within arm's reach) and *far* (extrapersonal space, beyond arm's reach). The majority of implementations we examined involve interaction within arm's reach, often by direct input (e.g. [12]) or with a handheld device (e.g. [68]). Some systems allow interaction with distant objects, particularly those for immersive virtual worlds or for outdoor use (e.g. Augmented Viewport [37]). Other researchers have explored the human body as an interface (e.g. [32]).

Input mode falls coarsely into two camps, *indirect* and *direct*. Indirect input includes cursors, ray-casting and variations of these methods. Direct input includes input using direct touch by hand, fingertip or stylus as well as virtual 'touch' with intangible surfaces (e.g. [13,29]).

Tangibility defines whether an information space is mapped to a surface that can be touched. Our frame work classifies implementations as either *tangible* or *intangible*. *Tangible* interfaces often leverage surfaces in the nearby environment, such as a wall (e.g. [12]) or device screen (e.g. [68]) and benefit from haptic feedback. *Intangible* designs typically make use of 'in-air' gestures (e.g. [29]) for user input.

Visibility describes the amount of visual representation available in an interface and also determines the degree to which spatial memory relies upon proprioception. Our framework uses three levels of visibility, *high*, *intermediate* and *low*. *High* visibility means that the information space is largely or fully visible. *Intermediate* visibility means some type of viewing constraint is present, for instance if only a small section of the workspace may be seen at one time (e.g. [68]). *Low* visibility implies that information management relies very little or not at all on visual feedback (e.g. [29]).

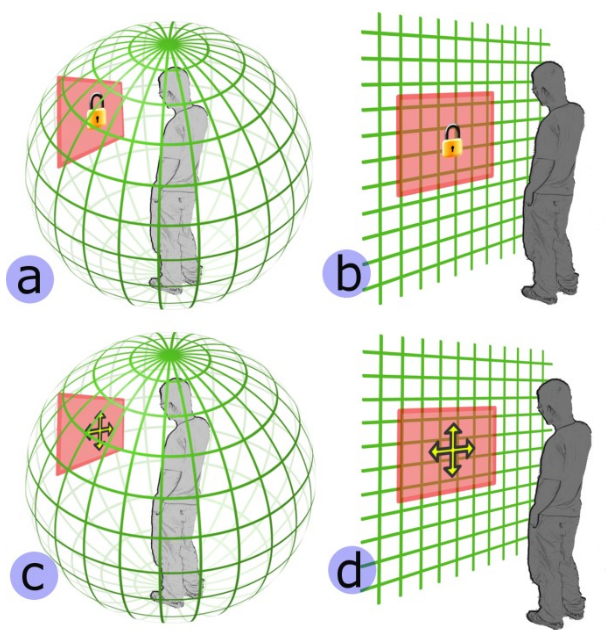

Figure 2. Four general *Reference Frames* for Ethereal Planes: (a) fixed-egocentric, (b) fixed-exocentric, (c) movable-egocentric and (d) movable-exocentric.

Input mode		*direct*		*indirect*
Tangibility		*tangible*	*intangible*	
Proximity	*on-body*	Skinput [32], OmniTouch [31]		
	near	Peephole displays [68], Cao et al. [12]	Touching the void [13], Imaginary interfaces [29]	Sidesight [10], Windows on the world [22]
	far			Virtual shelves [41], Augmented Viewports [37]

Table 2. Example combinations between proximity, input mode and tangibility categories of *Spatial Manipulation*.

Discretization specifies whether an information space is *continuous* or composed of *discrete* units. The majority of designs in our survey use *continuous* space. Examples of *discrete* mappings are the body-centric browser tab mappings described by Chen et al. [14] and the bins Wang et al. [64] placed around a mobile device for sorting photos.

Dimensional Interdependencies

While the dimensions of a design space are ideally orthogonal, dependencies between dimensions are rarely entirely absent. As a case in point, some choices in the Ethereal Planes dimensions will have implications for others. We clustered the dimensions by their closest dependencies into groups we call *Reference Frame*, *Spatial Manipulation* and *Spatial Composition* (Table 1). Here we discuss some of the tradeoffs between design choices within each of the three groups.

Reference Frame – Perspective and *movability* together encompass the concept of a spatial reference frame. Combinations of these two dimensions are summarized in Figure 2. Different reference frames are better suitable for different types of applications. In a mobile scenario, an *egocentric* perspective is more useful, since it will move along with a user on-the-go. In collaborative scenarios, *exocentric* space is more appropriate, since users will benefit from a shared, world-based reference frame, as is the case with a real-world, wall-fixed whiteboard. *Exocentric* frames are also useful for situating information spaces in the contexts where they are most practical [24]. However, in free space interactions, Hinckley et al. [36] note that *egocentric* coordinate systems are easier for users to comprehend and manipulate than *exocentric* frames.

Fixed information spaces are useful in situations where spatial memorability is important, for example in the placement of application shortcuts [41]. Once learned, objects in *fixed* spaces can also be recalled with the aid of proprioception [30,41,68]. *Movable* workspaces, conversely, are better for short-term memorability such as when the information contents are short-term, volatile or highly dynamic.

Spatial Manipulation – The three dimensions of *proximity*, *input mode* and *tangibility* are related to the manipulation of information spaces and of data and objects within them. Table 2 provides examples of relevant combinations between these dimensions. For various reasons, some combinations have no existing counterparts in our Ethereal Planes-related literature. With indirect input, for example, the concept of *tangibility* becomes less relevant, thus we do not include *tangibility* under the indirect column of the table. Conversely, it is difficult to imagine direct input with *far* proximity, thus no examples appear in our survey (although this does not mean that some conception of such a concept cannot be realized in future).

Input mode is dependent on *proximity*: whereas *indirect* input allows interaction with surfaces that are beyond reach, *direct* input is intuitive when the interface lies within reach. *Direct* input is practical with *on-body* surfaces since it leverages proprioception. Leveraging available surfaces, whether body or other, also assists motor precision [42].

Tangibility is influenced by the implementation technology. Projection-based interfaces are often *tangible*, since a projection surface is required. Stereoscopic displays (i.e. Caves, some HWDs) often use *intangible*, virtual surfaces, although information spaces are sometimes intentionally set to coincide with physical surfaces [61]. In free space, researchers have found that *indirect* input is faster, less fatiguing and more stable [2,36,62] than *direct* input. However, *direct* input is intuitive and can make use of expressive gestures, thus may be desirable even without the aid of a *tangible* surface. Our survey turned up many designs using *direct* input both with (e.g. [12,32]), and without (e.g. [13,29]) *tangible* surface contact.

Spatial Composition – Together, *visibility* and *discretization* contribute to the way information is organized spatially. One important factor related to these dimensions is spatial memory. Spatial memory is important in many of the interface designs considered in our survey, particularly when the information spaces are not confined within the boundaries of a typical display screen (e.g. [68]). Table 3 shows examples of different pairings between *visibility* and *discretization*. The majority of interfaces represent information visually, however some present little or no visual information. Spatial memory can be built either purely visually, or by muscle memory, although many designs leverage some combination of both (e.g. [32,68]). Designs with little or no visual feedback are more likely to rely highly on proprioception for object recall (e.g. [29,41]). *Discrete* spatial mappings are commonly used with interfaces with *intermediate* or *low* visibility. When little or none of the interface can be seen, designers can instead leverage spatial memory or proprioception, (e.g. Virtual Shelves [41]). In such cases, *discretization* is often leveraged to make recall manageable.

Visibility \ ***Discretization***	*continuous*	*discrete*
low	Imaginary interfaces [29]	Virutal shelves [41], Piles across space [64], mSpaces [17], body-centric browser tabs [14]
intermediate	Peephole displays [68]	Skinput [32], Chameleon [26]
high	Pen light [57], Mouse light [58]	

Table 3. Example pairings between the visibility and discretization categories of *Spatial Composition*.

FRAMEWORK APPLICATIONS

We created our Ethereal Planes framework to guide our own research and also to assist future designers. Here we discuss how our framework can be used to categorize and compare existing designs as well as aid the creation of new designs.

Categorizing Existing Designs

A fundamental aspect of any framework is its descriptive capacity. To show how Ethereal Planes can be used to describe existing designs, we apply it to the works from our literature review. For each design, we assigned dimensional values and classified the results, which provides us with a methodical system to contrast and compare these different designs. We acknowledge that our framework does not provide an absolute partitioning in which designs fit cleanly into the dimensional values. Rather there are many cases where different values apply to multiple presented concepts or the chosen values are open to interpretation. However, the goal of our framework is not to provide a set of arbitrary sorting bins, but to make the designer aware of important design choices and help them weigh the potential benefits of these choices.

Several distinct categories of similar designs emerged from our analysis, each of which we describe in detail below. Although these five categories represent only a small geometric region of the full design space, we found that the majority of reviewed designs (30 of 34) are a very good fit to one of them. As with the assignment of dimensional values, these categories are not absolute, thus we include minor variations that fit closely to the overall character of the group. A few more diverse exceptions are discussed at the end of this section and in section.

Peephole – In the first and largest of our categories, we group concepts that build on the *spotlight* and *peephole* metaphors. These designs allow interaction through 'peephole windows' that are moved around the surface of a 2D workspace. Both are conceptually similar with their main difference being the technology used: Whereas *peephole* interaction implies the use of spatially aware mobile devices, the *spotlight* metaphor typically refers to projection-based environments. The common moniker of 'peephole' interaction was coined by Yee [68], but is a direct descendant of Fitzmaurice's Chameleon. The common theme motivating these designs is to expand the workspace beyond the limited boundaries. To prevent getting lost in a large, mostly invisible space, the workspace remains world-fixed while the device user navigates the content within. Whereas the original Chameleon 26 implementation used the *discretized* space of a spreadsheet application, most variations use *continuous* 2D space. Several other variations, not discussed here, explore 2D 'image-plane' representations of 3D space. Variations from our research include: Touch Projector [7], mSpaces [17], Chameleon [26], Pass-them-around [43], Peephole displays [68], dynamically defined information spaces [12], PenLight [57], MouseLight [58], Augmented Surfaces [51], PlayAywhere [66], Lightspace [67], Bonfire [39] and X-Large virtual workspaces [40].

Floating – This group contains various instantiations of virtual windows that appear to *float* in mid-air. A common goal of these designers is to import the familiar characteristics of ubiquitous 2D applications into an immersive environment. *Floating* windows have often been used to implement auxiliary input controls such as panels, dialog boxes and menus, in immersive virtual reality environments 18. Since mid-air displays are *intangible*, designers often use *indirect* input modes such as mice [22,37] or ray-casters [2]. Chan et al. [13] provide an interesting exploration of *direct* interaction with *intangible* displays. Other variations include: Windows on the World [22], Wearable Conferencing Space [6], Friction Surfaces [2] and Augmented Viewport [37]. Most of these implementations use *exocentric* information spaces, however some HWD implementations [6,22] provide the option of *egocentric* floating windows for mobile users.

Off-Screen – This category includes designs that allow *indirect* input in the 'off-screen' region that surrounds a device's periphery. As in the peephole concept, *off-screen* designers address the problem of limited screen space by extending the theoretical plane of a device's screen into surrounding space. However, these systems are easily portable, allowing the surrounding workspace to be conveniently repositioned. They also avoid occlusion with *indirect* input, and are useful for navigational operations such as panning and zooming. We generalize this category as *exocentric* because two of the included designs (SideSight [10] and Portico [4]) use a device placed on a surface. However, the third example (off-screen pan and zoom [38]) is *egocentric*, since it uses a handheld device.

On-body – Another convenient tangible surface is the human *body*, used by the designs in this category. In many instances, a hand or arm doubles as a convenient projection surface in lieu of a wall or table, and is a convenient, always-available place to store buttons or task shortcuts. Body parts have the primary benefit of assisting target acquisition with proprioception, as evidenced in Harrison et al.'s Skinput [32]. Variations on this theme include

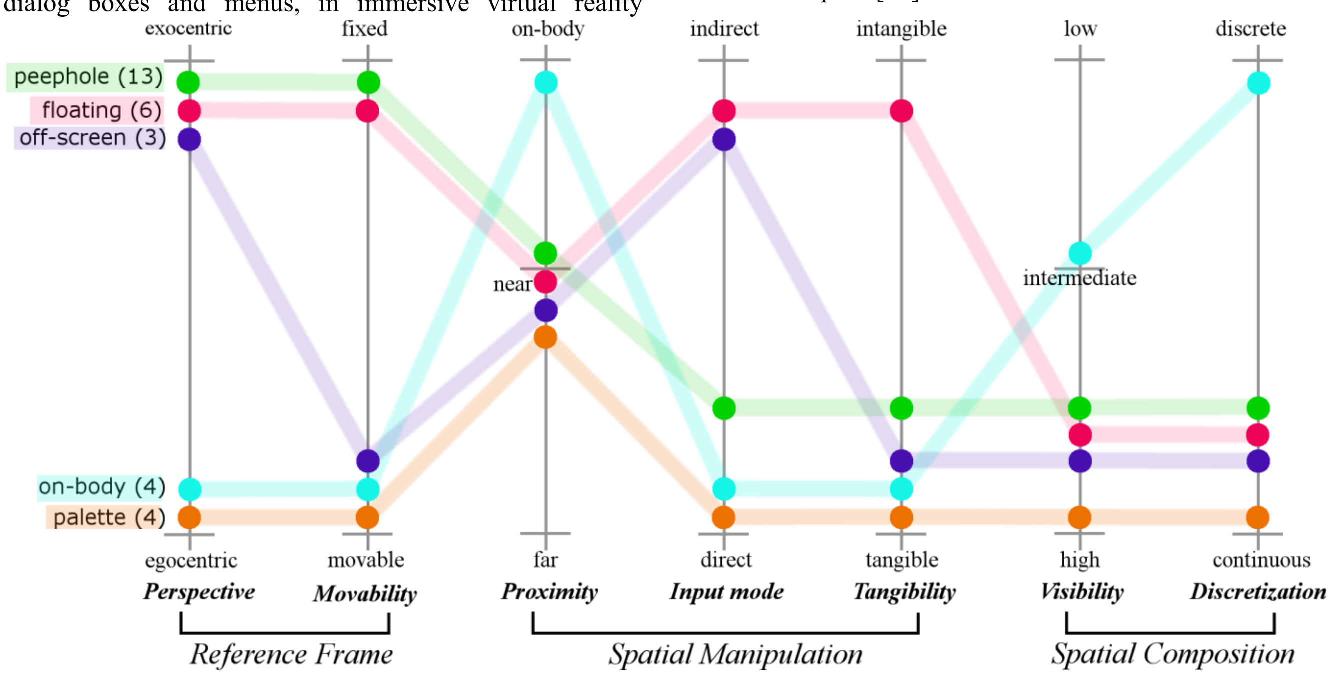

Figure 3. A parallel coordinates graph showing the main design categories found in our analysis of existing designs. Each category is plotted along the seven dimensions of the Ethereal Planes framework. (Best viewed in colour)

Imaginary Phone [30], OmniTouch [31] and Chen et al.'s Body-centric prototype [14].

Palette – These designs align the information space with a handheld *palette*, such a paddle or transparent sheet. This use of a handheld plane allows bimanual interaction, which can facilitate task performance [42]. Handheld tangible surfaces have commonly been used in immersive environments since *tangible* surfaces provide increased speed and control over *intangible* floating surfaces [42]. Variations include the Personal Interaction Panel [61] and various similar implementations [19,42,55].

In Figure 3 we provide a visual summary of the major design categories in a parallel coordinates graph. This graph shows the values of each category along the seven design dimensions. This figure fulfills several purposes: 1) It enables easy comparison between the patterns, revealing where they are similar and where they differ. 2) It shows clustering within the dimensions, including commonly occurring values (e.g. *near proximity* - *high visibility*) and commonly joined pairs (e.g. *exocentric-mixed* - *direct-tangible*). 3) Is makes clear areas of the design space that are under-utilized (e.g. *far proximity* - *intangible*).

For example, one particular design that defied easy classification is the Virtual Shelves implementation described by Li et al. [41]. With the Virtual Shelves interface, selectable objects, such as icons, are distributed in an *egocentric* sphere around the user. The user relies on spatial memory to make selections using a ray-casting metaphor, thus the objects are conceptually at a *far proximity*. This design combines some dimensional values not found in any of the main categories (Figure 4), such as an *egocentric-fixed* reference frame and *low visibility* with *discrete* space. The parallel coordinates visualization makes it easy to see that this design creates a unique pattern in the Ethereal Planes design space.

Filling Gaps, Tweaking and Combining

Beyond classification and comparison of existing designs, one purpose of a framework is to inspire and guide new creations. To show the generative potential of Ethereal Planes, we discuss several analytic processes that can be undertaken with our framework. Based on the work of Robinett [54], we explore three primary operations that can be used to transform our prior set classifications into ideas for new designs, by identifying *gaps* in the matrix, by '*tweaking'* (altering) existing designs or by *combining* two or more of them.

The first way to think about new designs is *filling gaps*; to look for valid combinations that have not been tried. By Robinett's method, our framework dimensions can be viewed as a seven-dimensional matrix, where each cell is a different combination of chosen values. Theoretically, this matrix has 288 unique design patterns. This number seems remarkable, considering that we were able to classify a large number of designs into only a handful of patterns. What then is the explanation for this difference? One

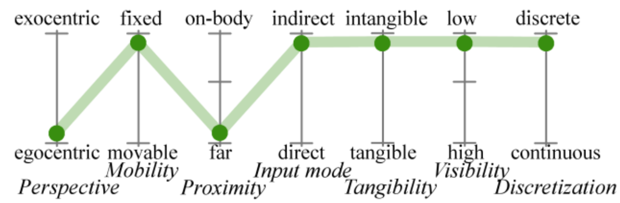

Figure 4. The Virtual Shelves design of Li et al. [41] holds a unique position in the design space from the major categories we identified in Figure 3.

primary reason is the number of interdependencies between the framework dimensions. Because the dimensions are not purely orthogonal, many of the possible combinations may be considered invalid. For instance, *direct* input with *far* information spaces seems impractical. However, the Ethereal Planes design space is still relatively unexplored and perceived dependencies may in fact be a result of attachment to prior paradigms. For instance, the most common reference frame types in the explored literature are *fixed-exocentric* and *movable-egocentric*, which correspond respectively to the most common types of real-world displays: desktop monitors and mobile devices. As designers gain more experience with mixed reality applications, some of the combinations that appear invalid may be explored with new and unconventional concepts. For example the *direct-far* combination mentioned above may be solved by the introduction a mechanism for controlling stretchable virtual limbs. On the other hand, *indirect-on-body* interaction might be found useful when looking at one's self in a mirror. In this manner, the Ethereal Planes framework is useful for plotting existing designs across the design dimensions, providing a methodical tool to help designers to identify new ground and inspire unique creations.

A second method for creating new designs is *tweaking*; rather than create a new combination from scratch, we can change one or two dimensions of existing patterns and imagine the resulting implications. In fact, one such example we identified in our literature review is the Imaginary Interfaces design of Gustafson et al. [29]. It is similar in nature to the *palette* category, however the user can 'draw' objects such as letters or mathematical functions with their fingertip on an intangible and invisible surface. This unusual design breaks the conventions of previous patterns by combining *low* visibility with a *continuous* workspace (Figure 5). Although only two dimensions are changed, the result introduces some significant design challenges, many of which are addressed in this novel work.

One other way to generate new ideas is to *combine* two or more existing patterns. An example of this type was also identified in our reviewed designs, in the AD-Binning implementation of Hasan et al. [33]. This interface extends the interaction plane of a mobile device screen into space around the device for making *discrete* item selections. This design has many dimensional values in common with

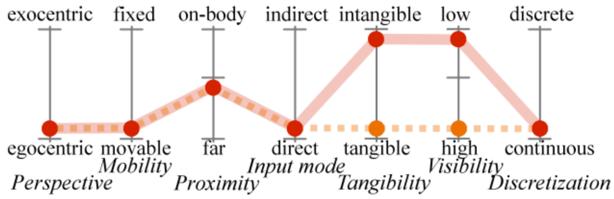

Figure 5. The Imaginary Interfaces design of Gustafson et al. [29] (solid path) varies from the *palette* category (dashed path) only in the *tangibility* and *visibility* dimensions.

palette category (*egocentric*, *movable*, *near proximity*, *direct input*), but also some in common with Virtual Shelves (*intangible, invisible, discrete* space). Combining these dimensions creates a new hybrid pattern, as seen in Figure 6. A similar fit to the framework was found in the Piles Across Space implementation of Wang et al. [64], which was designed for sorting photos into virtual piles around a desktop monitor. Designers of future interfaces can benefit from a design space that provides a conceptual workspace for trying new combinations.

One particular instance where combining existing designs can be useful is to support multiple interface 'modes' within a compound design. For example imagine a sketching application with read and write modes. Suppose a series of sketches are distributed in an egocentric sphere, floating around the user, which can be viewed using a mobile screen. When editing the sketches in write mode, the user uses the display as a *peephole*, since it provides a *tangible* surface to assist drawing in *continuous* space. To make drawing easier, the sketches are mapped to a single stationary (*exocentric*) plane, so the user doesn't need to change the device orientation. When viewing the sketches in read mode, however, the user can simply hold the device in one place and use her second hand as a pointer; the user knows the *discrete* location of each sketch in the *egocentric* sphere and whichever one she points to appears on the display. A single dimension can also act as a 'mode switch' within a single design. Imagine for instance an image browsing application. The user can have both a collaborative mode and a personal mode. To support sharing, the collaborative mode uses *exocentric* space, whereas the personal mode is placed in *egocentric* space.

Example: Designing the Personal Cockpit

To provide a final example of our framework's utility, we discuss a case where the Ethereal Planes framework was

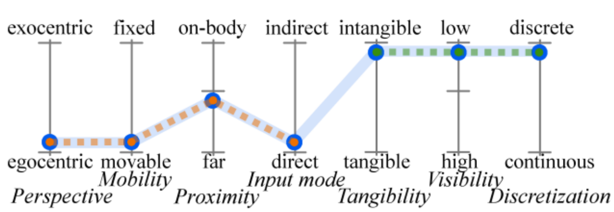

Figure 6. The AD-Binning design of Hasan et al. [33] (solid path) shares some dimensional values with the *palette* category (orange) and others with the Virtual Shelves design (green).

applied to an actual design. This case occurred during our work on the Personal Cockpit [3], a multi-display interface intended for use on HWDs (Figure 7). Here we briefly describe our implementation and walk through the seven design dimensions; along the way, we present our design choices, explain how they were influenced by the framework dimensions and provide some possible alternative choices for future implementations.

The Personal Cockpit is a spatial user interface for HWDs, intended for use with everyday mobile applications. Our design leverages free space around the user, allowing the user to partition content into multiple virtual windows that appear to float around the user's body. As an improvement over view-fixed windows available on current displays, our design allows faster task switching. We implemented the Personal Cockpit in a Cave environment, in which we emulated a HWD's limited field of view (FoV), and refined our design with several user studies. (For full details of the design, we refer readers to the referenced paper.)

Reference frame: The perspective of an information space is, to some extent, platform dependent. We have seen, for example, that designs leveraging the *peephole* metaphor use *exocentric* space to mitigate the limited display space of mobiles and projectors. An *exocentric* reference frame allows users to take advantage of proprioception for building spatial memory and helps to prevent them from getting lost in a large workspace.

With an ideal see-through HWD we would allow users to move virtual windows (2D information spaces) around freely in their environment. However, current devices require rendered content to fit within a limited FoV of about 40° or less (e.g. [63]). Since viewing content with this limitation is analogous to shining a projector's 'spotlight', we use *fixed* reference frames to maximize memorability. We allow the user to choose between *egocentric* and *exocentric* perspectives for different situations: *egocentric* windows are necessary for mobile use, whereas *exocentric* windows can be mapped to existing surfaces around the home or office to minimize occlusion and allow *tangible*, *direct* input. We nonetheless allow some *movable* exceptions to fixed windows: although windows will remain primarily fixed, users may want to periodically customize their arrangement, much as one

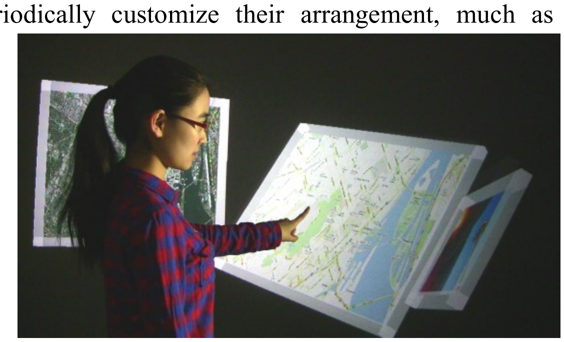

Figure 7. The Personal Cockpit [3] is a user interface design for using everyday applications on head-worn displays.

would rearrange icons on their mobile's home screen from time-to-time. For this purpose, we put handles on the windows, allowing them to be moved or resized using pinch gestures [45]. Also, users can move data objects from one window to another, or open a new window by dropping an application icon in mid-air.

Spatial Manipulation: We opted to explore *direct* input in our design to create an intuitive experience for users. Whereas some mechanism for *indirect* input makes sense with view-fixed displays (e.g. [27]), *direct* input is a good fit for the spatially-situated windows of the Personal Cockpit and may reinforce the user's sense or spatial awareness through proprioception. The use of *direct* input requires windows to be placed within arm's reach, in the *near* region. Unlike a peephole display, whose tangible surface aligns with the information space, the floating windows in our design are *intangible*. Because the lack of tangibility is known to present issues for direct input [13], we were required to mitigate these in our design. First, to provide depth feedback, we introduced a cursor that indicates whether a user's finger is in front of, intersecting, or behind a window. Second, the handles for moving or resizing windows are invisible by default, but change colour to indicate affordance for grasping when a hand is near (by turning green) and feedback when pinched (blue).

Spatial Composition: The information spaces in the Personal Cockpit are implemented as virtual windows, which are visible to the wearer of a HWD. Since these windows can be used to view rich application content, each window contains a *continuous* workspace. However, we also make the workspace *discrete* in a sense, since individual tasks are partitioned into different windows. Because the HWD's limited FoV allows only one window to be fully viewed at a time, our multi-window design has only *intermediate visibility,* however users will build up their spatial memory after repeated instances of switching between *fixed* windows. To reinforce visual spatial memory with proprioception, we place the body-fixed layout at a constant distance of 50 cm from the user's right shoulder. To make use of additional *egocentric* space around the user, the design could be expanded to include additional items placed fully out of normal viewing range. For example, a set of shortcut triggers could be placed at a region 90° to either side. Since the user will not often want to turn their head so far these items have a *low visibility*, supported by *discrete* space for easy recall**.**

FRAMEWORK EXTENSIONS

We acknowledge that there are limitations to our Ethereal Planes framework which may make it seem incomplete in certain contexts. However, we view Ethereal Planes as a core template that can be modified to suit a designer's needs, rather than a final product that fits all circumstances. Here we briefly discuss several potential extensions of our framework. These extensions include ideas that we initially attempted to introduce into our list of framework dimensions, but warrant deeper consideration at a higher level than is possible with the initial framework we introduce in this paper. Each of these topics requires several dimensions of its own that could constitute a separate layer of a more complete framework. In each case, these dimensions must be drawn from an additional body of literature and must be considered at a higher level than the basic interaction concepts of our initial framework.

Multi-modal interaction: Our input dimension takes into account only the paradigms of pointer selection and direct manipulation. This dimension could be expanded to include other input modes, particularly voice. The *visibility* dimension could similarly be expanded to consider non-visual output modes such as audio output. Such extensions would allow our framework to be extended to the design of interfaces for people with motor-skills or visual disabilities.

Co-located Collaboration: One of the applications of our framework is for collaborative scenarios. HWDs connected by network can be configured to allow multiple people to view the same virtual workspace from different perspectives [1]. Our framework could be extended by taking into consideration the large body of research on multi-surface environments. The modified framework should include aspects pertaining to the movement of content between surfaces and consideration of public vs private content [28].

Beyond 2D Surfaces: Our current framework focuses on 2D surfaces, although it could be extended to handle 3D objects. Such an extension should include additional dimensions to handle manipulation and viewing (grasping, rotation) of 3D objects. It should also include dimensions that take into account occlusion caused by the object's relative orientation or clutter from multiple objects.

CONCLUSION

We presented our Ethereal Planes framework for describing existing and new designs that use 2D information spaces in 3D mixed reality environments. From a bottom-up review of existing designs, we inferred our framework's seven dimensions – *perspective*, *movability*, *proximity*, *input mode*, *tangibility*, *visibility* and *discretization*. We provided a description of each of these dimension. We demonstrated how our framework can be used to describe, contrast and compare existing designs by grouping these into five representative categories that emerged from our analysis. We also show how our framework can assist in the development of new systems through operations such as filling gaps, tweaking or combining existing designs and discuss the framework's application during our design of the Personal Cockpit [3]. We provide examples of potential extensions to our framework to accommodate the specific needs of future designers.

ACKNOWLEDGMENTS

We acknowledge support from a NSERC Discovery Grant and a NSERC PGS scholarship for work on this project. We thank the anonymous reviewers for their helpful input.

REFERENCES

1. Agrawala, M., Beers, A.C., McDowall, I., Fröhlich, B., Bolas, M. and Hanrahan, P. The two-user responsive workbench: support for collaboration through individual views of shared space. *SIGGRAPH'97* (1997), 327-332.
2. Andujar, C. and Argelaguet, F. Friction surfaces: Scaled ray-casting manipulation for interacting with 2D GUIs. *Proc. EGVE '06*, Eurographics (2006), 101-108.
3. Ens, B., Finnigegan, R. and Irani, P. The Personal Cockpit: A spatial window layout for effective task switching on head-worn displays. *Proc. CHI '14*, ACM (2014), 3171-3180.
4. Avrahami, D., Wobbrock, J.O. and Izadi, S. Portico: Tangible interaction on and around a tablet. *Proc. UIST '11*, ACM (2011), 347-356.
5. Beaudonuin-Lafon, M. Instrumental interaction: An interaction model for designing post-WIMP user interfaces. *Proc. CHI '00*, ACM (2000), 446-453.
6. Billinghurst, M., Bowskill, J., Jessop, M. and Morphett, J. A wearable spatial conferencing space. Proc. ISWC '98, IEEE (1998), 76-83.
7. Boring, S., Baur, D., Butz, A. Gustafson, S. and Baudisch, P. Touch projector: Mobile interaction through video. *Proc. CHI '10*, ACM (2010), 2287-2296.
8. Bowman, D.A. and Hodges, L.F. Formalizing the design, evaluation and application of interaction techniques for immersive virtual environments. *Journal of Visual Languages and Computing 10* (1999), 37-53.
9. Brin, S. and Amirparviz, B. Laser alignment of binocular head mounted display. Patent No. 20130038510, Filed Aug. 9th, 2011, Iss. Feb. 14th, 2013.
10. Butler, A., Izadi, S., and Hodges, S. SideSight: multi-"touch" interaction around small devices. *Proc. UIST '08*, ACM (2008), 201-204.
11. Buxton, B. and Fitzmaurice, G. HMDs, Caves & Chameleon: A human-centric analysis of interaction in virtual space. *Computer Graphics 32*, 4 (1998), 69-74.
12. 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.
13. 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.
14. Chen, X.A., Marquardt, N., Tang, A., Boring, S. and Greenberg, S. Extending a mobile device's interaction space through body-centric interaction. *Proc. MobileHCI '12*, ACM (2012), 151-160.
15. Cockburn, A. and McKenzie, B. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. *CHI '02* (2002), 203-210.
16. Cockburn, A., Quinn, P., Gutwin, C., Ramos, G, and Looser, J. Air pointing: Design and evaluation of spatial target acquisition with and without visual feedback. *Int. Journal of Human-Computer Studies*. 69, 6, Academic Press (2011), 401-414.
17. Cuchard, J., Löchtefeld, 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.
18. de Haan, G., Griffith, E.J., Koutek, M. and Post, F.H. Hybrid interfaces in VEs: Intent and interaction. *Proc. EGVE '06*, Eurographics (2006), 109-118.
19. de Haan, G., Koutek, M. and Post, F.H. Towards intuitive exploration tools for data visualization in VR. *Proc. VRST '02*, ACM (2002), 105-112.
20. Döring, T., Sylvester, A. and Schmidt, A. A design space for ephemeral user interfaces. *Proc.TEI '13*, ACM (2013), 75-82.
21. Elias, L.J. and Saucier, D.M. *Neuropsychology : clinical and experimental foundations.* Pearson (2006).
22. 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.
23. Fisher, S.S., McGreevy, M., Humphries, J. and Robinett, W. Virtual environment display system. *Proc. I3D '86*, ACM (1986), 77-87.
24. Fitzmaurice, G.W. Situated information spaces and spatially aware computers. *Communications of the ACM 36, 7*, ACM (1993), 39-49.
25. Fitzmaurice, G.W., Ishii, H. and Buxton, W. Bricks: Laying the foundations for graspable user interfaces. *Proc. CHI '95*, ACM (1995), 442-449.
26. Fitzmaurice, G.W., Zhai, S. and Chignell, M.H. Virtual reality for palmtop computers. *Proc. TOIS '93*, ACM (1993), 197-218.
27. Google Glass. http://www.google.com/glass/start/
28. Greenberg, S., Boyle, M., and Laberge, J. PDAs and shared public displays: Making personal information public, and public information personal. *Personal Technologies* 3, 1 (1999), 54-64.
29. Gustafson, S., Bierwirth, D. and Baudisch, P. Imaginary interfaces: Spatial interaction with empty hands and without visual feedback. *Proc. UIST '10*, ACM (2010), 2-12.
30. Gustafson, S., Holz, C. and Baudisch, P. Imaginary phone: Learning imaginary interfaces by transferring spatial memory from a familiar device. *Proc. UIST '11*, ACM (2011), 283-292.
31. Harrison, C., Benko, H., and Wilson, A. D. OmniTouch: Wearable multitouch interaction everywhere. *Proc. UIST '11*, ACM (2011), 441-450.
32. Harrison, C., Tan, Desney and Morris, D. Skinput: Appropriating the body as an input surface. *Proc. CHI '10*, ACM (2010), 453-462.
33. Hasan, K., Ahlström, D. and Irani, P. AD-Binning: Leveraging around device space for storing, browsing

and retrieving mobile device content. *Proc. CHI '13*, ACM (2013), 899-908.

34. Holmes, N.P. and Spence, C. The body schema and multisensory representation(s) of peripersonal space. *Cognitive Processing 5*, 2 (2004), 94–105.
35. Hincapié-Ramos, J.D., Voida S. and Mark, G. A design space analysis of availability-sharing systems. *Proc. UIST '11*, ACM (2011), 85-95.
36. Hinckley, K., Pausch, R., Goble, J.C. and Kassell, N.F. A survey of design issues in spatial input. *Proc. UIST '94*, ACM (1994), 213-222.
37. Hoang, T.N. and Thomas, B.H. Augmented viewport: An action at a distance technique for outdoor AR using distant and zoom lens cameras. *ISWC '10* (2010), 1-4.
38. Jones, B., Sodhi, R., Forsyth, D., Bailey, B. and Maciocci, G. Around device interaction for multiscale navigation. *Proc. MobileHCI '12*, ACM (2012), 83-92.
39. Kane, S.K., Avrahami, D., Wobbrock, J.O., Harrison, B., Rea, A.D., Philipose, M. and LaMarca, A. Bonfire: A nomadic system for hybrid laptop-tabletop interaction. *Proc. UIST '09*, ACM (2009), 129-138.
40. Kaufmann, B. and Hitz, M. X-large virtual workspaces for projector phones through peephole interaction. *Proc. MM '12*, ACM (2012), 1279-1280.
41. Li, F.C.Y., Dearman, D. and Truong, K.N. Virtual shelves: Interactions with orientation aware devices. *Proc. UIST '09*, ACM (2009), 125-128.
42. Lindeman, R.W., Sibert, J.L. and Hahn, J.K. Towards usable VR: An empirical study of user interfaces for immersive virtual environments. *CHI '99* (1999), 64-71.
43. Lucero, A. Holopainen, J. and Jokela, T. Pass-them-around: Collaborative use of mobile phones for photo sharing. *Proc. CHI '11*, ACM (2011), 1787-1796.
44. Mine, M.R., Brooks, F.P. and Sequin, C. Moving objects in space: Exploiting proprioception in virtual-environment interaction. *SIGGRAPH '97* (1997), 19-26.
45. Piekarski, W. and Thomas, B.H. Tinmith-Metro: New outdoor techniques for creating city models with an augmented reality wearable computer. *Proc. ISWC '01*, IEEE (2001), 31-38.
46. Plaisant, C., Carr, D. and Shneiderman, B. Image-browser taxonomy and guidelines for designers. *Software 12,* 2 (Mar. 1995), 21-32.
47. PMDTechnologies. http://www.pmdtec.com
48. Poupyrev, I., Weghorst, S., Billinghurst, M. and Ichikawa, T. Egocentric object manipulation in virtual environments: Empirical evaluation of interaction techniques. *Comp. Graphics Forum 17*, 3 (1998), 41-52.
49. Pousman, Z. and Stasko, J. A taxonomy of ambient information systems: Four patterns of design. *Proc. AVI '06*, ACM (2006), 67-74.
50. Price, B.A., Baecker, R.M. and Small, I.S. A principled taxonomy of software visualization. *Journal of Visual Languages & Computing 4,* 3 (Sept. 1993), 211-266.
51. Rekimoto, J. and Saitoh, M. Augmented surfaces: A spatially continuous work space for hybrid computing environments. *Proc. CHI '99*, ACM (1999), 378-385.
52. Ren, G. and O'Neill, E. 3D selection with freehand gesture. *Comp. and Graphics 37*, 3 (2013), 101-120.
53. Ritchey. T. Fritz Zwicky, 'Morphologie' and policy analysis. Paper presented at *16th EURO Conference on Operational Analysis,* Brussels (1998).
54. Robinett, W. Synthetic experience: A taxonomy, survey of earlier thought, and speculations on the future. Technical report. University of North Carolina (1992).
55. Schmalstieg, D, Encarnação, L.M. and Szalavári, Z. Using transparent props for interaction with the virtual table. *Proc. I3D '99*, ACM (1999), 147-153.
56. SoftKinetic. http://www.softkinetic.com
57. Song, H., Grossman, T., Fitzmaurice, G., Guimbretiere, F., Khan, A., Attar, R. and Kurtenbach, G. PenLight: combining a mobile projector and a digital pen for dynamic visual overlay. *Proc. CHI '09* (2009), 143-152.
58. Song, H. Guimbretiere, F., Grossman, T. and Fitzmaurice, G. MouseLight: Bimanual interactions on digital paper using pen and a spatially-aware mobile projector. *Proc. CHI '10,* ACM (2010), 2451-2460.
59. Spindler, M., Büschel, W. and Dachselt, R. Use your head: Tangible windows for 3D information spaces in a tabletop environment. *Proc. ITS '12* (2012), 245-254.
60. Sutherland, I.E. A head-mounted three dimensional display. *Proc. AFIPS '68*, ACM (1968), 757-764.
61. Szalavári, Z. and Gervautz, M. The personal interaction panel: A two-handed interface for augmented reality. *Computer Graphics Forum 16*, Eurographics (1997), 335-346.
62. Teather, R.J. and Stuerzlinger, W. Pointing at 3D targets in a stereo head-tracked virtual environment. *Proc. 3DUI '11*, IEEE (2011), 87-94.
63. Vuzix Corporation. http://www.vuzix.com.
64. Wang, Q., Hsieh, T and Paepcke, A. Piles across space: Breaking the real-estate barrier on small-display devices. *Int. J. Hum.-Comput. Stud. 67*, 4, Elsevier (2009), 349-365.
65. Wickens, C.D. and Baker, P. Cognitive issues in virtual reality. In *Virtual Environments and Advanced Interface Design*, Furness, T.A. and Barfield, W. (Eds.). Oxford (1995), 514-542.
66. Wilson, A.D. PlayAnywhere: A compact interactive tabletop projection-vision system. *Proc. UIST '05*, ACM (2005), 83-92.
67. Wilson, A.D. and Benko, H. Combining multiple depth cameras and projectors for interactions on, above and between surfaces. *Proc. UIST '10* (2010), 273-282.
68. Yee, K. Peephole Displays: Pen interaction on spatially aware handheld computers. *Proc. CHI '03* (2003), 1-8.