VRRobot: Robot Actuated Props in an Infinite Virtual Environment

Emanuel Vonach* Clemens Gatterer† Hannes Kaufmann‡

Interactive Media Systems Group, TU Wien, Vienna, Austria

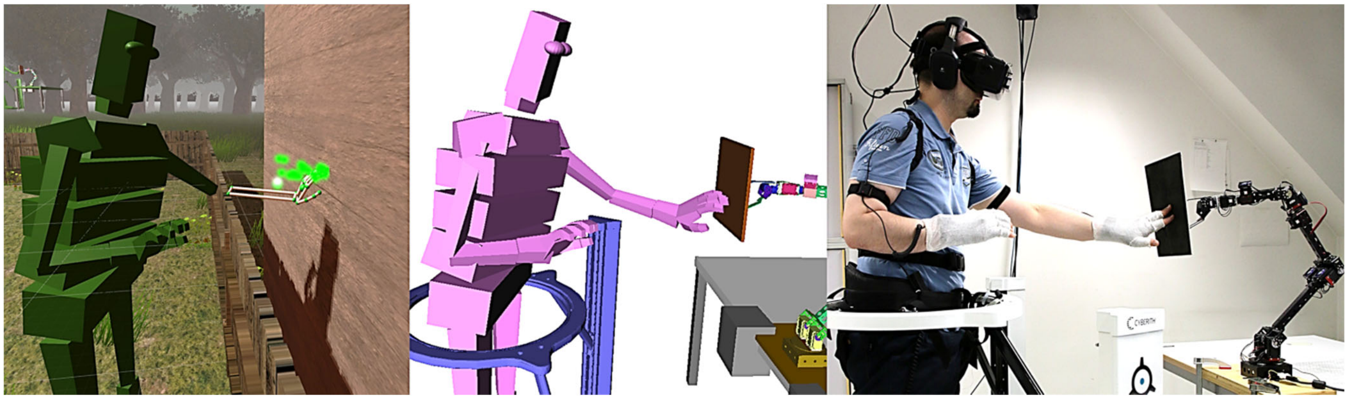

Figure 1: Left: User touches an object in the virtual world. Middle: The correct relative position to the user in the robot environment is known. Right: A physical prop is positioned accordingly in the real world to provide haptic feedback.

ABSTRACT

We present the design and development of a fully immersive virtual reality (VR) system that can provide prop-based haptic feedback in an infinite virtual environment. It is conceived as a research tool for studying topics related to haptics in VR and based on off-the-shelf components. A robotic arm moves physical props, dynamically matching pose and location of an object in the virtual world. When the user reaches for the virtual object, his or her hands also encounter it in the real physical space. The interaction is not limited to specific body parts and does not rely on an external structure like an exoskeleton. In combination with a locomotion platform for close-to-natural walking, this allows unrestricted haptic interaction in a natural way in virtual environments of unlimited size.

We describe the concept, the hardware and software architecture in detail. We establish safety design guidelines for human-robot interaction in VR. Our technical evaluation shows good response times and accuracy. We report on a user study conducted with 34 participants indicating promising results, and discuss the potential of our system.

Keywords: Prop-based virtual reality, encounter-type haptics, passive haptic feedback, fully immersive virtual reality.

Index Terms: H.5.2 [Information interfaces and presentation]: User Interfaces—Input devices and strategies; Haptic I/O; H.5.1 [Information interfaces and presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities;

1 INTRODUCTION

In recent years considerable progress towards fully immersive virtual reality environments has been made, including the introduction of new head mounted displays (e.g. Oculus Rift, HTC Vive), novel devices for close-to-natural locomotion (e.g. Cyberith Virtualizer, Virtuix Omni) and a range of affordable options for tracking a user's motions (e.g. Perception Neuron, Microsoft Kinect). However, even with a state-of-the-art VR setup and highly realistic rendering, a user is constantly reminded that the presented world is only virtual as soon as he or she tries to interact with an object within the simulation.

Despite haptics being a very active field of research, the approaches for general purpose haptic interfaces proposed in literature have severe limitations for fully immersive human scale VR (see section 3). Haptic feedback is a key aspect in a number of application scenarios for VR, especially if the objective is a transfer of experience from the virtual to the physical world [1], for example in case of the training of emergency service personnel or the treatment of psychological disorders (e.g. claustrophobia). A number of results reported in previous work (e.g. [2]) suggest that haptic feedback can improve task performance significantly as well as the perceived presence in a virtual environment. In order to increase the effectiveness and applicability of VR for a variety of application areas, we need a way to provide general purpose haptic feedback that is as close as possible to the way we interact with real world objects.

There are two concepts in previous works that show the potential to accomplish this goal. The first one is to position a haptic device at a desired place with a robotic arm and wait for the user to encounter it. This kind of feedback is referred to as an encounter-type haptic device and has the potential for free hand interaction and real touch sensations [3]. Furthermore different forces can be simulated, for example the resistance when pushing an object [4] or complex devices like a control panel [5]. However, previous systems based on this approach have severe restrictions, for example a limited working volume or a constrained way of interaction (i.e. with a single finger or a handle). This limits the applicability for a number of scenarios where the feedback should be close to the way we interact with real world objects.

Another approach is a method called passive haptics [6]. Real physical objects are employed to provide haptic feedback through

*e-mail: emanuel.vonach@tuwien.ac.at
†e-mail: clemens@gatterer.at
‡e-mail: hannes.kaufmann@tuwien.ac.at

their inherent properties, like shape, texture, etc. Although it has been shown before, that this approach can significantly increase the realism of a virtual experience [7], the physical props are not controlled by a computer and therefore would have to be available in every form and on every position required by the simulation. Moreover as soon as the virtual world would move, the physical and virtual objects would not be co-located anymore.

In this work we present a fully immersive VR system, incorporating aspects of encounter-type devices and prop-based haptic feedback without the described limitations (Figure 1). As illustrated in Figure 2 a head mounted display (HMD), a surround sound headset and a locomotion platform for close-to-natural walking allow the user to completely immerse in a virtual environment. The finger movements and the full body pose of the user are registered as well. A robotic arm is employed for actuating physical props for haptic feedback, correctly matching pose and location of an object in respect to the user's current position in the virtual world.

2 Contribution

One of the major contributions of this paper is the development of the system utilizing robot actuation of arbitrary physical props for haptic feedback in an infinite immersive virtual environment. In contrast to previous work (see section 3) the haptic feedback is not limited to one finger or hand, and the user is not impeded by an external structure like an exoskeleton. In fact a virtual object can be touched with any body part if the user wants to. The finger tips are completely free which allows to feel the inherent properties of the touched object. Moreover, the interactive environment is not limited to a small volume, the presented concept and the combination with a locomotion platform allow haptic interaction and free movement in a close-to-natural fashion in a virtual world of quasi unlimited size.

We address the newly arising challenges related to this approach, i.e. constantly updating the relative position of a physical object according to the current position of the user in VR, or taking the full body pose of the user into account for robot path planning. Furthermore, we demonstrate a novel way of robot-haptic interaction, as the robot arm initiates contact with a walking user, instead of purely reacting to his or her touch. We also present results from a user study with 34 participants affirming the capabilities of our system and the possibilities of our approach.

Finally, we contribute practical safety design guidelines for researchers in human-robot interaction in VR. The system is built up for the most part from relatively inexpensive off-the-shelf components, which facilitates replicability. We intended this system as a flexible research tool for studying topics related to haptic VR, and hope to inspire others to follow and expand on this approach.

The paper is structured as follows: Related work is presented and discussed in section 3. In section 4 we report on the concept and the development of our system. We specify the hardware and software components and establish safety design guidelines for human-robot interaction in VR. An extensive evaluation of the technical possibilities is presented in section 5. Results from a user study conducted with the system are given in section 6.

3 Related Work

Providing haptic feedback in VR is a very active and broad field of research. We will focus on closely related work, i.e. robots for haptic interaction and passive haptics.

Figure 2: Our fully immersive VR system: a) Cyberith Virtualizer, b) Oculus Rift DK2 with Leap Motion in front and 7.1 audio headset, c) Perception Neuron motion suit with over-gloves, d) perpendicular for initialization, e) Crustcrawler Pro robotic arm.

For a more extensive overview on haptic devices we recommend the works of Bowman et al. [8] and Burdea & Coiffet [9], as well as a more recent survey on large workspace haptic devices by Gosselin et al. [10].

A popular approach to provide haptic feedback for arbitrary virtual objects are single point haptic devices, e.g. the popular PHANToM models [11], [12]. However, the limitation to a single feedback point limits the applicability for a number of scenarios where the feedback should be close to the way we interact with real world objects. In a pioneer work Hirota & Hirose [13] suggested to use a real object to simulate virtual ones and called the concept Surface Display. In a later prototype [14] they showed a way to simulate curved surfaces on a small contact area. Both examples allowed only very restricted movements with a single finger, while our approach provides unrestricted haptic interaction with both hands and free movement in the virtual world.

McNeely [15] speculated about the use of robots to provide force feedback for VR and called the concept Robotic Graphics. The user does not have to hold a haptic device all the time, but the system places it at the desired location and waits for the user to encounter it. This kind of feedback is also referred to as an encounter-type haptic device. Since this approach has the potential to provide free hand and real touch sensations, we believe there is great potential in revisiting this approach for application in VR systems. Yokokohji et al. [4] demonstrated this potential in an early experiment: They suggest that a plate mounted on a robot arm could be optically tracked and visually registered with an HMD. The user's hand could be segmented from a video image and in that way he or she would be able to feel the side of a virtual cube when touching the plate. They realized a prototype with an LCD panel in front of the assembly and a plate that could be manipulated over a sensor knob on it. In this implementation the position of the physical object defines the position of a virtual overlay and the robot arm resists with a certain force to simulate the physical behavior. Later [16] optical hand tracking with a marker on the hand was added and incorporated into a path planning algorithm to avoid collisions. In that way two objects are simulated with the same sensor knob, depending on the distance to the hand. In both examples the position of the virtual objects is only changed over the sensor knob, while in our system the user interacts directly with matching physical props in order to feel the inherent properties of the touched object. In more recent work [5] the group optimized the path planning algorithm to reduce device duty. They presented a conceptual sketch to use a tool with three different buttons mounted on the robot arm to simulate a virtual control panel, but it was not implemented.

Shin et al. [17] present a one degree-of-freedom system to allow a highly realistic haptic simulation of a refrigerator door. However the system is specialized for this specific use case and not suitable for more generic application scenarios. The TouchMover [18] constitutes a robotic assembly actuating a touch screen with stereoscopic rendering. Moving the screen towards or away from the user the system can render realistic physical behavior and large forces of a presented virtual object, but it is restricted to one dimension and only a single touch point. For the simulation of arbitrary objects Yokokohji et al. [3] conceptualized a device with three actuated surface patches configured to accommodate the most likely grasping behaviors with three fingers. If mounted on a robotic arm the authors suggest it would be possible to allow grasping an arbitrary virtual object with thumb, index and middle finger. However, they also state it would be extremely difficult to track the fingers precisely and interaction would be limited to three fingers. In our system the haptic feedback is not limited to one dimension or individual fingers, a presented virtual object can be explored in its entirety with both hands.

Tachi et al. [19] experimented with a specially shaped object mounted on a robot arm they called Shape Approximation Device. It provided tilted surfaces to represent the edges of more complex virtual objects by orienting the appropriate face of the object towards the user. A mechanical exoskeleton is required to acquire the arm and finger position of the user and only a single contact point can be provided. Another haptic VR system proposed by Ortega & Coquillart [20] employs a system of motorized strings attached to a hand-held physical prop. The prop was more complex and could be handled with the whole hand. While these systems allow more complex objects, the necessary mechanical exoskeleton or the motorized strings restrict user movements severely. With our system the user is not impeded by an external structure and can move and turn freely.

In recent work Yamaguchi et al [21] use a drone to align a physical sheet of paper with a virtual object in an immersive environment. With a stick-shaped grasping object it is possible to touch the paper and thereby perceive haptic feedback. This allows a reasonable working volume, but the positioning is imprecise, the presentable forces are limited and the use of a grasping object negates the potential for free hand interaction. In their Snake Charmer Araujo et al. [22] employ a robot arm which can be fitted with a variety of endpoints with different surface textures on each side. Tracking the hands of a user it is possible to switch between different object positions. The surface with the appropriate material is oriented towards the user and aligned with the virtual object, so the user can feel different surface characteristics of virtual objects. However, the interaction volume is only around 50 cm^3, while our system provides an interactive volume covering most of arm's reach in virtual environments of infinite size.

Moreover, the situation in our approach is different to all the previously presented works. Up to now the virtual and the physical world were aligned, the mapping between the worlds was fixed and the physical object only had to switch between different predefined positions. With our approach the relation between the virtual and the physical world is constantly changing. We dynamically update the position of the physical object according to the current position of the user in the virtual world, which is not restricted to a previously known set of coordinates.

Apart from encounter-type haptic devices, there are also a number of systems that employ passive physical props for haptic feedback in VR. The virtual objects are not simulated but the presented physical prop already resembles the desired physical properties as close as possible. In that way it is possible to provide very realistic robust feedback, including texture, shape, and intuitive manipulation. The use of physical props was introduced by Hinckley et al [23] as an intuitive interface for neurosurgeons on a conventional monitor. Since then several authors have proven the effectiveness and realism of virtual environments augmented with physical props for haptic feedback, for example in simulating a virtual pit [24] or in creating a life-sized virtual environment with Styrofoam blocks in combination with projectors [25].

Since the physical props are passive, they would have to be available in every form and on every position required by the simulation. This means that scaling of the virtual environment is problematic, because of the required number of props as well as physical space. Cheng et al. [26] suggested to employ people instructed over projectors and sound to actuate physical props. This allows the scaling of the virtual world independent of the number of props, since they only have to be available where the user wearing an HMD can actually interact with them. The physical objects can be reused for unlimited virtual objects. Realism and presence are reported to be very high. However, space requirements remain to be extensive and employing a number of trained persons for augmentation for every use of the system can also be costly and inefficient.

We propose a fully immersive VR system (Figure 2) that differs from all previously published works. For haptic feedback we build upon the benefits of encounter-type devices as well as prop-based haptic feedback without suffering from existing limitations regarding the interaction volume, restricted movements, physical space requirements or scalability. We also address new challenges concerning collision avoidance and safety as well as constantly updating the relative position of the physical object according to the current position of the user in VR. In contrast to related work our system is more generic and allows free movement through virtual environments of unlimited size with close-to-natural walking. Our concept has the potential of full haptic coverage within arm's reach, while interaction is possible with bare fingers or any body part.

4 Concept & Implementation

4.1 Hardware

In order to enable close-to-natural and unrestricted movement through an arbitrarily sized virtual world, we employ the Cyberith Virtualizer omni-directional locomotion platform (Figure 2a). The user is fixated around the hips in the middle of a ring with a harness. Due to low friction it is possible to walk by slipping over the base plate with socks or overshoes while being held in place. Optical sensors measure direction and speed of the moving feet. The ring can turn freely and sensors within measure the directional orientation of the body. Together with data from the base, this provides correct relative body movement including side steps or backwards movements in a virtual world, uncoupled from the gaze direction. Furthermore sensors in the pillars reporting the height of the ring enable jumping or crouching.

For an immersive stereoscopic first person view of the virtual world we employ the Oculus Rift Development Kit 2 (DK2) HMD (Figure 2b). Translational head movements are optically tracked with a camera and infrared LEDs embedded in the HMD. Directional sound is presented with a 7.1 surround sound headset with passive noise cancelling to prevent any outside interruption. The cables are tied together on a retractable strap centrally above the Virtualizer to allow unencumbered 360° turning.

All necessary processing is performed on a 4 GHz Intel Core i7 CPU with 32 GB RAM and a GeForce GTX 980 Ti graphics card.

4.1.1 Robotic Arm

In order to actuate physical props for haptic feedback, the CrustCrawler Pro-Series robotic arm is employed in this project. It is custom-assembled with the intention to provide an appropriate length for a suitable work envelope while at the same time maintaining a reasonable tradeoff regarding the possible payload.

The current configuration is presented in Figure 3. From the bottom up it is consisting of a turntable with one Robotis Dynamixel MX-106 smart actuator inside for rotation and two parallel MX-106 on top, two parallel MX-64, two MX-64 (rotation & tilt), two MX-28 (rotation & tilt) and finally a gripper based on one AX18A. Aluminum girders of different lengths connect all components, resulting in an overall length of 1170 mm (960 mm reach). The robot arm is powered with maximal 222 W at 12 V. A summary of the technical specifications of the individual actuators is presented in Table 1.

We calculated that the motors would stall at a maximum load of 1.04 kg (@12 V). However, in practice if the actuated weight approaches the maximum, there is the risk that the motors overheat under longer continuous load or fast accelerations, causing an emergency shutdown. For this reason it is advisable to stay well below the maximum. As a precaution we attached additional cooling fans to the motors that are expected to bear most of the load. Each actuator provides its own motor controller. All of them are daisy-chained into a serial bus which is connected with a USB2Dynamixel adapter via USB to a PC providing all necessary control operations.

4.1.2 Human-Robot Interaction

In order to control a virtual avatar as well as for safety reasons (see section 4.3), the user is equipped with a Perception Neuron inertial motion suit (Figure 2c). 17 inertial sensors report the pose of every part of the body with 120 Hz. In order to exactly match the size of the virtual avatar to each user, the body height, the shoulder width as well as the hands, the upper and the lower arms are measured by hand.

Especially when looking at the hands while interacting with a virtual object, precise registration of the user's fingers in the virtual world is crucial. For that reason we mounted a Leap Motion on the front panel of the DK2 (Figure 2b). It tracks the hands and fingers with up to 200 Hz while keeping them free of any additional hardware.

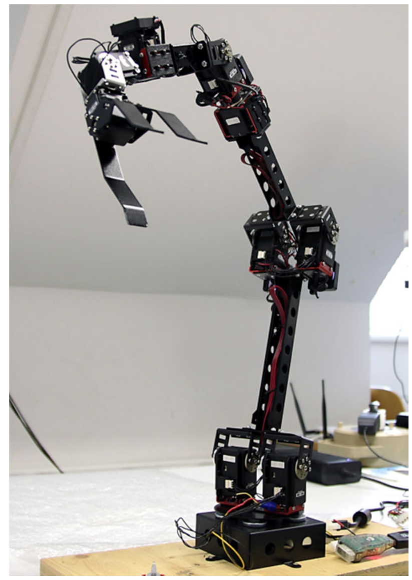

Figure 3: Robot arm configuration for our system.

Table 1: Technical specifications of actuators (@12 V / no load).

Type	Torque	Speed	Resolution
AX18A	1.8 Nm	582°/s	0.29°
MX-28	2.5 Nm	330°/s	0.088°
MX-64	6 Nm	378°/s	0.088°
MX-106	8.4 Nm	270°/s	0.088°

We implemented our system to switch between the inertial tracking provided by the motion suit and the more precise optical tracking by the Leap Motion when looking at the hands. During development it turned out, that the sensor mountings of the Perception Neuron suit absorb IR light to an extent that interferes with proper recognition of the hands by the Leap Motion. On this account we employ tight-fitting and breathable white over-gloves to cover the sensor mountings (Figure 2c).

The provided poses are relative to the head position reported by the camera tracking of the DK2, which is registered with the virtual environment on startup. For that reason it is important to position the HMD exactly in the correct distance from the camera during initialization. In order to facilitate this process we mounted an adjustable perpendicular (Figure 2d) to mark the exact position for the face plate of the DK2 and hold it in place during startup.

Finally, the key component is the employed 7-axis Crustcrawler Pro-Series robotic arm (Figure 2e), which enables our VR system to actuate physical props for haptic feedback. As soon as the user is in range of a virtual object he or she might want to interact with, the robot arm can select an appropriate physical prop and present it, dynamically matching pose and location of the object in respect to the user's current position in the virtual world. In that way when the user reaches for the virtual object, his or her hands also encounter it in the real physical space. Since the user's hip is fixated in the Virtualizer, the geometric relation to the robot arm in physical space is known. Thus the position of the physical prop only has to be updated when the user performs an action that changes his or her location in the virtual world, e.g. takes a step. Any other movement, e.g. head, upper body or legs, does not change the geometric relation in physical space. This is an advantage, because people tend to remain in the same location while exploring an object.

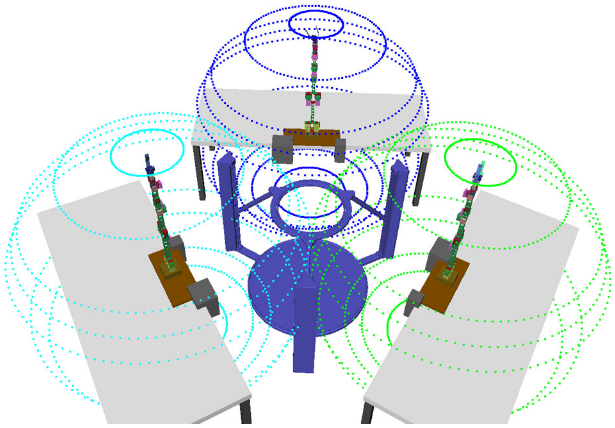

Figure 4: Three robot arms providing 360° haptic coverage.

Figure 4 shows the principle of employing three robotic arms with overlapping operating radii arranged around the locomotion platform, which would allow the exploration of infinite virtual worlds with 360° haptic coverage. However, for the moment we only employ one arm in order to proof the validity of our concept.

4.2 Software

4.2.1 Robot Control

The individual motor controllers on the actuators have to be coordinated over a main controller, which we implemented in C++. This controller runs on a PC and communicates with the individual actuators over the asynchronous serial Communication Protocol 1.0 of the Dynamixel SDK. In that way all available parameters can be read or set, for example the desired joint angle as well as limits for speed, angle or torque. On that basis we implemented the controller with a wide variety of functions: It is possible to move the robot arm in the work envelope via xyz-coordinates and define the orientation of the end effector. If desired, velocity and acceleration can be adjusted. We also provide commands for different pre-defined movements, for example opening and closing the gripper or taking an initial pose.

The joint angles of the individual motors required for a specific pose are calculated via Inverse Kinematic (IK). Since this has to be done each time the arm must take a new pose, this process has to be fast and optimized for a specific robotic arm. For that purpose we integrated the Open Robotics Automation Virtual Environment (OpenRAVE) [27]. This framework provides an environment for simulation and analysis of kinematic information related to the development of motion planning algorithms for example used in industrial automation.

A 3D model defining the used robot arm with all necessary information has been created. This includes the physical dimensions as well as the geometric and kinematic structure, e.g. the available joints as well as their angle and speed limitations. In the same way the relevant physical environment is defined, in our case the table, the Virtualizer, the physical props and an adjustable humanoid model representing the user. In order to avoid any collisions with itself, all the information about the user or parts of the environment are taken into account every time the IKs of a movement path for the robot arm are calculated. A calculated path by OpenRAVE is sampled as 15 ms segments by our controller which updates the corresponding actuator angles accordingly. The controller is designed to handle multiple robots at once, as presented in the principle illustrated in Figure 4. While the path planning itself runs in a single thread (since all calculations are interdependent), handling the movement along the path segments of a calculated path is done in parallel to the planning. We also integrated a Windows Communication Foundation (WCF) server to support future distributed operation over network.

4.2.2 Game Engine Integration

For general and convenient use we provided a C# wrapper for most of the controller functions and integrated it into a package for the game development engine Unity3D. There is a template for a virtual scene that includes models of the physical setup. The user's avatar is fully customizable to different body dimensions. Its position and orientation in the virtual world is determined by the analog signal from the Virtualizer, describing the movement of the user in the physical world. The Perception Neuron Unity SDK is integrated to apply the motions captured by the motion suit to the limbs and posture of the avatar. The models in the Unity scene are always in synchronization with their counterparts in OpenRAVE, which are used for path planning of the robotic arm. The scene is configured to render the virtual world on the Oculus Rift HMD and incorporates the Leap Motion Orion SDK to provide finger tracking. The developed Unity package also provides a template that contains all required components to fit any object in the virtual world with a haptic representation. As soon as the user's avatar approaches a specified distance to such an object, the robot arm is assigned exclusively to it. If necessary the arm can pick up the appropriate physical prop, and then takes a standby position outside reach, in a pose that is most likely ideal depending on the approach vector of the user.

When the avatar closes to the user's reach, the robot arm presents the physical prop attached to the gripper, matching the correct position relative to the user to the corresponding object in the virtual world. If the path the robot arm takes to this position would last longer than a certain timespan (typically 2 s), we display a loading icon at the location of the virtual object to signal that the physical representation is not ready yet. For shorter paths this is not necessary. If a specific path is desired or if a required position is known in advance, it is possible to calculate a path in advance, which speeds up the response time of the arm since the IKs don't have to be calculated at runtime. When the avatar moves out of range, the robot arm moves to an idle position and access is released to all objects again.

4.3 Safety Guidelines

With every system involving robots or robotic arms, especially in close collaboration with humans, safety is always of particular importance. We have taken this concern very seriously and assessed possible risks for a user of the system thoroughly. In order to ensure the user's safety we provided precautionary measures and inherent safety on several levels:

User monitoring:
Before entering the system, the user's body dimensions are measured and the avatar is tailored to fit. The motion suit reports the current body pose with 120 fps (latency ~20 ms) and the planner calculates the paths for the robot arm accordingly, ensuring a safety margin of about 5 cm to the user. A quick second collision check confirms a clear path before the actual movement of the arm starts, stopping immediately if a possible collision is detected. The path is rechecked for a specified timespan to give the user the opportunity to move away, should the collision persist the system stops with an error.

Safe design:
We designed the system to provide passive safety. The low weight (including turntable about 1.9 kg) of the robot arm ensures low

inertia forces. The choice of the length of the arm and its distance to the Virtualizer prevents it to reach the area around the head in a normal upright pose, as long as the user is not extremely small. The minimum and maximum angles of the joints are set on each motor controller in a way to prevent the user's hands from being trapped or pinched between them.

Low speed:
We limit the maximum rotational speed of the individual actuators to 40°/s. In that way in case of a collision, there is not enough inertia to hurt the user due to the low weight and speed.

Torque limits:
If despite all safety measures a collision should occur or the user is somehow caught between joints, it would cause the measured torque in one of the actuators to exceed the designated safety limits. In that case it would signal an overload error and cut power to the motor immediately.

Emergency switch:
In all our experiments there is always a human operator present. Should the operator observe an unexpected situation or an emergency, it is possible to cut power to the whole system by flicking a single switch.

5 Technical Evaluation

In order to evaluate the robot's reachability, we calculated the possible IK solutions via OpenRAVE for points within reach of our robot arm. Reachability depends on the robot's kinematic constraints. Figure 5 illustrates reachable points in a plane through the center of the work envelope, whereas the red lines indicate the possible orientations of the end effector at each point.

Further characteristics of our system that affect the ability to present a convincing haptic representation of an object include the accuracy and repeatability with which a target point within the work volume can be reached. We followed the procedure described by Nicolescu et al. [28] to approach 44 measurement points which are evenly distributed with a distance of 150 mm along the edges of an imaginary cube with a side length of 600 mm positioned within the work envelope.

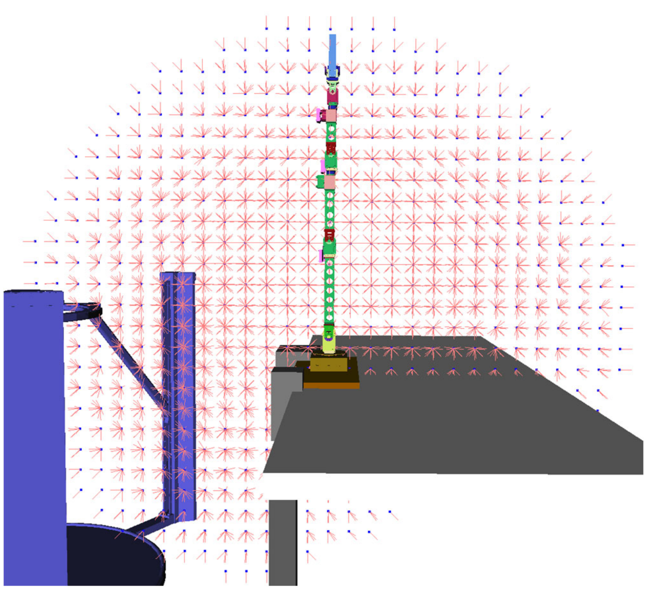

Figure 5: Reachability of the robotic arm.

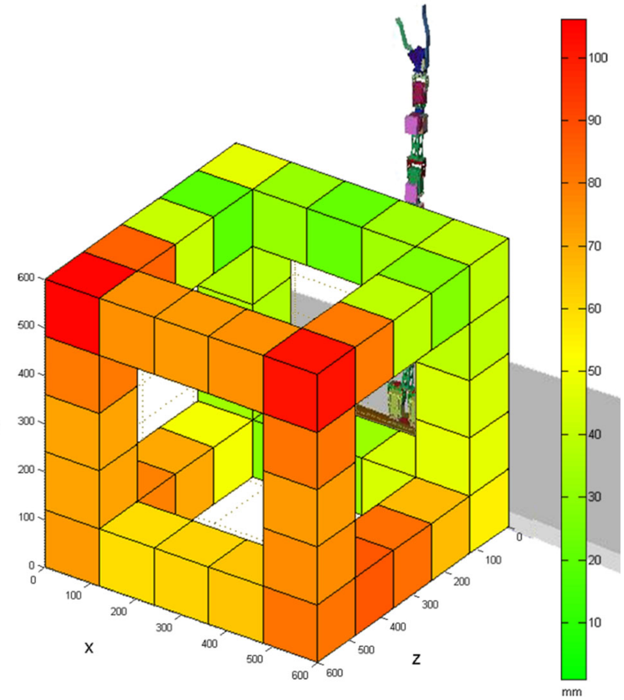

Figure 6: Error map for accuracy in work envelope.

Between every target on the cube the arm returned to the home position, which means standing completely upright. For each nominal coordinate set that the robot arm was instructed to reach, we measured the actual coordinates it arrived at in real space and read out the angles reported back by the encoders in the motor controllers. For measuring the actual coordinates, we employed the commercial tracking system ARTTRACK1 by Advanced Realtime Tracking GmbH. Our system uses 3 infrared cameras to provide the position of passive retro-reflective markers with 60 Hz and an accuracy of around 1 mm over a room-sized tracking volume and less than 0.5 mm in the center [29]. We tracked a single marker at the tip of the arm's end effector, and calibrated the system with the center inside the measurement cube directly before the procedure to achieve the highest possible accuracy. In that way we measured a mean accuracy (displacement between nominal and real coordinates of the end effector) of our system for each coordinate axis of X = 15.33 mm (SD 14.50), Y = 46.20 mm (SD 24.43) and Z = 11.30 mm (SD 23.03). The mean absolute accuracy is 57.81 mm (SD 22.20). Since the accuracy is not uniform over the work envelope we presented an error map relating the accuracy with the position in the work volume graphically (Figure 6).
The difference between the nominal coordinates and the coordinates derived from the reported actuator angles results in the mean error of the position encoders and the used IK model for each coordinate axis of X = 0.22 mm (SD 5.26), Y = 23.40 mm (SD 12.58) and Z = -2.93 mm (SD 11.07), absolute 25.65 mm (SD 14.38). The difference between the coordinates derived from the actuators and the measured real coordinates shows the mean error caused by the kinematic chain for each coordinate axis of X = 15.11 mm (SD 11.35), Y = 22.80 mm (SD 14.18) and Z = 14.23 mm (SD 13.01), absolute 36.68 mm (SD 10.22).

Since the accuracy depends to a large extent on the current load, the repeatability is often used as a measure for the quality of a robotic arm. It is a lot less susceptible to the current load because it does not describe the actual error of a position, but the difference in error every time this position is reached.

Table 2: Mean timings for paths of different length.

Path length	Planning	Travelling	Response
Short (150 mm)	0.83 s (SD 0.69)	2.75 s (SD 2.05)	3.58 s (SD 2.72)
Medium (600 - 671 mm)	2.43 s (SD 2.96)	3.73 s (SD 1.12)	6.17 s (SD 2.89)
Long (848 - 1039 mm)	1.6 s (SD 0.42)	4.10 s (SD 1.75)	5.70 s (SD 2.12)

The absolute repeatability for our system, derived from the radius of the sphere containing all 44 real measurements of the home position, is 16.26 mm.

Finally the response time is an important property of our system, since it describes the total time until the haptic feedback is ready for a certain location. The response time is the sum of the time for planning a path to a desired pose and the travel time to move along the path to the destination. For that reason it is highly dependent on the length and complexity of the path as well as the speed limit. In order to estimate typical response times in a realistic setting, we used the imaginary cube from the accuracy measurements (without returning to the home position), and approached all 44 points moving from point to point in a combination of short (150 mm), medium (600 - 671 mm) and long (848 - 1039 mm) distances.

Although it could move much faster, we limited the arm to the same low speed we employ for safety reasons during normal use. The mean timings we observed with this procedure are presented in Table 2. The profound role of the complexity of the path shows in the minimum and maximum response times for the short paths: The minimum response time was 0.74 s, while the maximum response time was 12.43 s. However, both were rather extreme cases which are not typical.

6 Subjective Evaluation

In order to make a statement about the subjective experience when using our system, we conducted an experiment with 34 participants (age 14 - 58, 12 female). On a scale of 1 to 5 (5 is the highest), all of them reported regular computer use of at least 3 (mean 4.5), however only 14 stated any prior VR experience and 2 of these higher than average (mean 1.7). We wanted to investigate the acceptance of the system as well as how convincing the presented objects and the interaction was compared to real life. In the following we will present the design and procedure of the experiment and report on results.

6.1 Experiment Design & Procedure

In our experiment the participants were randomly assigned to one of two groups. Everyone experienced two short VR scenarios with our system, depending on the group one of the scenarios with (condition 1: haptic) and the other one without haptic feedback (condition 2: non-haptic).

The first scenario took place in a virtual garden with a 5 m by 2.65 m wooden wall at one side (Figure 1: Left). The participant's task was to approach the wall and to touch a white sparkling spot that showed on its surface. When the spot was touched it changed color to green and was held for two seconds. Then it disappeared and another spot showed on another location on the wall. This was repeated until 7 points distributed over the wall had been touched. For the haptic condition (group 1) the robot arm moved to each location and provided physical feedback for the area around the sparkling spots with a lightweight 300 mm by 300 mm wooden panel fixated on the gripper.

Figure 7: Guiding avatar and approaching pedestrian in city.

In the second scenario the test subject had to walk along the pavement of a city (altogether 123 m), following a guiding avatar which ensured a consistent walking speed and path for all tests (Figure 7). Over the last 76 m a total of 10 pedestrians approached from the opposite direction, 4 of them programmed to avoid and the other 6 to make sure to bump into the user when passing by. In this scenario the gripper of the robot arm was fitted with a boxing glove to provide physical feedback for the shoulders of the pedestrians in the haptic condition (group 2). The glove added 200 mm of soft padding to the length of the arm, thereby permitting direct contact without compromising safety.

For both scenarios the robot arm was in upright home position before beginning, regardless of the condition. The physical prop was fixated after the participant had put on the HMD and the headphones, so he or she was unaware of the two experimental conditions. There were two experimenters present, one was guiding the test subjects through the experiment, and the other one was operating the system and the emergency stop. Every test was recorded on video and the screen output was captured.

The procedure for a test session started with a verbal introduction by the experimenter. The test subject was informed that the robot arm would provide haptic feedback in certain situations and that the experiment would help to evaluate this kind of virtual experience and the system. He or she was made aware of the possibility of a remaining risk and it was made clear that the experiment can be aborted at any time if desired. After that the participant had to sign a consent form, confirming that the risks were understood and accepted. Subsequently a preliminary questionnaire and a simulator sickness pre-questionnaire (SSQ) [30] had to be answered. After that the participant's body dimensions were measured. The HMD and the motion suit were put on and both were calibrated. The user entered the Virtualizer and was encouraged to accustom him- or herself with walking on the locomotion platform. In the meantime the second experimenter fixated the appropriate physical prop on the gripper. After a short explanation of the task, the first scenario, the garden, was started. After completing the task, the participant could continue to explore the environment if desired. Then he or she put off the HMD and the headphones and was asked to answer the SSQ again. Next the HMD was put on again and the procedure was repeated for the second task and scenario, the city. After that the test subject put off all devices and was asked to fill out the SSQ once more as well as an extensive questionnaire. Typically a complete test session including the preparations and the questionnaire lasted about 1 hour, each VR scenario around 5 minutes.

6.2 Results

The questionnaire administered after our experiment primarily used Likert scales with up to 6 items each, arranged in a way to have at least two questions between corresponding items. Altogether the questionnaire contained almost 60 questions, excluding the SSQ and the general questions.

The SSQ had to be answered before, in between and after each scenario. The resulting symptom variable scores were weighted and summed in accordance with Kennedy et al. [30] to obtain a total simulator sickness score after each scenario. The mean differences after the first scenario were -35.94 (SD 148.99) and after the second scenario 35.62 (SD 175.71). Considering that in sum the scores remained almost constant, there are no specific difficulties indicated concerning simulator sickness. A number of participants reported verbally that they actually felt better after using the system, aside from some sweating from the over 100 m virtual walk. Only one test subject did voice slight symptoms after the test and also had relatively high SSQ scores.

We were interested if the system could be conceived as dangerous, so we asked on a 5-point scale (1 = Not at all, 5 = Very much) if the participants trust the system they were using as well as if they feel uneasy interacting with the robot arm. They reported a high mean trust with 4.53 (SD 0.88) and low mean uneasiness 1.33 (SD 0.65). We also hypothesized that after using our system, the trust would increase and the uneasiness would decrease and therefore asked the same questions before the test as well. However, the mean difference in trust rating was 0.00 (SD 0.98) and in uneasiness -0.36 (SD 1.11). A Wilcoxon signed-rank test showed no statistically significant difference in trust ($Z = -1,805$, $p = .071$) nor in uneasiness ($Z = -0,110$, $p = .912$).

We gave the respondents the opportunity to expand on their choice: Out of six answers, all confirmed that they felt safe in the system and three added as a reason that the arm did not feel too stiff.

In order to find out how walking in the system was conceived, we used a 7-point Likert scale with 6 items based on questions by Slater et al. [31]. After each scenario we asked how complicated or simple, unnatural or natural and difficult or straightforward moving in the virtual world was (7 always the most positive). The reliability was tested with Cronbach's α ($\alpha = .891$) and the resulting mean was 4.33 (SD 1.17).

We also wanted to know how convincing virtual objects presented in our system are and formulated the hypothesis that they would be more convincing with haptic feedback than without. For that purpose the participants were assigned to either experience haptic feedback for the wall in the garden (group 1) or haptic feedback for the people in the virtual city (group 2). The dependent variables were measured on 7-point Likert scales with 2 items each, the reliability was tested with Cronbach's α. The items were how credible and how realistic the wall/people felt compared to a real object/person (1 = Very noncredible/unrealistic, 7 = Very credible/realistic).

The results are presented in Figure 8 and show mean ratings for the haptic conditions for the wall of 4.78 (SD 1.50, $\alpha = .906$) and for the people of 4.30 (SD 1.53, $\alpha = .813$). However, compared to the non-haptic conditions for the wall of mean 3.84 (SD 2.04, $t(32) = -1.534$, $p = .135$) and for the people of 3.78 (SD 1.53, $t(31) = 1.056$, $p = .299$) the differences were not statistically significant in both cases.

Furthermore we investigated how the act of touching the wall and the incidence of colliding with people was conceived by the users of our system compared to real life. We also formulated the hypothesis that they would be more similar with haptic feedback than without.

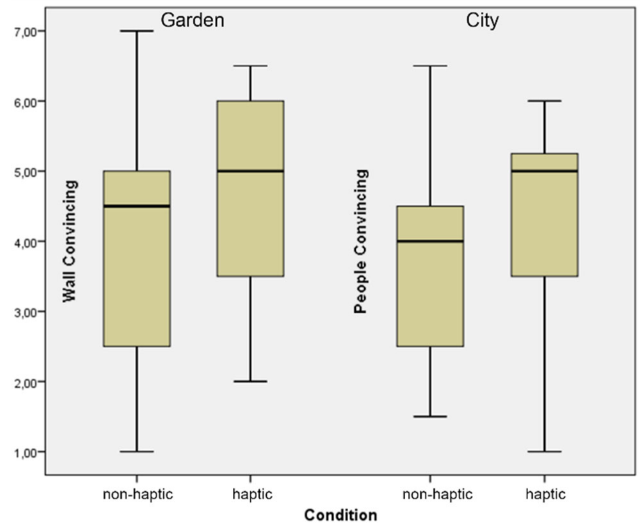

Figure 8: How convincing were the virtual objects.

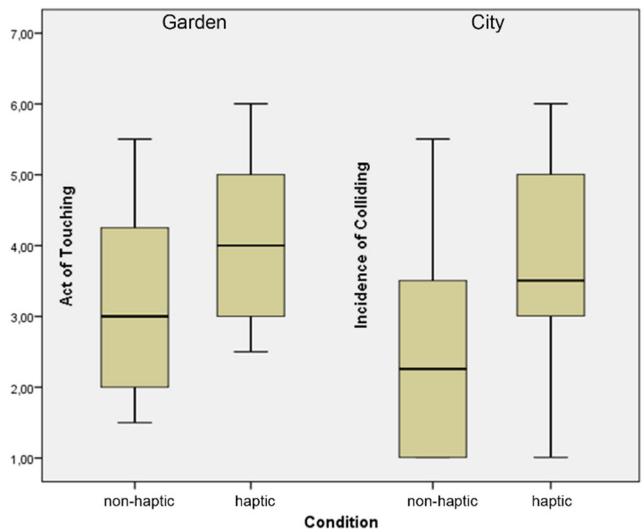

Figure 9: How similar were the interactions compared to real life.

The items were how natural and how similar the act of touching/colliding with people seemed compared to such situation in real life (1 = Very unnatural/Not at all similar, 7 = Very natural/similar). The results are presented in Figure 9 and show mean ratings for the haptic conditions for the act of touching of 4.03 (SD 1.32, $\alpha = .460$) and the colliding with people of 3.63 (SD 1.57, $\alpha = .739$). Compared to the non-haptic conditions for the act of touching with mean 3.09 (SD 1.32, $t(32) = -2.189$, $p = .036$) and the colliding with people 2.58 (SD 1.39, $t(31) = 2.045$, $p = .049$) the differences were statistically significant in both cases. However, the low reliability for the act of touching indicates that the participants discerned between the act of touching being natural and being similar to real life.

Finally we asked the participants to report the most positive and the most negative incidents experienced during the experiments. Altogether 30 persons stated positive incidents and 20 negative ones. The haptic feedback was especially praised 10x (5x wall, 5x people), on the other side there was 1x complaint that the actual touch of the robot arm for the collision with the virtual people lasted too long. The matching of hands and limbs was complimented 7x, while 3x a negative incident was reported in that context. The walking in the virtual world was stated positively 4x and 9x negatively (4x the reason was that it felt not natural). Additionally 4x it was complained that the virtual people could not be evaded (which they were programmed not to) and 2x that they walked through oneself and did not bounce back.

7 Discussion & Limits

We demonstrated the capabilities of our system in the technical as well as the subjective evaluation. 34 users experienced two fully immersive virtual worlds in which they could move freely and physically interact with the environment. In our confined lab space they walked a distance of over 100 m through a virtual city and rated the form of movement very positively. At the same time it was reported several times as a negative incident, mostly because it did not feel as natural as expected. Moreover, we did not observe any noteworthy difficulties especially considering that almost all of the participants used a locomotion platform for the first time.

Our established safety guidelines proved adequate for the system, which is also reflected by the fact that very high trust in the system was reported and the participants indicated very low uneasiness interacting with physical props provided by a robotic arm. The fact that there was no significant improvement in trust as well as uneasiness after using the system could be due to the already high ratings before use.

In our experiments we showed that the system can provide haptic feedback for virtual objects of completely different nature. In case of the wall the participants interacted intentionally and with their hands and fingers, two users even used both hands at the same time. Whereas in the city they interacted unintentionally with their shoulders with the virtual people. The results from the subjective evaluation show that the virtual wall as well as the people were quite convincing. Although a tendency is indicated that the haptic feedback contributed to make the virtual objects convincing, the difference to the non-haptic condition was not significant. This would suggest that the experience was also quite convincing without haptic feedback, which is not surprising considering that almost none of the participants had any comparable VR experience before.

The act of touching the virtual wall was reported as relatively natural, whereas colliding with people was rated lower in comparison. A possible reason could be that the users were irritated by the virtual pedestrians walking right through the user's avatar, which was also stated as a negative incident twice. Nevertheless, the results supported our hypothesis that the provided haptic feedback contributed significantly to make the act of touching as well as the colliding more similar to real life.

Although the measured mean accuracy of the robotic actuation seems not particularly high, it has to be considered that it is worse if far from the center of the working volume. However, most interaction takes place in front of the user at middle height, where he or she can easily reach. The mean accuracy in this region is around 25 mm, which is enough for low fidelity objects like a wall. The results also show that a considerable part of the impreciseness is caused by the kinematic chain, for example bending of the frame. Moreover, the repeatability of 16 mm indicates that it might be possible to improve the accuracy by calibration.

The timings presented in the technical evaluation represent a situation without any prior knowledge of the situation, where the actual response times can be unpredictable, because they depend highly on the complexity of the required path. However, this is usually not the case and the response time can be reduced considerably. For example for the wall in the garden the arm takes an ideal waiting position on approach of the user while he or she is not in reaching distance. After that there are often only minor adjustments necessary which can be accomplished below 1 s. Another way is to calculate the path in advance or in parallel to a previous movement when possible, which reduces the time for the path calculation to a fraction of a second. In the virtual city an ideal waiting pose and pre-calculated paths permitted a fixed response time of 2.00 s for a medium distance (521 mm).

There are also limitations of what is possible with our current implementation. Although the employed robot arm is reasonably priced (~4100 USD) compared to other arms, the permissible loads as well as the gripping power are not very high. For that reason we limit the weight of physical props to around 100 g for reliable operation. Furthermore, for obvious reasons it is only possible to present one object at a time. On the other side according to Hoffman [32] even a single haptic interaction can improve a user's assumptions about the solidity of all objects in the virtual world.

Additionally if a user wants to make fast movements or direction changes, the times for path calculation and moving safely to a destination do not allow for adequate response times. During our experiments we also observed one participant who wanted to walk further towards the virtual wall while being in touch with it, which is not supported due to our safety measures.

Finally, since every inertial system is prone to drift, the inertial motion suit would have to be recalibrated for extended use of the system.

8 Conclusion & Future Work

We presented the concept, the hardware and the software architecture as well as safety guidelines for a fully immersive VR system with haptic feedback based on robotic actuation of physical props. We demonstrated the capabilities of our system in our experiments. On less than 4 m by 4 m of lab space our participants experienced virtual worlds where they walked over 100 m in a close-to-natural fashion, they touched a virtual wall with bare fingers and bumped into virtual pedestrians. The results revealed in our evaluation are very encouraging and demonstrate the use of our system as a versatile research tool for haptic experiences.

However, we also see room for improvement: For example the response times for the haptic feedback could be considerably improved by dynamically adapting the speed of the robotic arm, so our speed limit would only apply for a certain safety zone around the user and the arm could move faster while out of reach. We want to get the response times fast enough to allow following the current position of the user's hand to provide feedback wherever he or she wants to interact. Currently we are also investigating alternative robot arms providing more load, more accurate actuators and more rigid links. We also consider adding a circular rail around the Virtualizer as additional axis, which would eliminate the necessity of traversing the arm's own center as well as permit 360° haptic feedback with only one robot. Furthermore the possible drift as well as the additional preparation time introduced by the inertial motion suit led us to investigate possible optical solutions to provide a reliable measure when the drift gets unacceptable and at some point replace the inertial suit altogether.

In the future we want to take advantage of the available physics simulation in the game engine, which would allow us to augment the virtual objects with physical properties. If the virtual object is pushed away, the physical object matches its behavior since the position and orientation is synchronized. We also consider including features like vibration and dynamic movements to simulate tactile properties and surface geometry for example the bark of a tree. Finally it will be possible to present high fidelity objects like a control panel or extend the actuated physical props with electronic components ranging from buttons to medical sensors, which opens up a whole range of new opportunities.

ACKNOWLEDGEMENTS

This work was partially sponsored by the Department of the Navy, Office of Naval Research (Foreign Research Grant). We would like to express special thanks to all the people who participated in our experiment.

REFERENCES

[1] S. J. Biggs and M. A. Srinivasan, “Haptic interfaces,” *Handbook of Virtual Environments*, pp. 25–74, 2002.

[2] E.-L. Sallnäs, K. Rassmus-Gröhn, and C. Sjöström, “Supporting presence in collaborative environments by haptic force feedback,” *ACM Trans. Comput. Interact.*, vol. 7, no. 4, pp. 461–476, Dec. 2000.

[3] Y. Yokokohji, N. Muramori, Y. Sato, and T. Yoshikawa, “Designing an encountered-type haptic display for multiple fingertip contacts based on the observation of human grasping behavior,” in *Proc. of HAPTICS '04*, 2004, pp. 66–73.

[4] Y. Yokokohji, R. L. Hollis, and T. Kanade, “What You Can See is What You Can Feel - Development of a Visual/Haptic Interface to Virtual Environment,” in *Proc. of IEEE Virtual Reality '96*, 1996, pp. 46–53.

[5] M. Yafune and Y. Yokokohji, “Haptically rendering different switches arranged on a virtual control panel by using an encountered-type haptic device,” in *Proc. of IEEE WHC '11*, 2011, pp. 551–556.

[6] R. W. Lindeman, J. L. Sibert, and J. K. Hahn, “Hand-held windows: towards effective 2D interaction in immersive virtual environments,” in *Proc. of IEEE Virtual Reality '99*, 1999, pp. 205–212.

[7] B. E. Insko, “Passive haptics significantly enhances virtual environments,” University of North Carolina at Chapel Hill, 2001.

[8] D. A. Bowman, E. Kruijff, J. J. LaViola Jr, and I. Poupyrev, *3D user interfaces: theory and practice*. Addison-Wesley, 2005.

[9] G. C. Burdea and P. Coiffet, *Virtual Reality Technology*, 2nd ed. New York, NY, USA: John Wiley & Sons, Inc., 2003.

[10] F. Gosselin, C. Andriot, J. Savall, and J. Martín, “Large Workspace Haptic Devices for Human-Scale Interaction: A Survey,” in *Haptics: Perception, Devices and Scenarios*, Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 523–528.

[11] T. Massie and J. Salisbury, “The PHANTOM Haptic Interface : A Device for Probing Virtual Objects,” in *Proc. of ASME Dynamic Systems and Control Division*, 1994, vol. 55, pp. 295–301.

[12] A. J. Silva, O. A. D. Ramirez, V. P. Vega, and J. P. O. Oliver, “PHANToM OMNI Haptic Device: Kinematic and Manipulability,” in *Proc. of CERMA '09*, 2009, pp. 193–198.

[13] K. Hirota and M. Hirose, “Development of surface display,” in *Proc. of IEEE Virtual Reality '93*, 1993, pp. 256–262.

[14] K. Hirota and M. Hirose, “Simulation and presentation of curved surface in virtual reality environment through surface display,” in *Proc. of IEEE Virtual Reality '95*, 1995, pp. 211–216.

[15] W. A. McNeely, “Robotic graphics: a new approach to force feedback for virtual reality,” in *Proc. of IEEE Virtual Reality '93*, 1993, pp. 336–341.

[16] Y. Yokokohji, J. Kinoshita, and T. Yoshikawa, “Path planning for encountered-type haptic devices that render multiple objects in 3D space,” in *Proc. of IEEE Virtual Reality '01*, 2001, pp. 271–278.

[17] S. Shin, I. Lee, H. Lee, G. Han, K. Hong, S. Yim, J. Lee, Y. Park, B. K. Kang, D. H. Ryoo, D. W. Kim, S. Choi, and W. K. Chung, “Haptic simulation of refrigerator door,” in *Proc. of HAPTICS '12*, 2012, pp. 147–154.

[18] M. Sinclair, M. Pahud, and H. Benko, “TouchMover,” in *Proc. of ITS '13*, 2013, pp. 287–296.

[19] S. Tachi, T. Maeda, R. Hirata, and H. Hoshino, “A Construction Method of Virtual Haptic Space,” in *Proc. of ICAT '94*, 1994.

[20] M. Ortega and S. Coquillart, “Prop-based haptic interaction with co-location and immersion: an automotive application,” in *Proc. of HAVE ’05*, 2005, vol. 2005, pp. 23–28.

[21] K. Yamaguchi, G. Kato, Y. Kuroda, K. Kiyokawa, and H. Takemura, “A Non-grounded and Encountered-type Haptic Display Using a Drone,” in *Proc. of SUI ’16*, 2016, pp. 43–46.

[22] B. Araujo, R. Jota, V. Perumal, J. X. Yao, K. Singh, and D. Wigdor, “Snake Charmer,” in *Proc. of TEI ’16*, 2016, pp. 218–226.

[23] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell, “Passive real-world interface props for neurosurgical visualization,” in *Proc. of CHI ’94*, 1994, vol. 30, pp. 452–458.

[24] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks, “Walking > walking-in-place > flying, in virtual environments,” in *Proc. of SIGGRAPH ’99*, 1999, pp. 359–364.

[25] K.-L. Low, G. Welch, A. Lastra, and H. Fuchs, “Life-sized projector-based dioramas,” in *Proc. of VRST ’01*, 2001, p. 93.

[26] L.-P. Cheng, T. Roumen, H. Rantzsch, S. Köhler, P. Schmidt, R. Kovacs, J. Jasper, J. Kemper, and P. Baudisch, “TurkDeck: Physical Virtual Reality Based on People,” in *Proc. of UIST ’15*, 2015, pp. 417–426.

[27] R. Diankov, “Automated Construction of Robotic Manipulation Programs,” Carnegie Mellon University, 2010.

[28] A. Nicolescu, M. Ivan, C. Avram, and T. Dobrescu, “Volumetric accuracy experimental evaluation and 3D error map generation for a Kawasaki FS 10 E articulated arm industrial robot,” in *Recent Advances in Robotics, Aeronautical and Mechanical Engineering*, Wseas LLC, 2013, pp. 63–68.

[29] C. Trübswetter, “Analyzing and Monitoring Tracking Accuracy of an ART System,” TU München, 2003.

[30] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal, “Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness,” *Int. J. Aviat. Psychol.*, vol. 3, pp. 203–220, 1993.

[31] M. Slater, M. Usoh, and A. Steed, “Taking steps: the influence of a walking technique on presence in virtual reality,” *ACM Trans. Comput. Interact.*, vol. 2, no. 3, pp. 201–219, Sep. 1995.

[32] H. G. Hoffman, “Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments,” in *Proc. of IEEE Virtual Reality '98*, 1998, pp. 59–63.