A Universal Virtual Locomotion System: Supporting Generic Redirected Walking and Dynamic Passive Haptics within Legacy 3D Graphics Applications

Category: Research

Abstract

In this paper we introduce a virtual locomotion system that allows navigation within any large-scale virtual environment (VE) by *real walking*. Based on the results of a user study we have quantified how much users can unknowingly be redirected in order to guide them through an arbitrarily sized VE in which virtual paths differ from the paths tracked in the real working space. Furthermore we introduce the new concept of dynamic passive haptics. This concept allows to map *any number* of virtual objects to real *proxy objects* having similar haptic properties, i. e., size, shape and surface structure, such that the user can sense these virtual objects by touching their real world counterparts. This mapping may be changed dynamically during runtime and need not be one-to-one. Thus dynamic passive haptics provides the user with the illusion of interacting with a desired virtual object by redirecting her/him to the corresponding proxy object. Since the mapping between virtual and proxy objects can be changed dynamically, a small number of proxy objects suffices to represent a much larger number of virtual objects.

We describe the concepts in detail and discuss their parameterization which has been derived from the initially conducted user study. Furthermore we explain technical details regarding the integration into legacy 3D graphics applications, which is based on an interceptor library allowing to trace and modify 3D graphics calls. Thus when the user is tracked s/he is able to explore any 3D scene by natural walking, which we demonstrate by 3D graphics applications from different domains.

Keywords: Virtual Realty, Virtual Locomotion Interface, Generic Redirected Walking, Dynamic Passive Haptics

1 Introduction

Walking is the most basic and intuitive way of moving within the real world. Keeping such an active and dynamic ability to navigate through large-scale virtual environments (VEs) is of great interest for many 3D applications demanding locomotion, such as in urban planning, tourism, 3D entertainment etc. Although these domains are inherently three-dimensional, most applications do not support VR-based user interfaces, least of all *real walking* through the 3D scenes is possible. At most, such applications can display the virtual scene stereoscopically.

In many existing VR systems, the user navigates by means of hand-based input devices in order to specify direction, speed as well as acceleration and deceleration of movement ([23]). Although advanced visual simulation often requires a good sense of locomotion in order to increase the user's immersion into the virtual world, most of these systems do not provide a real sense of walking. An obvious approach to enable users to explore a virtual world by real walking is to transfer the user's movements to corresponding movements in the virtual environment by means of a simple one-to-one mapping. Apparently this technique has the drawback that the users' movements are restricted by the limited range of the tracking sensors and a rather small workspace in the real word. Therefore, concepts for virtual locomotion interfaces are needed that enable walking over large distances in the virtual world while physically remaining within a relatively small space. Many hardware-based approaches have been presented to address this issue [1, 14, 15]. Unfortunately, most of them are very costly and support only walking of a single user, and thus they will probably not get beyond a prototype stage. Fortunately, cognition and perception research suggests that more cost-efficient alternatives exist. It has been known for decades that visual perception usually dominates proprioceptive and vestibular senses. Thus, if the visualization stimulates the user appropriately, it should be possible to guide her/him along a path in the real world which differs from the path the user perceives in the virtual world. For instance, if the user wants to walk straight ahead for a long distance in the virtual world, small rotations of the camera *redirect* her/him to walk unconsciously in circles in the real world. If the induced rotations are small enough, the user receives the impression of being able to walk in the virtual world in any direction without restrictions.

Besides an intuitive navigation, also the perception of a VE is essential in order to achieve a high degree of immersion. While graphics and sound rendering have matured so much that a realistic synthesis of real world scenarios is now possible, generation of haptic stimuli is still a vast area for research. A tremendous effort has been undertaken to support *active* haptic feedback by specialized hardware which generates certain haptic stimuli [5]. However, the necessary instrumentation and wiring restrict the user's interaction space. *Passive* haptic feedback is provided by physical models which represent objects of the virtual world. These models are arranged in the real world in correspondence to their virtual counterparts. By touching a physical model the user gets the impression of interacting with the corresponding virtual object. Since the interaction space is limited only few physical models can be supported, thus restricting the number of virtual objects that can be sensed by the user. Moreover, the presence of physical models in the interaction space prevents the exploration of other parts of the virtual world not represented by the current physical setup. Thus exploration of large scale environments and support of passive haptic feedback seem to be mutually exclusive.

In this paper we address this challenge by combining redirected walking and passive haptics. By visual stimuli the user is guided to a corresponding physical model if he approaches an object in the virtual world; otherwise s/he is guided around the physical models to avoid collisions. With this approach *any number* of virtual objects can be sensed by means of real *proxy objects* having similar haptic properties, i. e., size, shape and surface structure. The mapping from virtual to real objects need not be one-to-one. Since the mapping as well as the visualization of virtual objects can be changed *dynamically* during runtime, a small number of proxy objects suffices to represent a much larger number of virtual objects. By redirecting the user to a preassigned proxy object which hereupon represents a virtual counterpart the user gets the illusion of interacting with a desired virtual object.

Based on the previously mentioned concepts, we propose a new universal virtual locomotion system, which enables users to explore any 3D graphics application by means of real walking; the user's immersion is not disturbed by limited interaction space or physical objects present in the real environment. Furthermore, virtual objects encouraging the user to touch them virtually are mapped to proxy objects objects having similar size and surface structure in order to provide dynamic haptic feedback. Our approach can be used easily in any fully-immersive VR-setup providing user tracking as well as stereoscopic projection, no special hardware is needed in order to support walking or haptics. The technical realization is based on a 3D graphics interceptor library, which allows to trace and/or modify 3D graphics calls of legacy 3D applications. Thus the camera parameters can be changed to support the illusion of real walking, and the sizes of virtual objects can be adapted to the sizes of their physical counterparts. With the presented approach the user can navigate within arbitrarily sized VEs by remaining in a comparably small physical space, where passive haptics is supported. We believe that these features make immersive exploration of VEs more natural and thus ubiquitously available, e. g., when navigating in existing applications such as Google Earth or multiplayer online games.

The remainder of this paper is structured as follows. Section 2 summarizes previous related work. In Section 3 we present a pilot study we have conducted in order to quantify how much users can be *redirected* and how much virtual and proxy objects may differ without the user noticing the discrepancy. Based on the results of this study we have developed the virtual locomotion system supporting *generic redirected walking* as well as *dynamic passive haptics*, which is described in Section 4. Section 5 gives some technical details about how these concepts can be integrated into arbitrary 3D applications and discusses the integration into some example applications. Section 6 concludes the paper and gives an overview about future work.

2 Previous Work

Currently locomotion and perception in virtual worlds are in the focus of many research groups. Early hardware-based technologies such as treadmills or similar devices allow users to walk through VEs ([3]). Most of these approaches do not support omni-directional walking, i. e., the user is not able to change the physical walking direction easily. Hence, various prototypes of interface devices for walking have been developed, including torus-shaped omni-directional treadmills, motion foot pads, and robot tiles ([1, 2, 13, 14, 15]). All these systems have in common that they are very costly and hardly scalable since they support only walking of a single user. For multi-walker scenarios, it is necessary to instrument each user with a separate device. Moreover, most of the described technologies are only applicable to HMD setups, other systems such as CAVEs or curved projection walls are not supported. Although these hardware interface devices represent enormous technological achievements, most likely they will not get beyond a prototype stage in the foreseeable future. Hence the demand for alternative approaches is tremendous. As a solution to this challenge, traveling by exploiting walk-like gestures has been proposed in many different variants, giving the user the impression of walking, for example by walking-in-place, while physically remaining almost at the same position ([12, 22, 20]). However, as a matter of fact real walking is a more presence-enhancing locomotion technique than any other navigation metaphor ([22]).

Redirected walking ([18]) is a promising solution to the problem of limited tracking space and the challenge of providing users with the ability to explore the VE by walking. With redirected walking, the virtual world is imperceptibly rotated around the center of the users head. Thus, when the user explores the potentially infinite VE, s/he unknowingly walks along curved paths within the limited tracking area. This approach is also applied in robotics when controlling a remote robot by walking ([7]). In our system we have extended these redirection concepts by combining motion compression respectively gain, i. e., scaling the real distance a user walks, rotation compression respectively gain, which displays the real turns smaller or larger, and different amounts of curvature, which bend the user's walking direction such that s/he walks on a curve. The phenomenon that users do not recognize small differences between a path in the VE and a path in the real world is based on principles from perceptive psychology. Perception research has identified essential differences between cognition as well as estimation of features in VR in contrast to their counterparts in the virtual world. For example, many researchers have described that distances in virtual worlds are underestimated in comparison to the real world ([10, 11]). Furthermore, it has been discovered that users have significant problems to orient themselves in virtual worlds ([19]). In this context Brooks et al. have investigated how visual perception can dominate over *proprioception* ([4]). They have introduced a shift between the visual representation of the user's arm and its pose in the physical space; up to a certain degree users have not noticed this shift.

Haptic feedback is often supported by exploiting expensive haptic hardware, such as Phantom devices ([5]) or specialized data gloves. Unfortunately no device exists, which can be worn comfortably without any wires and provides at least a sufficient sense of touch. Passive haptic feedback has been used effectively to provide a sense of touch ([9]). The main idea is to replicate counterparts of virtual objects, e. g., virtual walls, virtual tables etc., in the physical space and to arrange them correspondingly. It has been shown that this increases the immersion in the VE significantly ([22]). A first approach to combine passive feedback with redirected walking has been introduced by ([16]): The idea is to rotate the VE such that virtual objects line up with their real-world counterparts. But in the proposed setup walking was restricted to simple straight paths, and the passive feedback was limited to a single symmetric object. Moreover, the passive feedback object had exactly the same size as the virtual ones. Overall, considerable efforts have been undertaken in order to allow a user to walk through a larg-scale VE while presenting continuous passive haptic stimuli, but until now this challenge could not be addressed sufficiently.

3 Pilot Study for Generic Redirected Paths and Dynamic Passive Haptics

A major challenge for our approach is to address two orthogonal VE issues: lack of passive haptic feedback and a limited tracking area. Solutions to both problems are usually mutually exclusive, since their requirements conflict with each other.

As described in Section 1, in order to enhance the user's immersion it is essential that s/he can walk in the entire scene without any constraints enforced by limited tracking space, and that s/he can touch obstacles presented in the virtual world. However, the more proxy objects for passive haptics populate the environment, the more the user's movement in the physical space is restricted. Hence, when using redirection concepts, it has to be ensured that users are guided in such a way that they do not collide with objects of the physical environment or with each other. Consequently users should be able to sense virtual objects when they approach them, i. e., they should be able to sense and touch a proxy object of similar size and surface characteristic.

Hence we have to redirect users on certain paths in order to allow them to sense proxy objects according to the correlation between the real and the virtual environment. In this section we present a pilot study in which we have evaluated how much paths and objects in both worlds can differ, in order to determine limits and thresholds for redirecting users and for employing proxy objects without letting users notice the difference between both worlds.

3.1 Experimental Design

3.1.1 Test Scenario

In our experiments movements of the users are restricted to a $10 \times 7 \times 2.5$ meters tracking range. In the center of this area we placed a square table of size $2 \times 2 \times 1$ meters. The user's path always leads him clockwise or counterclockwise around the table respectively a virtual block displayed in the VE (see Figure 1 and 2).

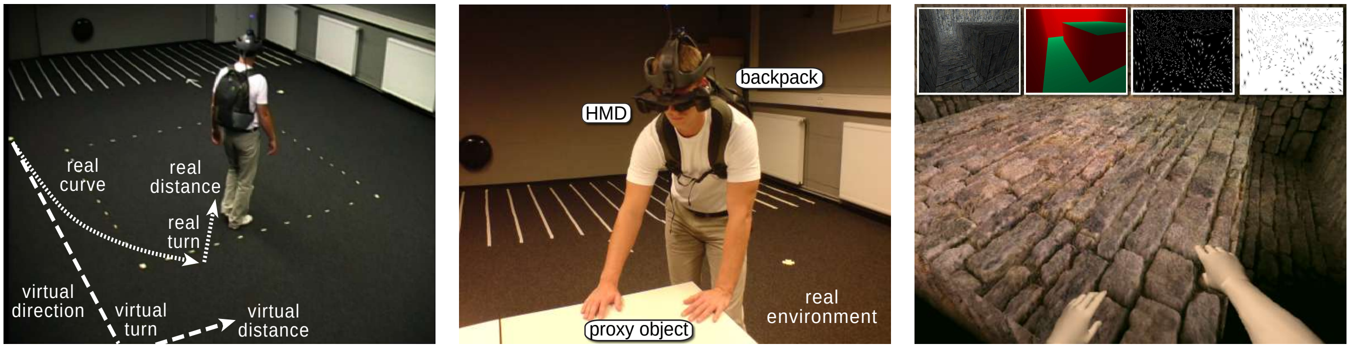

(a) **(b)** **(c)**

Figure 1: Virtual Locomotion Scenario: (a) a user walks through the real environment on a different path with a different length in comparison to the perceived path in the virtual world. (b) the user touches a real proxy object which is considerably smaller than (c) the virtual object seen from the user's perspective (alternative visualizations are overlaid).

As illustrated in Figure 2 the virtual room in which the user walks measures $x_1 \times y_1 \times z_1$ meters, a square block of size $x_2 \times y_2 \times z_2$ meters is displayed in the center. The room and the block can be scaled uniformly. The visual representation of the virtual environment can change continuously between different levels of realism (see Figure 1 (c) (overlays)).

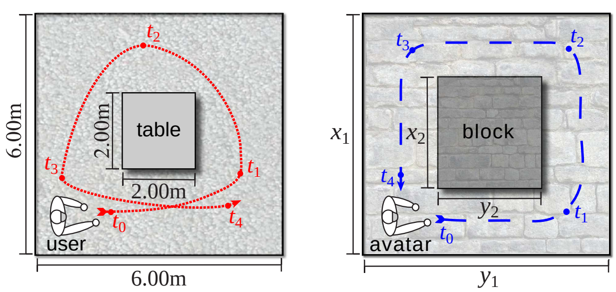

(a) Real setup **(b)** Virtual scene

Figure 2: Illustration of a user's path during the experiment showing (a) path through the real setup and (b) virtual path through the VE and positions at different points in time $t_0, ..., t_4$.

3.1.2 Participants

A total of 8 (7 male, 1 female) subjects participated in the study. Three of them had experience with walking in VR environments using an HMD setup. We arranged them into three groups: 3 "expert users" (EU Group) who knew about the objectives and the procedure before the study, and 3 "aware users" (AU Group) who knew that we would manipulate them, but had no knowledge about how the manipulation would be performed. Subjects of both groups were asked to report if and how they realized any discrepancy between actions performed in the real world and the corresponding transfer in the virtual world. 2 "naive users" (NU Group) had no knowledge about the goals of the experiment and thought they had to report any kind of tracking problems.

3.2 Tasks

We have performed some preliminary interviews and tests with the participants in which we revealed their spatial cognition and body awareness by means of distance perception and orientation tests: Subjects had to walk certain distances blindfolded, which they had viewed before, and they had to rotate by particular angles and rotate back blindfolded. One objective of the study was to see if and how body awareness may be affected by using our virtual locomotion system over time. Therefore, we have performed the same tests before, during and after the experiments.

As illustrated in Figure 2, during the main experiment the participants had to walk either clockwise or counter-clockwise around the virtual block in the center of the virtual room. In order to support generic redirected walking concepts as well as dynamic passive haptic strategies, we have modulated the real and the virtual environment by means of the following independent variables:

3.2.1 Independent Variables

- **Rotation compression resp. gain factor** s_{rot} describes the compression respectively gain of a user's head rotations, i. e., when the user rotates the head by α degrees, the virtual camera is rotated by $s_{rot} \cdot \alpha$ degrees.
- **Amount of curvature** s_{cur} denotes the bending of a real path, i. e., when the user moves, the camera rotates continuously such that the user walks along a curve in order to stay on a straight path in the virtual world. The curve is determined by a segment of a circle with radius r, where $s_{cur} := \frac{1}{r}$. The resulting curve is considered for a straight distance of 1 meter. Hence in the case that no curvature is applied $r = \infty$ and $s_{cur} = 0$, whereas if the curvature causes that the user has rotated by $90°$ clockwise after 1 meter, the user has covered a quarter circle and therefore $s_{cur} = 1$.
- **Motion compression resp. gain factor** s_{mot} denotes the scaling of translational movements, i. e., 1 unit of physical motion is mapped to s_{mot} units of camera movement in the same direction.

In order to quantify certain parameters for dynamic passive haptics we have applied the following independent variables to the virtual objects and the virtual environment:

- **Object compression resp. gain factor** s_{obj} denotes a uniform scaling transformation applied to virtual objects or the entire VE. It would also be possible to use different scaling factors in order to perform non-uniform scaling transformations.
- **Level-of-detail (LoD)** represents the degree of realism for the virtual scene, i. e., the scene can be altered from having relief-mapping based on high-quality textures, to flat, textured walls, and flat-shaded walls (see Figure 1 (c) (overlays)).
- **Optical flow** gives indications about the motion of objects within a visual representation. In order to quantify how optical flow influences our concepts we apply textures to the

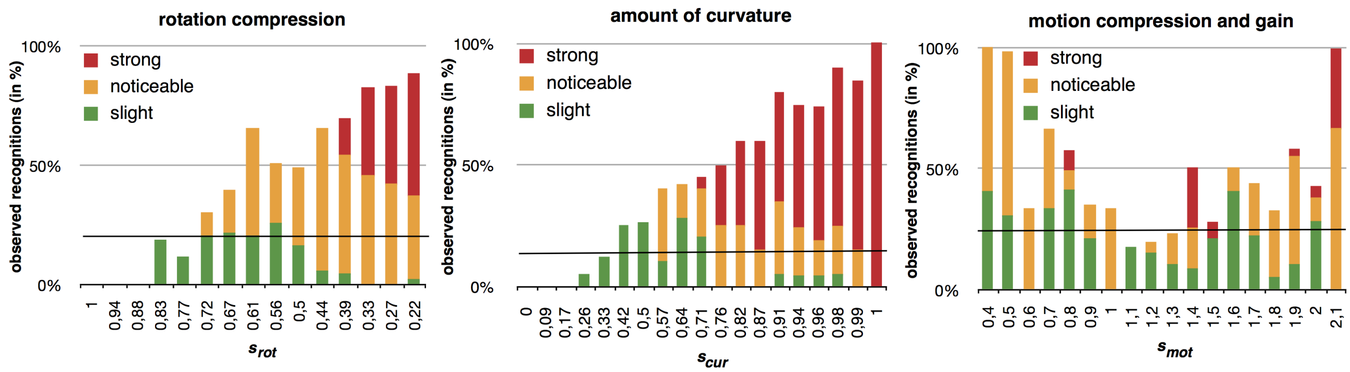

Figure 3: Evaluation of the generic redirected walking concepts for (a) rotation compression factors s_{rot}, (b) amount of curvature s_{cur} and (c) motion compression resp. gain factors s_{mot}. The bars indicate how much users have perceived the manipulated walks. The EU, AU and NU groups are combined in the diagrams. The horizontal lines indicate the thresholds as described in Section 4.1.

surfaces of the virtual room which only contain small circles. The number, size and lifetime of the circles can be changes: many circles which do not disappear provide rich optical flow, whereas less circles with short lifetime provide only low optical flow (see Figure 1 (c) (overlays)). of the virtual room, where

We use the above variables in our generic redirected walking concepts and evaluate how they can be modified without the user noticing any changes. The sequence and also the amount of change has been chosen differently for each subject in order to reduce any falsification caused by learning effects. After a training period we used random series starting with different amounts of discrepancy between real and virtual world. We slightly increased the difference each 5 to 20 seconds randomly until subjects report visual-proprioceptive discrepancy–this means that the perceptual threshold has been reached. Afterwards we have performed further tests in order to verify the subjective value.

All variables have been logged and subjective comments have been recorded in order to reveal in how far subjects perceive a difference between the virtual and the real world. The amount of difference has been evaluated on a four-point Likert scale where (0) means no distortion, (1) means a slight, (2) a noticeable and (3) a strong perception of the discrepancy.

3.3 Setup

The experiments have been performed in a laboratory environment. Technical aspects about this setup are described in this section.

Visualization Hardware

In the experiments we used an Intel computer (host) with dual-core processors, 4 GB of main memory and an *n*Vidia GeForce 8800 for system control and logging purposes. The participants were equipped with a HMD backpack consisting of a laptop PC (slave) with a GeForce 7700 Go graphics card and battery power for at least 45 minutes (see Figure 1). The scene was rendered using DirectX and our own software with which the system maintained a frame rate of 30 frames per second. The VE was displayed on two different head-mounted display (HMD) setups: (1) a ProView SR80 HMD with a resolution of 1240×1024 and a large diagonal optical field of view (FoV) of 80°, and (2) eMagin Z800 HMD having a resolution of 800×600 with a smaller diagonal FoV of 45°. During the experiment the room was darkened in order to reduce the user's perception of the real world. In addition we attached a completely opaque towel to the HMDs, such that the users were not able to perceive any information from the real environment.

Tracking System and Communication

We used the WorldViz Precise Position Tracker, an active optical tracking system which provides sub-millimeter precision and sub-centimeter accuracy. With our setup the position of up to eight active infrared markers can be tracked within an area of approximately $10 \times 7 \times 2.5$ meters. The update rate of this tracking system is about 60 Hz, providing real-time positional data of the active markers. The positions of the markers are sent via wireless LAN to the laptop. For the evaluation we attached a marker to the HMD, but we have also tracked hands and feet of the user. Since the HMD provides no orientation data, we use an Intersense InertiaCube2 orientation tracker that provides full 360° tracking range along each axis in space and achieves an update rate of 180 Hz. The InertiaCube2 is attached to the HMD and connected to the laptop in the backpack of the user.

All computers including the laptop on the back of the user are equipped with wireless LAN adapters. We use a dual-sided communication: data from the InertiaCube2 and the tracking system is sent to the host computer where the observer logs all streams and oversees the experiment. In order to apply generic redirected walking and dynamic passive haptic concepts, i. e., altering the variables explained in Section 3.2.1, the experimenter can send corresponding control inputs to the laptop. The entire weight of the backpack is about 8 kg, which is on the one hand quite heavy, while on the other hand no wires disturb the immersion. Hence, no assistant must walk beside the user in order to watch over the wires. Sensing the wires would give the participant a cue to orient physically, an issue we had to avoid in our study. Communication between the user and the experimenter has been performed via a dual head set system only. In addition acoustic feedback within the virtual scene can be applied via the head set, but was only used for ambient noise in the experiment.

3.4 Analyses of the Results

Rotation Compression and Gain

We have tested 147 different rotation compression and gain factors. Figure 3 (a) shows the corresponding factors applied to a 90° rotation. The bars show how many users have perceived the manipulation at different intensities. It points out that when we scaled a 90° rotation down to 80°, which corresponds to a compression factor $s_{rot} = 0.88$, none of the participants has realized the compression. Even with a compression factor $s_{rot} = 0.77$ subjects have rarely (11%) recognized the discrepancy between the physical rotations and the corresponding camera rotations. If this factor is applied users are forced to physically rotate nearly 30° more when they perform a 90° virtual rotation without recognizing the discrepancy. Even when a rotation compression factor $s_{rot} = 0.5$ is applied only 50% of the rotations were underestimated. The subjects have adapted to rotation compression respectively gain factors very fast and they sense them as correct rotations. We have performed a *virtual blindfolded turn test*. The subjects were asked to turn 135° in the virtual environment, where a rotation compression factor of $s_{rot} = 0.7$ has been applied, i. e., subjects had to turn physically about 190° in order to achieve the required virtual rotation. Afterwards they were asked to turn back to the initial position. When only a black image was displayed the participants rotated on average 148° back. This is a clear indication that the users sense the compressed rotations close to a real 135° rotation and hence adapt well to the applied rotation compression factor.

Amount of Curvature

Overall we have tested 165 distances where we applied different amounts of curvature as illustrated in Figure 3 (b). When s_{cur} satisfies $0 < s_{cur} < 0.17$ the curvature has not been recognized by the subjects. Hence after 3 meters we were able to redirect subjects up to 15° left or right while they were walking on a segment of a cycle with radius of approximately 6 meters. As long as $s_{cur} < 0.33$, only during 12% of the walked distances subjects perceived the difference between real and virtual paths. Furthermore, we observed the more slowly participants walk the less they observe that they walk on a curve instead of straight ahead. When they increased the speed they began to careen and realized the bending of the walked distance.

Motion Compression

We have tested 216 distances to which different motion compression respectively gain factors have been applied (see Figure 3 (c)). As mentioned in Section 1, users tend to underestimate distances in VR environments. Consequently subjects underestimate the walk speed when a motion gain factor below 1.2 is applied, i. e., the gaits are mapped one-to-one or are even slightly accelerated. On the opposite, when a motion compression factor satisfies $s_{mot} > 1.6$ subjects recognize the scaled movements immediately. Between these thresholds some subjects overestimated the walk speed whereas others underestimated it. However, most subjects stated the usage of such a factor only as slight or noticeable. In particular the more users tend to careen the less they realize the application of a motion gain or compression factor. This may be due to the fact that when they move the head sidewards the motion compression factor also applies to corresponding motion parallax phenomena. This could lead to the issue that users adapt more easily to the scaled motions thereupon.

For this experiment we have also performed a *virtual blindfolded walk test*. The subjects were asked to walk 3 meters in the VE, where motion compression respectively gain factors between 0.7 and 1.4 have been applied. Afterwards they were asked to turn, review the walked distance and walk back to the initial position, while only a blank screen has been displayed again. Without any factors applied to the motion, users walked back on average 2.7 meters. For each motion compression factor the participants have walked too short, which is a well-known effect because of the described underestimation of distances but also due to safety reasons; since after each step participants are less oriented and thus tend to walk shorter so that they do not collide with any objects. However, on average they have walked 2.2 meters for a motion gain factor $s_{mot} = 1.4$, 2.5 meters for $s_{mot} = 1.3$, 2.6 meters for $s_{mot} = 1.2$, and 2.7 meters for $s_{mot} = 0.8$ as well as for $s_{mot} = 0.9$. When the motion compression factor satisfies $0.8 < s_{mot} < 1.2$, i. e., only small compression has been applied, user have walked 2.7 meters back on average. Hence for these motion compression respectively gain factors user have adapted to the VE perfectly.

Dynamic Passive Haptics

To map virtual to real objects, we considered only the bounding boxes of the proxy and the virtual objects. At the end of the series we have tested the described approach of dynamic passive haptics. Therefore, when a subject arrived at corner of the virtual room we have applied an object scaling factor s_{obj} to the VE, which has also been used as motion compression respectively gain factor, i. e., $s_{mot} = s_{obj}$. For instance, when the scaling factor satisfies $s_{obj} = 2.0$, the scene is uniformly scaled by 2.0, but also the user's translational movements have been scaled with this factor. Hence subjects needed required the same time to walk around the virtual block, although its visual representation has doubled size. The evaluation shows that an object compression respectively gain factor between 0.7 and 1.4 could not been recognized when a corresponding motion factor is applied. When $1.4 < s_{obj} < 1.6$ users have slightly recognized the scaling, whereas scaling factors above 1.6 have been stated as clearly noticeable. For orientation mismatches between virtual and proxy objects the same results as for redirected walking concepts seem to be applicable, but this has to be examined in further studies.

Further Observations

Although, some participants had walked over 45 minutes using the HMD setup, cyber sickness has occurred to two participants only. In a post-test two days later these two participants had cyber sickness again, although no generic redirected walking concepts have been applied. Moreover, users feel uncomfortable due to the weight of the setup and the tightness under the HMD. However, they definitely prefer to wear a backpack than to take care about wires during walking. The effect of different visual appearances had no significant influence on the evaluation, although in poor optical flow environments subjects tend to realize redirected walking less than in contrast to environments which are populated with rich optical flow. In addition there was no significant difference in the evaluation between the EU group, AU group or the NU group; even when we tested the experts, they hardly recognized the application of generic redirected walking or dynamic passive haptics strategies.

4 Universal Virtual Locomotion Interface

In this section we describe the parameterization of our universal virtual locomotion system with respect to the results obtained from the pilot study described in Section 3.

4.1 Virtual Locomotion Strategies

In order to provide the user with passive haptic feedback it is essential to know where or what a user will approach. For instance, when the user moves towards a virtual building, in particular a wall, s/he should be redirected to a corresponding proxy object in the physical world that is associated to the virtual wall and provides passive feedback, e. g., a wall in the physical space. To redirect a user the virtual camera has to be manipulated such that the user will unknowingly follow a certain path. Based on the results from Section 3 we formulate some guidelines so that it is possible to redirect sufficiently such that the user neither perceives the redirected walking nor cyber sickness occur:

1. The *rotation compression resp. gain factor* is constrained to $0.75 < s_{rot} < 1.35$.
2. The *amount of curvature* is constrained to $s_{cur} < 0.33$.
3. *Motion compression resp. gain factor* must satisfy $0.85 < s_{obj} < 1.45$.
4. The application of the *object compression resp. gain factor* is restricted by $0.62 < s_{mot} < 1.45$ in order to coincide proxy objects with virtual objects.

As evaluated in a test of significance when these requirements are fulfilled less than 15% of all walks are perceived as manipulated

(p-value < 0.05). Besides the application of passive haptic feedback, another essential objective is to keep the user in the tracking area and to prevent collisions with physical objects that are not in the view of the user in the virtual world. In our approach the main path on which users walk is a circle in a squared tracking area and an ellipse in a rectangular tracking area. When the user potentially leaves the tracking area, we determine the angle of intersection between the user's path and the boundary of the tracking area. Then corresponding camera modifications are performed such that the user is redirected in or on the circle with respect to the guidelines 1–4.

4.2 Dynamic Path-Prediction

The association between proxy objects and their virtual counterparts has to be done manually in our approach. Therefore, we use an XML-based description to assign virtual objects to proxy objects and to define their orientation. Currently the geometries of virtual and proxy objects are only defined by means of 3D bounding boxes.

In order to redirect the user to a proxy object the virtual object to which the user approaches has to be predicted. Therefore, we have developed a simple path prediction approach, where the predication depends on the combination of the current view and walk direction. While the walk direction determines the predicted path, the view direction is used for verification. When the projections of both vectors to the ground differ more than 45°, no reliable prediction can be made and the user is only redirected when a collision in the physical space may occur or when the user might leave the tracking area. When the angle between both vectors is sufficiently small ($< 45°$) the walk direction defines the predicted path. A corresponding line segment is tested for intersections with virtual objects that are defined in the XML-based description. Therefore a ray shooting similar to the approaches referenced in [17] is applied. If an intersection is determined with the path and an edge of the bounding box of a virtual object respectively proxy object, a corresponding path of continuous curve segments is composed such that the user is redirected accordingly; each curve segment follows the guidelines from Section 4.1. Indeed situations may occur where the curve segments cannot follow the requirements in order to redirect the user. In this case the parameterization of the curve segments needs to be altered.

Usually the bounding boxes between the virtual object and an assigned proxy object are different. As a result from the pilot study 3 we know that the size of virtual objects can be modified with respect to guideline 4 defined in Section 4.1 in order to let users perceive virtual and proxy objects congruently. Since the XML-based description contains the initial orientation between virtual and proxy objects, it is even possible to redirect a user to the desired proxy object from the correct side in order to enable a consistent passive feedback. Since we have only experimented with scenarios where path prediction for virtual and proxy objects are constrained to bounding boxes calculation of intersections is simple and can be updated for each frame. When the proxy objects are tracked, a change of their pose is updated in the XML-based description. Hence also dynamic scenarios, where the virtual but also the physical environment may change, are considered in our approach.

4.3 Security Consideration

Although one objective of the dynamic real-time path-prediction is to prevent the user from leaving the tracking area or colliding with physical objects, it might occur that s/he left the provided path. Hence within 100 cm of the boundary of the tracking area or of an object with which a collision might occur we perform a rigorous 90° rotation of the virtual scene with respect to the angle between the predicted path and the edge of the object's bounding box in order to force the user to walk around obstacles respectively to stay in the tracking space. Indeed, this drastic change of orientation is observed by the user. Furthermore, in the case that the user potentially might collide with an obstacle or leaves the tracking area unintentionally, we add multi-sensory feedback as warning cues. As visual feedback the VE displayed on the HMD is blended with red color the closer the user gets to the obstacle respectively boundary of the tracking area. In addition an acoustic warning signal is send to the user in this case.

4.4 Virtual Avatars

Several studies confirmed that the user association with a virtual body is essential to enhance presence ([8]). Therefore, we use a detailed polygonal model of a human body, consisting of two legs (each over 5,000 polygons) and two arms (each over 3,000 polygons) which can be associated to tracked user's extremities.

As mentioned in Sections 4.1 and 4.2 we apply the path prediction and the corresponding redirection concepts dynamically. Hence, moving virtual objects as well as moving proxy objects can be incorporated as long as they are included in the scene description. Thus, also multiple users can be associated to virtual avatars as long as they are tracked. Hence users can see virtual avatars of other users and interaction can be performed between users respectively avatars. Therefore, for example, a *virtual handshake* with passive haptics given by the user's real hand can be performed.

5 Integration into Legacy Application

The described concepts for generic redirected walking and dynamic passive haptics can be applied in any legacy 3D graphics application. There is a rising demand for applications dealing with complex 3D data sets, which would benefit from the possibility to navigate naturally through the data. Most of these applications do not provide the required interfaces; even other VR-based technologies such as stereoscopic displays or tracking interfaces are usually not supported. The objective of this work is to provide a universal virtual locomotion system that is not restricted to a test bed application, but which is usable for arbitrary graphics applications. Hence, we have developed a 3D interceptor library which allows us to trace and modify 3D content from any graphics application based on OpenGL or DirectX.

5.1 General 3D Interceptor Library

The central idea of the general 3D interceptor library is to intercept graphical function calls that define the visual appearance of the virtual scene. This procedure is well-known and applied widely ([6]). But when tracing graphics calls only the content, which is currently visible to the camera, can be captured. In common 3D graphics applications not all geometry is sent to the graphics card in order to increase performance, for example, view frustum culling, hidden surface removal and level of detail concepts are applied at the application stage. Hence when tracing function calls and performing a camera rotation before executing them, the virtual scene will disappear since no data is available for the current view frustum. For this reason this approach is only usable where small modifications of the camera are intended, for example for stereoscopic display, i. e., only small translations respectively shearing transformations are applied. In order to be able to manipulate the camera arbitrarily while all features of the scene camera, such as culling, LoD, dynamics etc., are maintained, we have extended the framework described in [21]. In the original version the framework allows to trace and modify graphics function calls, but also has the described limitations. Therefore we have developed and integrated the following extensions.

We assume that the user's head is tracked at a certain position with a given orientation in terms of a tracking coordinate system. This tracked data has to be applied to the virtual scene camera in order to explore the VE by head movements only.

5.1.1 Calculating Virtual Camera Parameters

As remarked, when integrating an additional transformation before executing the captured 3D graphics function calls, the entire application logic is not aware that the camera position has changed. Usually when using the navigation interface of an arbitrary application, for instance via mouse or keyboard, position and orientation changes of the virtual camera are limited between two successive frames. Thus, the potential positions reachable from the current position can be arranged in a grid-like manner. This makes it easy to identify the possible successive positions for the next frame. The same applies for a change of orientation. We exploit this restriction of the navigation interface. The main idea is to arrange the scene camera to the position and orientation which is closest to the tracked position/orientation. Hence, for each application we provide a mapping script that enables to emulate corresponding mouse or keyboard interactions which arrange the scene camera as close as possible to the tracked position and orientation. The position which is located most closely to the tracked position is determined by euclidean distance metrics, whereas the closest orientation is determined by that orientation, whose view frustum is most congruent with the view frustum reachable by the application interface.

After the desired grid position and orientation is determined, the virtual camera is moved to that position and rotated accordingly using the mapping script, but rendering is deferred. Now the essential parts of the virtual scene are captured and minor changes can be applied, i. e., the transformation between the calculated scene's camera's position/orientation and the tracked position/orientation. Since this transformation changes the position/orientation of the scene camera only marginally, almost the entire 3D content will be in the current view frustum, and will be displayed properly.

5.1.2 Non-tracked Virtual Avatar

With the 3D interceptor library we are not restricted to transformations of traced graphical function calls, but also further rendering calls can be added. Hence we add the virtual avatar described in Section 4.4 to the rendering process (see Figure 1 (c)). As stated in [8] even when the position and orientation of the extremities are not tracked it is beneficial to add static or moving (if the user moves) avatar elements to the scene in order to increase the user's immersion. Hence when the extremities are not tracked we orient the virtual torso orthogonally to the walk direction of the user. During the experiment none of the participants has recognized any accuracy errors caused by this approach. Furthermore, animated hand and leg movements are initiated when movement of the user's head is tracked. This is beneficial in particular when multiple users collaborate and users can see the avatars of other users walking.

When using our universal virtual locomotion system concepts users can explore large-scale VEs, but as in the real world, walking over large distances might be exhausting and moreover users want to exploit benefits provided by VR. Since our library enables us to manipulate any 3D graphics application based on OpenGL or DirectX, we can also apply other generic navigation concepts, which are usable across different applications. We have therefore integrated certain navigation concepts which allow to efficiently explore also large scale VEs, in particular, by virtual flying, gliding or using a rocket belt metaphor (see Figure 4 (a)), e. g., by using both-handed devices. Although the navigation interface differs from application to application, these concepts can be applied across different applications and provide a universal navigation interface.

5.2 Example 3D Graphics Applications

In this section we describe two example applications from different domains, i. e., a geospatial visualization and a collaborative 3D entertainment environment, in which we have integrated the concepts presented in this paper.

5.2.1 Google Earth

The widely used geographic visualization application Google Earth combines a search engine with satellite imagery, maps, terrain and 3D city models in order to collect geographic information. Several city and urban governments have supplied their data, or at least use Google Earth as visualization toolkit for their data. By now some city models include 3D buildings modeled up to a LoD 4, which corresponds to textured models of interiors. The user is able to virtually navigate world-wide and to explore specific features and landmarks within urban environments.

We have integrated our approaches in Google Earth which enables users to explore a virtual model in a natural way. In Figure 4 (b) the user travels through an interior model of a train station by real walking. When using our generic redirection concepts the user is able to walk through 3D city models without any restrictions in the physical space. In the future we will apply dynamic passive haptics in Google Earth, which means to provide proxy objects and to generate a corresponding XML-based description that we extract from the Google KML file format. Hence the user will be able to sense virtual walls displayed in Google Earth by means of proxy objects.

5.2.2 Second Life

The popularity of Second Life has inspired us to integrate the concepts into this virtual world which is entirely built and owned by its residents. When using our concepts a user can walk through the virtual world that is displayed in Second Life in the same way as walking through Google Earth. Furthermore, we can also apply the concepts described in Section 5.1.2 in order to allow multiple users to explore the VE simultaneously. Interaction concepts which might result from such a multi-user scenario have not been examined in the scope of this paper.

Besides executing additional rendering calls the 3D interceptor library supports even suppressing of certain graphics function calls. The Second Life application requires us to prevent the display of the user's avatar which is usually visualized in front of the virtual scene camera permanently. Instead we apply the user's virtual avatar as described in Section 5.1.2. Figure 4 (c) shows the view of a user in a Second Life scene. The image is captured from the user's perspective. The avatars seen from this view are avatars of other users which are not part of our system.

6 Discussion and Future Work

In this paper we have introduced a software-based universal virtual locomotion system, which has been designed according to guidelines which we have identified from the pilot study. With the system it is possible to explore VEs displayed in any 3D graphics application by real walking. Hence, the challenge of natural traveling in limited tracking space has been resolved sufficiently by redirected walking approaches, where dynamic passive feedback provides the user with the possibility to touch virtual objects respectively associated proxy objects.

Although participants of the user study as well as further users have stated the benefits and usability of universal locomotion system for different application domains, we have identified some issues which could be improved. Currently the XML-based description of the virtual and the real world defining associations for dynamic passive objects has to be done manually. Since many HMDs have a camera attached, computer vision algorithms may be exploited in order get some basic information about the real environment, whereas the depth image of the VE may be analyzed to retrieve information about the virtual world. Thus the assignment between identified virtual and proxy objects could be done automatically if a corresponding ontology exists.

Another drawback is the grid-like estimation of possible camera positions and orientations. The therefore used mapping script which supports full control over virtual cameras could be omitted if each application would provide an interface to set their scene camera. Therefore, we develop a generic interface, which accesses the virtual scene camera if supported by the 3D application. Additionally, the path-prediction algorithms are currently in a basic stage.

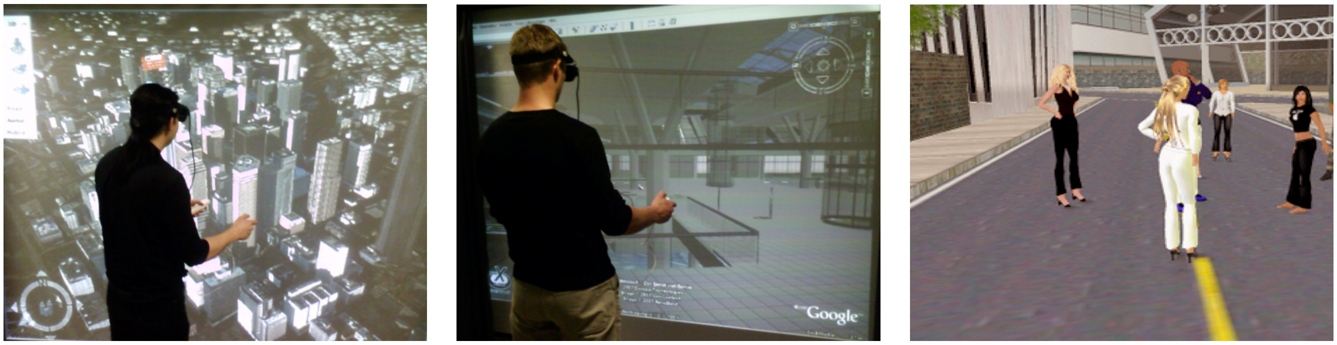

(a) (b) (c)

Figure 4: Example images illustrating the integration of the concepts into exemplary applications: (a) a user flying above a city model using the rocket belt metaphor and (b) a user walking through a model of train station displayed in Google Earth, as well as (c) the Second Life world from the user's perspective.

With increasing complexity of scenes the calculated paths tend to lack the guidelines evolved in this paper. However, improved approaches for the prediction exist ([7]), but have to be integrated into our system yet.

As mentioned above, the presented system enables us to develop and integrate further navigation and interaction concepts: In our system virtual objects can be scaled by a factor between 0.7 and 1.4 without the user recognizing the difference. We can further exploit this lack of human perception by modifying not only virtual objects, but also the extremities of the user's avatar, in particular the position of the arms and hands when sensing a virtual object similar to [4]. To further improve the redirection of the user even the visualization of the legs can be adapted. For instance, when the user should be directed on a curve bend to the left, we visualize a moving leg with a small shift to the right. In order to balance the user moves the leg to the left accordingly.

The application of the concepts presented in this work raise further interesting questions, in particular multi-user scenarios in which several users interact simultaneously provide great potential. In these cases even virtual avatars of certain users might be associated to other users; this procedure would be similar to the approach of dynamic proxy objects.

Overall the introduced universal locomotion system provides a promising approach which definitely enhances the user's presence in virtual worlds. Many existing applications from different domains could potentially benefit from the possibility to naturally explore virtual environments in an immersive way.

References

[1] L. Bouguila and M. Sato. Virtual Locomotion System for Large-Scale Virtual Environment. In *Proceedings of Virtual Reality*, pages 291–292. IEEE, 2002.

[2] L. Bouguila, M. Sato, S. Hasegawa, H. Naoki, N. Matsumoto, A. Toyama, J. Ezzine, and D. Maghrebi. A New Step-in-Place Locomotion Interface for Virtual Environment with Large Display System. In *Proceedings of SIGGRAPH*, pages 63–63. ACM, 2002.

[3] D. Bowman, E. Kruijff, J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley, 2004.

[4] E. Burns, S. Razzaque, A. T. Panter, M. Whitton, M. McCallus, and F. Brooks. The Hand is Slower than the Eye: A Quantitative Exploration of Visual Dominance over Proprioception. In *Proceedings of Virtual Reality*, pages 3–10. IEEE, 2005.

[5] M. Calis. Haptics. Technical report, Heriot-Watt University Internal Report, 2005.

[6] Eyebeam. OGLE: The OpenGLExtractor.

[7] H. Groenda, F. Nowak, P. Rößler, and U. D. Hanebeck. Telepresence Techniques for Controlling Avatar Motion in First Person Games. In *Intelligent Technologies for Interactive Entertainment (INTETAIN 2005)*, pages 44–53, 2005.

[8] C. Heeter. Being There: The Subjective Experience of Presence. *Presence: Teleoperators and Virtual Environments*, 1(2):262–271, 1992.

[9] B. Insko, M. Meehan, M. Whitton, and F. Brooks. Passive Haptics Significantly Enhances Virtual Environments. In *Proceedings of 4th Annual Presence Workshop*, 2001.

[10] V. Interrante, L. Anderson, and B. Ries. Distance Perception in Immersive Virtual Environments, Revisited. In *Proceedings of Virtual Reality*, pages 3–10. IEEE, 2006.

[11] V. Interrante, B. Ries, J. Lindquist, and L. Anderson. Elucidating the Factors that can Facilitate Veridical Spatial Perception in Immersive Virtual Environments. In *Proceedings of Virtual Reality*. IEEE, 2007.

[12] V. Interrante, B. Riesand, and L. Anderson. Seven League Boots: A New Metaphor for Augmented Locomotion through Moderately Large Scale Immersive Virtual Environments. In *Proceedings of Symposium on 3D User Interfaces*, pages 167–170. IEEE, 2007.

[13] H. Iwata. The Trous Treadmill: Realizing Locomotion in VEs. *IEEE Computer Graphics and Applications*, 9(6):30–35, 1999.

[14] H. Iwata, Y. Hiroaki, and H. Tomioka. Powered Shoes. *SIGGRAPH 2006 Emerging Technologies*, (28), 2006.

[15] H. Iwata, H. Yano, H. Fukushima, and H. Noma. CirculaFloor. *IEEE Computer Graphics and Applications*, 25(1):64–67, 2005.

[16] L. Kohli, E. Burns, D. Miller, and H. Fuchs. Combining Passive Haptics with Redirected Walking. In *Proceedings of the International Conference on Augmented Tele-Existence*, volume 157, pages 253 – 254. ACM, 2005.

[17] M. Pellegrini. Ray Shooting and Lines in Space. *Handbook of discrete and computational geometry*, pages 599–614, 1997.

[18] S. Razzaque, Z. Kohn, and M. Whitton. Redirected Walking. In *Proceedings of Eurographics*, pages 289–294. ACM, 2001.

[19] B. Riecke and J. Wiener. Can People not Tell Left from Right in VR? Point-to-Origin Studies Revealed Qualitative Errors in Visual Path Integration. In *Proceedings of Virtual Reality*, pages 3–10. IEEE, 2007.

[20] M. C. Schwaiger, T. Thümmel, and H. Ulbrich. Cyberwalk: Implementation of a Ball Bearing Platform for Humans. In *Proceedings of 12th Internation on Human-Computer Interaction*, pages 926–935, 2007.

[21] F. Steinicke, T. Ropinski, G. Bruder, and K. Hinrichs. Interscopic User Interface Concepts for Fish Tank Virtual Reality Systems. In *Proceedings of the Virtual Reality*, pages 27–34, 2007.

[22] M. Usoh, K. Arthur, M. Whitton, R. Bastos, A. Steed, M. Slater, and F. Brooks. Walking > Walking-in-Place > Flying, in Virtual Environment. In *International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 359 – 364. ACM, 1999.

[23] M. Whitton, J. Cohn, P. Feasel, S. Zimmons, S. Razzaque, B. Poulton, and B. M. und F. Brooks. Comparing VE Locomotion Interfaces. In *Proceedings of Virtual Reality*, pages 123–130. IEEE, 2005.