A Mixed Reality Space for Tangible User Interaction

Martin Fischbach*, Christian Treffs*, David Cyborra†, Alexander Strehler†
Thomas Wedler†, Gerd Bruder†, Andreas Pusch†, Marc E. Latoschik*, Frank Steinicke†

*Human-Computer Interaction
Universität Würzburg
Am Hubland
97074 Würzburg
Tel.: +49 931 31 86314
martin.fischbach@uni-wuerzburg.de

†Immersive Media Group
Universität Würzburg
Oswald-Külpe-Weg 82
97074 Würzburg
Tel.: +49 931 31 85868
gerd.bruder@uni-wuerzburg.de

Abstract: Recent developments in the field of semi-immersive display technologies provide new possibilities for engaging users in interactive three-dimensional virtual environments (VEs). For instance, combining low-cost tracking systems (such as the Microsoft Kinect) and multi-touch interfaces enables inexpensive and easily maintainable interactive setups. The goal of this work is to bring together virtual as well as real objects on a stereoscopic multi-touch enabled tabletop surface. Therefore, we present a prototypical implementation of such a mixed reality (MR) space for tangible interaction by extending the smARTbox [FLBS12]. The smARTbox is a responsive touch-enabled stereoscopic out-of-the-box system that is able to track users and objects above as well as on the surface. We describe the prototypical hard- and software setup which extends this setup to a MR space, and highlight design challenges for the several application examples.

Keywords: Multi-touch, tangible interaction, stereoscopic display

1 Introduction

Research on multi-touch display interaction has mainly focused on monoscopically rendered 2D or 3D visualizations, whereas the challenges that occur when the virtual content is displayed stereoscopically in the 3D space, but interaction remains constrained to a 2D surface, has rarely been considered so far. Over the last couple of years, stereoscopic feedback received much attention due to the rise of 3D motion pictures and 3D consumer technology (e. g., TVs, video game consoles, notebooks). First commercial hardware systems supporting stereoscopic display and (multi-)touch interaction have recently been launched (iliGHT 3D Touch[1], Nintendo 3DS[2], etc.), and interdisciplinary research projects provide first insights into how interaction with stereoscopic content displayed on a two-dimensional touch surface (iMUTS[3], InSTInCT[4], etc.) has to be designed.

[1]http://ilight-3i.com/en
[2]http://www.nintendo.com/3ds
[3]http://imuts.uni-muenster.de
[4]http://anr-instinct.cap-sciences.net/?q=node/34

More and more hardware solutions for the entertainment market allow the sensing of human gestures and postures performed both on touch-sensitive surfaces as well as in the 3D space (e. g. Leap Motion[5], Microsoft Kinect[6]). Although these technologies provide an enormous potential for a variety of novel applications and interaction concepts, merely first steps have been done in the area of multi-touch enabled stereoscopic surfaces. Such setups are still rarely used, because certain aspects hinder a wide acceptance [SBK+12]. For instance, the fact that stereoscopic content can be displayed behind or in front of the projection screen, whereas a touch can only be detected on (or very close to) the surface, induces a visuo-haptic conflict [SSK+09]. Another problem comes from accommodation-vergence conflicts that are present in most stereoscopic display environments and become even more important when users have to switch their current visual focus between their hands and the stereoscopic objects [VSBH11].

For a long time, human beings have shaped their physical environment to materialize concepts of form and function. Despite the diversity of existing interaction devices, the human hand remains by far the most natural and fundamental interface for interacting with real-world objects as well as with virtual objects. Seamlessly merging these real and the virtual worlds created within a computer is a challenging topic in current research. The result of superimposing the real world by computer-generated images is called augmented reality (AR) whereas the enhancement of virtual worlds using real-world data is called augmented virtuality (AV). The term mixed reality (MR) [MK94] encompasses both AR as well as AV. As pointed out by [MDG+95], the main issues of MR environments are consistency of geometry, time, and illumination. Therefore, one of the most important tasks is to match a superimposed object's location with the that of the corresponding real world object. Likewise, reflections, light sources and shadows must match in both worlds to obtain consistent global illumination; hence movements in the two worlds must be synchronized.

MR systems have been the subject of growing interest by designers due to the dual need of users to both benefit from computing power and interact with the real world. A first attempt to satisfy this requirement consisted in augmenting the real world with digital information: The central rationale for AR. Another approach is to make interaction with computer-generated content more natural by means of a transfer of objects and actions. Examples include input modalities [NC95] based on real objects, such as Fitzmaurice et al.'s bricks [FIB95], or Ishii & Ulmer's phicons [IU97]. Ishii has described this interaction paradigm as *Tangible User Interface (TUI)*. As a tangible is placed on a table that also acts as a display, various interactions can be performed. For instance, animated symbols appear, such as waveforms, circles, circular grids, or sweeping lines. While some symbols merely show what the particular tangible is doing, others can be used with the fingertip to control the respective module.

In this paper, we present how the smARTbox [FLBS12] can be used for tangible inter-

[5]http://www.leapmotion.com
[6]http://www.microsoft.com/en-us/kinectforwindows

action as a mixed reality (MR) design space. The smARTbox is a responsive touch-enabled stereoscopic tabletop system. For this work we augmented this setup to support tangible interaction with physical objects based on fiducial tracking. These physical objects serve as tangible proxies for corresponding virtual objects, which have been transferred to the virtual world by a simple 3D reconstruction process.

The overall goal is thus to bring together virtual and real objects using a stereoscopic multi-touch-enabled tabletop device to create a novel tangible mixed reality design space. The contributions of this paper can be summarized as:

- the description and implementation of a first stereoscopic multi-touch-enabled tabletop setup supporting fiducial tracking,
- the proposition of a tangible mixed reality design space, and
- the development of interaction concepts for 3D design and modeling tasks.

2 Related Work

Many approaches for extending multi-touch interaction techniques to 3D applications with monoscopic rendering have been proposed. For instance, Hancock et al. [HCC07] have presented the concept of shallow-depth 3D, i. e., 3D interaction with limited depth. They have developed interaction techniques for 3D object selection and manipulation. Extending the interaction space beyond the touch surface has been investigated by Hilliges et al. [HIW+09]. They have tested two depth sensing approaches to enrich multi-touch interaction on a tabletop setup with monoscopic projection. Grossman and Wigdor [GW07] provided an extensive review of the existing work on interactive surfaces and developed a taxonomy for classification of this research. This framework takes into account the perceived and the actual display space, the input space and the physical properties of an interactive surface.

Schöning et al. [SSK+09] have discussed the challenges of direct touch interaction with stereoscopically rendered scenes, and first prototypes have been introduced since then as described in Section 1. Interaction with objects with negative parallax on a multi-touch tabletop setup has further been addressed by Benko et al. [BF07] and Strothoff et al. [SVH11]. As mentioned above, problems arise in such setups because stereoscopic content can be displayed behind or in front of the projection screen, whereas a touch can only be detected on (or very close to) the surface [SSK+09]. The mapping between an on-surface touch point and the corresponding point on a virtual object is straightforward in the monoscopic case. But with stereoscopic projection, this mapping introduces problems [TS11]. Particularly, since there are different projections for each eye, the question is where users would touch the surface when they tried to touch a stereoscopic object. One may expect that the touch would be anywhere around the two projections of a stereoscopically displayed object. What Valkov et al. found is that users actually touch an intermediate point which is located between both projections with a significant offset to the users dominant eye [VSBH11]. However, as shown by previous work, users are insensitive to discrepancies between visual penetration and touch feedback when they try to touch stereoscopic objects displayed close to the surface [VSB+10]. In particular, it has been found that sensitivity to these visuo-tactile conflicts is lower, if objects are displayed with negative parallax, e. g., on top or above a tabletop surface.

Research efforts as well as commercialization of tangible and multi-touch user interaction has recently seen an exponential growth. For instance, Microsoft distributes the Windows-based platform Microsoft Surface[7]. Beside multi-touch tracking of fingers, the platform supports the recognition of physical objects by their footprints such as cell phone or digital cameras. The Tangible Media Lab led by Hiroshi Ishii has designed several interfaces that employ physical objects, surfaces, and spaces as tangible embodiments of digital information and processes[8]. AR-Jig, for example, is a handheld TUI for 3D digital modeling in an AR space [AI07]. AR-Jig has a pin array that displays a 2D physical curve coincident with a contour of a digitally displayed 3D form. It supports physical interaction with a portion of a 3D digital representation, allowing 3D forms to be directly touched and modified. Recompose has been proposed as a system for manipulating an actuated surface [LLDB11]. This actuated surface enables a number of interaction techniques exploiting the shared space of direct and gestural input. Another example of a prototype system for interactive construction and modification of 3D physical models is based on building blocks [MWC+12]. Similar to our system, theirs uses a depth sensing camera for acquiring and tracking the physical models.

The combination of multi-touch technology and TUIs has been employed in various applications. Two of them are highly relevant to the application scenario described in this paper. Urp (Urban Planning Workbench) [UI99] is one example of a TUI, which uses miniature physical models of architectural buildings to configure and control an underlying urban simulation of shadow, light reflections, or wind flow. Another notable interactive installation is instant city [9], which combines gaming, music, and architecture. The user can build 3D structures and set up a city with rectangular building blocks, which simultaneously results in the interactive assembly of musical fragments of different composers. However, as already mentioned above, none of the described setups presented so far is capable of supporting multi-touch and tangible user input combined with stereoscopic display.

3 Tangible Mixed Reality Tabletop Setup

This section describes the system setup used for the implementation of the developed tangible user interface concepts. It is an extension of the smARTbox setup [FLBS12], combining five key features in the frame of a portable box: (1) stereoscopic visualization, (2) multi-touch detection, (3) fiducial tracking, (4) full-body user tracking and (5) simple 3D reconstruction.

In the following a detailed walk trough the realization of this key features is presented.

[7]http://www.microsoft.com/surface
[8]http://tangible.media.mit.edu/
[9]http://www.instantcity.ch/

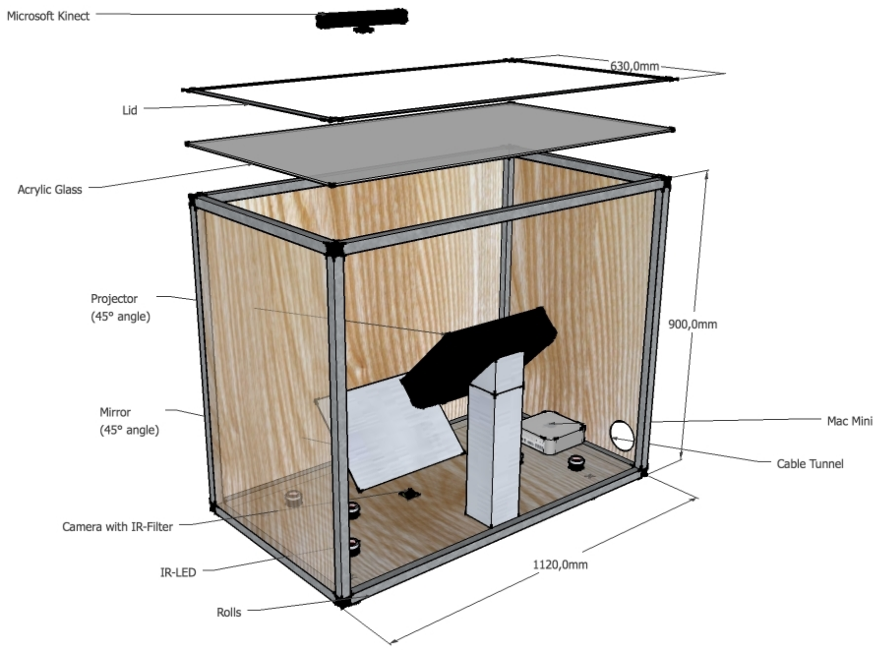

Figure 1: Stereoscopic multi-touch enabled tabletop setup (see also [FLBS12]).

A schematic of the utilized hardware setup is given in Figure 1. It was constructed at a total cost of less than USD 4,000. The utilized software setup is based on the software platform Simulator X [LT11], which provides a framework for input and output handling as well as for the application itself.

3.1 Stereoscopic Visualization

The top side of the table consists of a 62 cm x 112 cm back projection screen with a gain of 1.6, allowing a bright display for a large range of viewing angles. For stereoscopic visualisation an Optoma GT720 DLP projector with a resolution of 1280x800 pixels at a refresh rate of 120Hz is used. The image is projected from the base of the table to the back projection screen via a mirror that is mounted at an angle of 45°. Using active shutter glasses, the stereo images are displayed frame-sequentially to the eyes of a user with a rate of 60Hz per eye. The VE is rendered on an Intel Core i7 computer with 3.40GHz processors, 8GB of main memory, and an nVidia Quadro 4000 graphics card.

3.2 Multi-touch Detection

The back projection screen is enhanced to a horizontal touch-sensitive input surface via the Rear-DI principle [MTSH+10]: Six clusters of high-power infrared (IR) LEDs illuminate the entire screen surface from below. Thereby the acrylic glass serves as diffusing layer. In the case an object comes in contact with the surface, e.g., a finger or palm, it reflects the IR light back. The reflected light is recorded by a Point Grey Dragonfly®2 camera, equipped with an IR band-pass filter and a wide-angle lens. It captures 8-bit monochrome images with a resolution of 1024x768 pixels at a rate of 30 frames per second. A modified version of NUI Group's CCV (Community Core Vision) software [10] uses this video stream to detect the touch input. Detected touch points are provided to applications via TUIO [11], a protocol commonly used for input forwarding from tangible multitouch surfaces.

3.3 Fiducial Tracking

CCV's fiducial tracking functionality is utilized in order to identify real-world objects. Fiducials are markers, which work similar to the QR-Code system, but are easier structured, so that they can be differentiated even at larger distances. CCV is able to distinguish between 216 fiducials. It assigns one unique ID to each of them and tracks their orientation (around the vertical axis) and position (in the surface plane) in a 2D coordinate system, registered with the touch surface. This information is provided to applications via the TUIO protocol. Fiducials are used to track several real-world objects, in particular block models. They are applied as proxy objects or interaction devices (cf. Section 4).

3.4 Kinect Body Tracking

A Microsoft Kinect multi-sensor system is utilized in combination with the Flexible Action and Articulated Skeleton Toolkit (FAAST [SLR+11]) for real-time full-body tracking of users interacting with the setup. FAAST provides detected spatial data via the VRPN protocol [THS+01], and in particular allows to track the head position and orientation of a user at an acceptable level of precision. This tracking data is used to realize head-coupled rendering, and to generate stereoscopic views (see Figure 3).

3.5 3D Reconstruction Process

The 3D reconstruction process is based on the application of the simple open source RGB-Demo[12] software. It allows the 3D capturing of real-world objects and thus facilitates their representation in a VE. The RGBDemo reconstruction process starts with the capturing of the desired real-world object from different viewpoints. The software registeres these different datasets and reconstructs a corresponding 3D point mesh. The mesh is exported as a .ply-file and can manually be refined.

4 Tangible User Interaction

In this section, we describe the interaction concepts based on the described tangible MR tabletop surface.

[10] http://ccv.nuigroup.com/

[11] TUIO is an open framework that defines a common protocol and API for tangible multitouch surfaces (see http://www.tuio.org/)

[12] http://labs.manctl.com/rgbdemo/

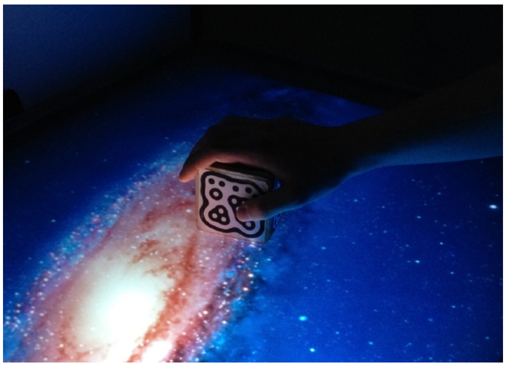

Figure 2: 3D interaction cube with fiducials attached to all six faces.

4.1 Tangible Proxy Objects

One essential idea of this work is to turn tangible real-world objects into manipulable virtual objects. There are several possibilities for reconstructing objects in 3D, from expensive setups with multiple infra-red sensing cameras to low-budget solutions like the Microsoft Kinect. In our setup we use a simple reconstruction based on the Kinect and RGBDemo application as described above. After a physical object has been digitized (produced by the reconstruction process or by traditional 3D modeling) and been brought to the virtual world, we have access to a 3D virtual model registered with the corresponding tangible proxy object.

Most objects users want to interact with on a tabletop surface provide some upright position. We simply place a specified fiducial marker beneath the physical proxy object to associate the two. When the marker is in contact with the table and has been identified by CCV, one instance of its associated virtual 3D object is imported into the VE and displayed on the surface next to the tangible object (see Figure 3 (left)). Now, the user can interact with the tangible proxy object, while its transformations (translations and rotations around the yaw axis) are applied to the virtual object in a one-to-one mapping. When the user uplifts the tangible object away from the surface, the virtual object remains in its last pose until it is selected and manipulated again. The user can generate several instances of each object, simply by re-placing the tangible object back to the surface.

Due to the Fiducials, the interaction with virtual objects is constrained to 2D translations and rotations around the yaw axis, which is enough for most tangible objects placed on a touch surface. In addition, simple pinching gestures at the tangible proxy prop could be used to scale the virtual objects accordingly.

4.2 3D Interaction Cube

In order to support 3D interaction with virtual objects in a more generic way, we constructed another simple, but effective and intuitive input device. The main control device of our setup for 3D view and object manipulation is an 8x8x8cm wooden cube (see Figure 2). This cube can be used as tangible input device for many operations. To track the cube from all sides when it is placed on the surface, we attached an explicit Fiducial to each face. Now, rotations

Figure 3: Tangible MR design space and a user interacting (left) with a virtual teapot using the physical proxy teapot, and (right) with both hands in a 3D city design application. All proxy objects are tracked at three degrees of freedom based on the Fiducials attached to their bottom sides.

of the cube to a different face can easily be tracked and corresponding information such as the Fiducial's ID and its alignment can be sent via the TUIO protocol. Hence, each of the six faces of the cube provide three degrees of freedom, i. e., translation on the surface and yaw rotations. In addition, overturning the cube to each direction can be detected. This provides great potential for many application scenarios.

We implemented a simple view manipulation technique. In modern 3D modeling applications, e. g., Blender or 3D Studio Max, usually orthogonal projections to the 3D scene are supported. The faces top, bottom, left, right, front and back can be perfectly mapped to the faces of our cube, which allows an easy way of changing the that view in a VE. Flipping the cube at 90 degree around any of its six axis (from which three are collinear) switches the view to the corresponding view using a smooth animation between the current and the new view (see Figure 3(right)).

5 Applications

In our opinion the proposed MR space for tangible interaction has enormous potential for a variety of different applications domains ranging from city or terrain modeling, architectural design, gaming, edutainment to simulation. In particular the horizontal surface allows to lay objects on the surface and to perform yaw rotations as described above in a very natural and constrained way. In architectural design one could use tangible proxies for different kind of furnitures, and arrange them in virtual rooms, which are stereoscopically displayed on the surface. In flow simulation, tangible objects may serve as blockers in order to manipulate the flow. One could play different kinds of games on our setup such as shuffle board.

So far, we have implemented one prototypical application from the domain of urban planing. Urban planning tasks are of great importance to civil works since both the environment and the inhabitants of a city are affected. The cityscape as well as the quality of life of the residents essentially depend on the appearance of buildings, road networks, planting, green spaces, and recreation areas such as parks and playgrounds. To facilitate a realistic impression of how a building area would visually integrate into the environment and to enable communication regarding development proposals, it is important to present intermediate results of the planning process properly as well as comprehensibly [SRHM06]. On the basis of a development plan, architects generate virtual 3D models and return three-dimensional visualizations of these areas to the urban planner. This procedure has the following two major shortcomings. First, the returned visualizations are static insofar as the 3D models cannot be explored interactively by the urban planners. Second, modifications to the 3D models, which, for instance, have been proposed after reviewing the 3D visualization, cannot be performed by urban planners. Instead, the architectural office has to be asked again to incorporate these modifications into the 3D model. During a planning task, this usually takes several iterations resulting in inefficiency as well as unnecessary expense [SRHM06].

Our system offers facilities to overcome these issues. After creating instances of a virtual 3D city such as buildings or trees, and registering these virtual objects with corresponding tangible proxies, the user can interact with the virtual objects by means of interacting with the tangible proxy objects. Again, placing a tangible object (with Fiducial) on the table, the registered virtual object is rendered to the VE with a small offset. The virtual 3D model can now be transformed by rotating and/or translating the tangible proxy object as explained above. Figure 3(right) shows how a user interacts with the 3D interaction cube and a tangible proxy object registered to a virtual building. The 3D interaction cube in the non-dominant hand is used for viewpoint manipulations, whereas the tangible proxy building is used to manipulate the position and orientation of buildings. Another tangible proxy can be used to manipulate virtual trees in this setup.

6 Conclusion

In this paper, we presented the concept of a tangible MR interaction space. Therefore, we described a prototypical technical setup and highlighted central interaction challenges. To our knowledge, this is the first stereoscopic multi-touch-enabled tabletop setup supporting tracking of Fiducials and interaction with tangible proxy objects.

This setup has a great potential for future interaction concepts in semi-immersive setups. For instance, one can study how far the co-located, but slightly offset interaction with a tangible proxy gives the sensation of interacting with the corresponding virtual object. Another question may be to which extent the tangible proxies can differ from the registered virtual object without confusing the user. One essential future task is to improve the setup, in particular in terms of its 3D reconstruction and tracking algorithms. 3D model acquisition, reconstruction, and tracking of physical objects have a long history in computer vision, with applications in areas such as VR, AR, object modeling, simulation, CAD/CAM, and cultural heritage, whereas our current reconstruction process is only a proof of concept. Finally, we will explore more application domains despite from the described city modeling application in order to learn more about the potentials and limitations of our proposed MR space for tangible user interaction.

References

[AI07] M. Anabuki and H. Ishii. AR-Jig: A Handheld Tangible User Interface for Modification of 3D Digital Form via 2D Physical Curve. In *Proceedings of the 6th IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR)*, 2007.

[BF07] H. Benko and S. Feiner. Balloon Selection: A Multi-Finger Technique for Accurate Low-Fatigue 3D Selection. *IEEE Symposium on 3D User Interfaces*, 7:79–86, 2007.

[FIB95] G. Fitzmaurice, H. Ishii, and W. Buxton. Bricks: Laying the Foundations for Graspable User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1995.

[FLBS12] M. Fischbach, M. Latoschik, G. Bruder, and F. Steinicke. smARTbox: Out-of-the-Box Technologies for Interactive Art and Exhibition. In *Virtual Reality International Conference (VRIC)*. ACM, 2012.

[GW07] T. Grossman and D. Wigdor. Going Deeper: a Taxonomy of 3D on the Tabletop. In *Proceedings of the Second Annual International Workshop on Horizontal Interactive Human-Computer Systems*, pages 137–144. IEEE Computer Society, 2007.

[HCC07] M. Hancock, S. Carpendale, and A. Cockburn. Shallow-depth 3d interaction: design and evaluation of one-, two- and three-touch techniques. In *Proceedings of Human Factors in Computing Systems (CHI)*, pages 1147–1156. ACM, 2007.

[HIW+09] O. Hilliges, S. Izadi, A. D. Wilson, S. Hodges, A. Garcia-Mendoza, and A. Butz. Interactions in the Air: Adding Further Depth to Interactive Tabletops. In *Symposium on User Interface Software and Technology (UIST)*, pages 139–148. ACM, 2009.

[IU97] H. Ishii and B. Ullmer. Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In *Proceedings on Human Factors in Computer Systems (CHI)*, pages 234–241. ACM, 1997.

[LLDB11] D. Leithinger, D. Lakatos, A. DeVincenzi, and M. Blackshaw. Recompose: Direct and Gestural Interaction with an Actuated Surface. In *SIGGRAPH Emerging Technologies*, volume 13. ACM, 2011.

[LT11] M. E. Latoschik and H. Tramberend. Simulator X: A Scalable and Concurrent Software Platform for Intelligent Realtime Interactive Systems. In *Proceedings of IEEE Virtual Reality (VR)*, pages 171–174, 2011.

[MDG+95] P. Milgram, D. Drascic, J. Grodski, A. Restogi, S. Zhai, and C. Zhou. Merging Real and Virtual Worlds. In *Proceedings of IMAGINA*, 1995.

[MK94] P. Milgram and F. Kishino. A Taxonomy of Mixed Reality Visual Displays. In *IEICE Transactions on Information and Systems, Special issue on Networked Reality*, 1994.

[MTSH+10] C. Müller-Tomfelde, J. Schöning, J. Hook, T. Bartindale, D. Schmidt, P. Oliverk, F. Echtler, N. Motamedi, P. Brandl, and U. Zadow. Building Interactive Multi-touch Surfaces. In *Tabletops - Horizontal Interactive Displays*, Human-Computer Interaction Series, pages 27–49. Springer London, 2010.

[MWC+12] A. Miller, B. White, E. Charbonneau, Z. Kanzler, and J. LaViola. Interactive 3D Model Acquisition and Tracking of Building Block Structures. In *Transactions on Visualization and Computer Graphics*, volume 18, pages 651–659. IEEE, 2012.

[NC95] L. Nigay and J. Coutaz. A Generic Platform for Addressing the Multimodal Challenge. In *Proceedings of Human Factors in Computing Systems (CHI)*, pages 98–105. ACM, 1995.

[SBK+12] F. Steinicke, H. Benko, A. Krüger, D. Keefe, J.-B. de la Riviére, K. Anderson, J. Häkkilä, L. Arhippainen, and M. Pakanen. The 3rd dimension of CHI (3DCHI): touching and designing 3D user interfaces. In *Proceedings of the 2012 ACM annual conference extended abstracts on Human Factors in Computing Systems Extended Abstracts*, pages 2695–2698. ACM, 2012.

[SLR+11] E. Suma, B. Lange, A. Rizzo, D. M. Krum, and M. Bolas. FAAST: The Flexible Action and Articulated Skeleton Toolkit. In *Proceedings of IEEE Virtual Reality (VR)*, pages 245–246, 2011.

[SRHM06] F. Steinicke, T. Ropinski, K. Hinrichs, and J. Mensmann. Urban City Planning in Semi-immersive Virtual Reality. In *Proceedings of the International Conference on Computer Graphics Theory and Applications (GRAPP)*, pages 192–199. INSTICC Press, 2006.

[SSK+09] J. Schöning, F. Steinicke, A. Krüger, K. Hinrichs, and D. Valkov. Bimanual Interaction with Interscopic Multi-Touch Surfaces. In *Proceedings of the 12th IFIP TC 13 International Conference on Human-Computer Interaction: Part II*, INTERACT, pages 40–53, Berlin, Heidelberg, 2009. Springer.

[SVH11] S. Strothoff, D. Valkov, and K. Hinrichs. Triangle Cursor: Interactions With Objects Above the Tabletop. In *Proceedings of the International Conference on Interactive Tabletops and Surfaces*, pages 111–119. ACM, 2011.

[THS+01] R. M. Taylor, T. C. Hudson, A. Seeger, H. Weber, J. Juliano, and A. T. Helser. VRPN: A Device-Independent, Network-Transparent VR Peripheral System. In *Proceedings of the Symposium on Virtual Reality Software and Technology*, pages 55–61. ACM, 2001.

[TS11] R. J. Teather and W. Stürzlinger. Pointing at 3D targets in a stereo head-tracked virtual environment. In *Proceedings of 3DUI*, pages 87–94. IEEE, 2011.

[UI99] J. Underkoffler and H. Ishii. Urp: A Luminous-tangible Workbench for Urban Planning and Design. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 386–393. ACM, 1999.

[VSB+10] D. Valkov, F. Steinicke, G. Bruder, K. Hinrichs, J. Schöning, F. Daiber, and A. Krüger. Touching Floating Objects in Projection-based Virtual Reality Environments. In *Proceedings of Joint Virtual Reality Conference (JVRC)*, pages 17–24, 2010.

[VSBH11] D. Valkov, F. Steinicke, G. Bruder, and K. Hinrichs. 2D Touching of 3D Stereoscopic Objects. In *Proceedings of Human Factors in Computing Systems (CHI)*, pages 1353–1362. ACM, 2011.