FP7-ICT-2007-C
Objective ICT-2007.8.0
FET-Open 222107

NIW

Natural Interactive Walking

Deliverable 6.1

PRESENCE STUDIES IN MULTIMODAL AUGMENTED FLOORS

Stefania Serafin, Rolf Nordahl, Luca Turchet, Niels C. Nilsson, AAU
Anatole Lécuyer, Gabriel Cirio, Maud Marchal, INRIA
Jeremy R. Cooperstock, McGill

Sep. 23, 2011
v. 1.0

Classification: PU

Contents

0.1 Introduction

The goal of this deliverable is to provide an overview of the studies performed during the third year of the NIW project regarding integration of technologies and evaluation of presence in multimodal virtual environments.

In the first chapter a simulation of multimodal rendering of fluids is presented. In the second chapter several innovative interfaces for navigation in virtual environments are presented. In the third chapter it is introduced a multimodal architecture to simulate walking in virtual environments, and an evaluation of the ability of this architecture to induce a sense of presence in VR is performed.

Overall, the different multimodal interfaces presented in this chapter provide a rich overview of the activities performed by the different partners of the project, often through collaborations among institutions.

1 Vibrotactile Rendering of Fluids

1.1 Introduction

Many common materials with which we interact on a daily basis can be simulated and displayed through the vibrotactile modality. Simulation may employ physically based models of the material to provide a high degree of realism. This approach has been applied to solids such as wood and metal [68] and aggregates such as gravel and snow [115]. Display can be achieved by the wide availability of off-the-shelf and easily built vibrotactile hardware (actuated floors [114], shoes [72], and hand-held transducers) supporting bodily interaction with virtual materials. The addition of such vibrotactile rendering to the complementary visual and acoustic modalities can enhance the quality of the display, increasing the degree of realism, and in turn, improve user immersion. In this manner, compelling virtual reality (VR) scenarios involving physically based virtual materials have been demonstrated using hand-based and foot-based interaction with visual and vibrotactile feedback [115, 117].

However, some materials, such as water and other fluids, have been largely ignored in this context. For VR simulations of real-world environments, the inability to include interaction with fluids is a significant limitation. The work described here represents an initial effort to remedy this, motivated by our interest in supporting multimodal VR simulations such as walking through puddles or splashing on the beach, as shown in Figure 1.1. Further applications include improved training involving fluids, such as medical and phobia simulators, and enhanced user experience in entertainment, such as when interacting with water in immersive virtual worlds.

To this end, the main contribution of this part is the introduction of the first physically based vibrotactile fluid rendering model for solid-fluid interaction. Similar to other rendering approaches for virtual materials [115, 117, 69], we leverage the fact that vibrotactile and acoustic phenomena share a common physical source. Hence, we base the design of our vibrotactile model on prior knowledge of fluid sound rendering. Since fluid sound is generated mainly through bubble and air cavity resonance, we enhanced our fluid simulator with real-time bubble creation and solid-fluid impact mechanisms, and can synthesize vibrotactile feedback from interaction and simulation events. Using this approach, we are exploring the use of bubble-based vibrations to convey fluid interaction sensations to users. We render the feedback for hand-based and, in a more innovative way, for foot-based interaction, engendering a rich perceptual experience of feeling the sensations of water.

Figure 1.1: A user feeling the water under his feet as a wave washes up on the beach

1.2 Related Work

Recent work has focused on the vibrotactile feedback of virtual materials. All the physically based approaches rely on the same principle: using a common model to generate acoustic and vibrotactile feedback, due to the common physical source of both phenomena in an interaction context. Visell et al. [115] use an aggregate model to simulate snow, gravel and sand, which is rendered through tiles equipped with vibrotactile transducers. Using the same devices, a fracture model is used to simulate the cracking ice of a frozen pond [117]. Papetti et al. [72] use various acoustic models from the Sound Design Toolkit [1] to generate the vibrations of crumpling materials. The signals are displayed through custom shoes equipped with loudspeakers and haptic actuators. Nordahl et al. [69] developed a similar approach with rigid, friction and aggregate models. They conducted a material recognition experiment, showing that some material properties could be conveyed through haptic rendering alone.

Following these aforementioned approaches for ground material rendering, we aim at leveraging real-time fluid sound synthesis algorithms to generate the relevant vibrotactile feedback. Those techniques that are physically based rely on the oscillation of air bubbles trapped inside the fluid volume [87] to produce sound. The first bubble sound synthesis technique was proposed in Van den Doel's seminal work [19] where, based on Minnaert formula [57], he provides a simple algorithm to synthesize bubble sounds based on a few parameters. However, the synthesis was not coupled to a fluid simulation. This is achieved by Drioli et al. [21] through an ad-hoc model for the filling of a glass of water, based on the height of the fluid inside the glass and on collision events. Moss et al. [60] propose a simplified, physically inspired model for bubble creation, designed specifically for real-time applications. It uses the fluid surface curvature and velocity as parameters for bubble creation and a stochastic model for bubble sound synthesis based on Van den Doel's work [19]. However, the model is designed for a shallow water simulator, which greatly reduces interaction possibilities by allowing only surface waves, precluding splashes and object penetration.

Inspired by the physically based fluid sound synthesis work of Moss et al. [60], and utilizing a particle-based fluid model [59], we develop an efficient bubble generation technique (Section 1.4) and introduce a novel vibrotactile model (Section 1.5). This enables rich body-fluid interactions with vibrotactile and multi-modal cues (Sections 1.6 and 1.7), resulting in positive user feedback (Section 1.8).

1.3 Overview of the Approach

When an object vibrates under an applied force, a pressure wave is generated at its surface, traveling to the subject's ears and mechanoreceptors. We motivate our vibrotactile approach, based on sound generation mechanisms, on the fact that acoustic and tactile feedback are both vibrations that share a common physical source.

By comparing film frames with the air-borne generated sound, Richardson [87] provides an explanation for the process of a projectile impacting and entering a fluid volume. The impact produces a "slap" and projects droplets, while the object penetration creates a cavity that is filled with air. The cavity is then sealed at the surface, creating an air bubble that vibrates due to pressure changes. Smaller bubbles can spawn from the fragmentation of the main cavity, as well as from the movement of the fluid-air interface, such as when the droplets return to the fluid volume.

Our vibrotactile model is therefore divided in three components, following the physical processes that generate sound during solid-fluid interaction [87, 27]: (1) the initial high frequency impact, (2) the small bubble harmonics, and (3) the main cavity oscillation. As a consequence, it is highly dependent on the efficient generation and simulation of air bubbles within the fluid. Hence, a real-time fluid simulator enhanced with bubble synthesis is required on the physical simulation side.

Figure 1.2 provides an overview of our approach. The physical simulator automatically detects the solid-fluid impacts and the creation of air bubbles caused by interaction between a solid (such as a foot, a hand or an object) and the fluid volume. For each of these events, it sends the corresponding

[1] http://www.soundobject.org/SDT/

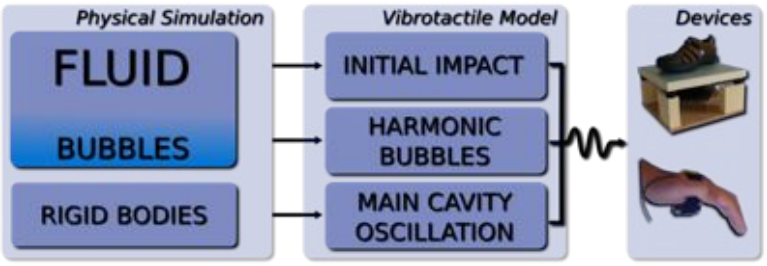

Figure 1.2: Overview of our approach: the physical simulation computes the different parameters that are fed to the 3-step vibrotactile model, producing the signal sent to the various vibrotactile displays.

message to the vibrotactile model, which synthesizes a vibrotactile signal according to the simulation parameters. The signal is then output through a specific vibrotactile device, such as an actuated tile for foot-fluid interaction, or a hand-held vibrator for hand-fluid interaction.

1.4 Fluid Simulation with Bubbles

The first building block of our approach is the fluid volume itself: we require a physically based real-time fluid simulation. Among existing fluid simulation techniques, the Smoothed-Particle Hydrodynamics (SPH) [59] model fulfills our requirements well, since the resulting fluid is unbounded, fast to compute and preserves small-scale details such as droplets. It is based on a set of particles discretizing the simulated media and conveying different physical properties, such as mass and viscosity. The smoothed quantity Q_i of a particle i at any position $\mathbf{x}$ in space is computed through the general formula:

$$Q_i = \sum_j Q_j V_j W(\mathbf{x} - \mathbf{x}_j, h) \tag{1.1}$$

where Q_j is the discrete quantity Q sampled for neighboring particle j at position $\mathbf{x}_j$, V_j is the volume of j, and W is the smoothing kernel of support h, where particles beyond h are not taken into account.

The motion of fluids is driven by the Navier-Stokes equations. Using the implementation of these equations in the SPH model [61], pressure and viscosity forces are computed at each time step. Rigid bodies are simulated as a constrained set of particles. For further details on SPH fluids, we refer the reader to references [14] and [61].

As previously explained, in order to achieve vibrotactile interaction with fluids we need to simulate the bubbles inside the fluid. Since we only seek bubble creation events resulting in bubble sound synthesis, a bubble has a very short life span within our model, and can be seen more as an event than as the actual simulation of a pocket of air. Hence, we adopt and simplify an existing SPH bubble synthesis algorithm [62] to obtain an efficient bubble creation and deletion mechanism.

A bubble is spawn when a volume of fluid entraps a volume of air. In order to detect this phenomenon within the SPH simulation, we compute an implicit color field c^p as in the method of Muller et al. [62]. This color field estimates the amount of neighboring particles (fluid, rigid and bubble) around any position in space, while its gradient ∇c^p estimates in which direction the surrounding particles are mainly located. At each timestep, we compute ∇c^p at each fluid particle position with:

$$\nabla c_i^p(\mathbf{x}_i) = \sum_j V_j \nabla W(\mathbf{x}_i - \mathbf{x}_j, h) \tag{1.2}$$

A fluid particle i triggers a bubble creation if the following conditions are fulfilled:

- the vertical component of ∇c_i^p is positive: the fluid particle has most of its surrounding particles above it, creating a pocket of air under it.
- the magnitude of the velocity of the particle is above a threshold: still or slow moving fluid

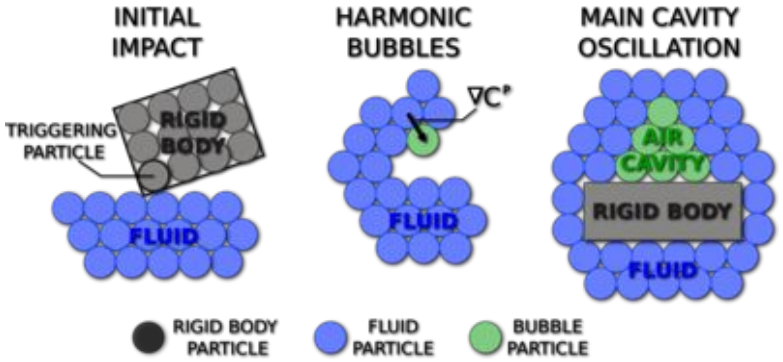

Figure 1.3: The three components of our vibrotactile model.

particles do not generate bubbles.

A bubble is destroyed when it is not entrapped by fluid or held by a rigid body anymore. Since we only use bubbles for triggering events, a bubble is also destroyed if it is alone in the surrounding media. To this end, we compute another implicit color field, c^b, which only considers bubble particles, thus estimating the amount of neighboring bubbles surrounding a point in space. A bubble i is destroyed if one of the following conditions is fulfilled:

- the vertical component of ∇c_i^p is negative: the bubble particle has most of its surrounding particles under it, and the air cannot be trapped anymore.
- the color field c_i^b is null: the particle is alone inside the media.

1.5 Vibrotactile Model

Our vibrotactile model receives the events from the physical simulation, and can synthesise a signal through 3 different components: the initial high frequency impact, the small bubble harmonics, and the main cavity oscillation, as illustrated in Figure 1.3.

1.5.1 Initial Impact

During rigid body impacts on a fluid surface, Richardson [87] observed a damped high-frequency and low amplitude sound immediately after the impact, later explained as a guided acoustic shock [54]. To the best of our knowledge, no model provides the equations for air pressure oscillations due to a rigid body impact on a fluid surface. Previous work has been able to model the phenomenon to some extent, only for very simple shapes and specific penetration cases [35]. Nevertheless, the short duration of the impact does not justify a computationally expensive implementation. Hence, similar to previous work [21], we follow a physically inspired approach exploiting the short and burst-like nature of the vibration.

Synthesis

The impact signal is synthesized in a three step approach. A burst of white noise is first generated, spanning on the vibrotactile frequency range with a given base amplitude A. The signal is then fed to a simple envelope generator in order to modulate its amplitude. The signal rises exponentially during an attack time t_a, from nil to the original amplitude A, followed by an exponential decay of release time t_r, mimicking the creation and attenuation of the short and highly damped impact. Last, the modulated signal excites an elementary resonator. For this, a second-order resonant filter is used, creating a resonance peak around a central frequency w_0. The impact signal is therefore approximated as a resonating burst of white noise, with parameters to control its amplitude (A), duration (t_a, t_r) and central frequency (w_0).

Control

An impact event is triggered when the distance between a rigid body particle and a fluid particle is below the smoothing radius. Since only the particles at the surface of the rigid body have to be taken into account to avoid false triggers, a new implicit color field c^r is computed only considering rigid body particles: particles belonging to the lowest level sets of c^r belong to the surface. Richardson [87] observed that, in general, the intensity of the impact sound between a rigid body and a fluid is proportional to v^3, where v is the speed of the body at the moment of impact. Hence, after detecting an impact, we can synthesize an impact signal of amplitude A proportional to v^3. We use the same manually set duration and central frequency parameters for all impacts.

1.5.2 Harmonic Bubbles

Small bubbles are generated by small pockets of air trapped under the water surface. Splashes and underwater cavity fragmentation are two causes for small bubble generation. By approximating all bubbles as spherical bubbles and relying on our SPH simulation enhanced with bubble generation, we can easily synthesize and control this component of the model.

Synthesis

Following van den Doel [19] and modeling the bubble as a damped harmonic oscillator, the pressure wave $p(t)$ of an oscillating spherical bubble is given by

$$p(t) = A_0 sin(2\pi t f(t)) e^{-dt} \tag{1.3}$$

A_0 being the initial amplitude, $f(t)$ the resonance frequency and d a damping factor.

Minnaert's formula [57] approximates the resonance frequency f_0 of a spherical bubble in an infinite volume of water by $f_0 = 3/r$, where r is the bubble radius. In order to account for the rising in pitch due to the rising of the bubble towards the surface, Van den Doel [19] introduces a time dependent component in the expression of the resonant frequency: $f(t) = f_0(1 + \xi dt)$, with $\xi = 0.1$ found experimentally. Taking into account viscous, radiative and thermal damping, the damping factor d is set to $d = 0.13/r + 0.0072r^{-3/2}$. As for the initial amplitude A_0, previous work [55] suggests, after empirical observations, that $A_0 = \epsilon r$ with $\epsilon \in [0.01; 0.1]$ as a tunable initial excitation parameter. For a detailed explanation of the different hypothese and equations, we refer the reader to [60] and [19] .

Control

Our bubble vibration synthesis algorithm allows the generation of bubble sounds based on two input parameters: the bubble radius r and the initial excitation parameter ϵ. Using our SPH simulation, we couple the vibration synthesis with bubble creation events and automatically select the aforementioned parameters.

If we wanted to simulate the fluid and the bubbles at a scale where the particle radius matches the smallest bubble radius that generates a perceivable vibration (12mm for a 250Hz frequency), we would require around a million particles for a cubic meter of fluid. Achieving real-time performances with this number of particles is currently highly challenging for common hardware. Since we cannot directly link the particle radius to the resonating bubble radius, we transpose the physically inspired approach of Moss et al. [60] to determine the radius and excitation parameters: power laws are used both for r and ϵ, within the ranges allowed by each parameter. When a bubble is created, values for r and ϵ are computed, and sent to the signal synthesis algorithm.

1.5.3 Main Cavity Oscillation

The main cavity is a large bubble that produces a characteristic low-frequency bubble-like sound. By modeling the cavity as a single bubble with a large radius, we can rely on our harmonic bubble

synthesis and control algorithms for this second component of our vibrotactile model.

Synthesis

As for the harmonic bubble component (Section 1.5.2), we use Equation 1.3 to synthesize the vibration produced by the oscillation of the main cavity. Since we will be using larger values for r, the resulting vibration will be of a much lower frequency, coherent with what we hear in real life. ϵ is set to 0.1 since no variability is desired.

Control

In order to detect the formation and collapsing of the main cavity during object penetration, we track the grouping of individual bubbles within our SPH simulation. Bubbles are spawned and stay alive when a cavity begins its closing and collapsing process, until they fill most of the cavity volume, as illustrated in Figure 1.3 and rendered in Figure 1.4 (right, bubbles in blue). At this point, there are bubbles within the cavity that are surrounded exclusively by other bubbles. These bubbles are detected when their color field c^b is above a threshold. If such a particle is detected, there is a potential cavity collapse.

Starting from the detected particle, we perform a search for neighboring bubbles to find the extent of the cavity. Bubble neighbors are added to the set of cavity bubbles, and the process is repeated on the new neighbors until no new neighbor is added. As the search is executed on the GPU, an iterative implementation is required, with one thread per new bubble neighbor, benefiting from our accelerated neighbor search algorithms. During our experiments, we required less than 5 search cycles to account for all the bubbles inside a cavity.

The total number N_b of cavity bubbles is proportional to the volume of the cavity. Since the cavity is modeled as a single large spherical bubble in the signal synthesis algorithm, its radius r can be deduced from the volume of the cavity. Hence, the number of cavity bubbles is mapped to the radius r of the spherical cavity, with user-defined minimal (r^{min}) and maximal (r^{max}) values: $[N_b^{min}, N_b^{max}] \rightarrow [r^{min}, r^{max}]$.

1.6 Vibrotactile Rendering

The vibrotactile model is implemented in PureData, while the SPH fluid and bubble simulation are implemented on GPU [14]. The communication between the SPH simulation and the acoustic model is handled through the Open Sound Control (OSC) protocol. Each time a bubble, cavity or impact event is detected in the fluid simulation, an OSC message is sent to the acoustic model with the corresponding parameters for sound synthesis.

We designed three scenarios representing three possible interaction conditions. For the graphic rendering, we used a meshless screen-based technique optimized for high frequency rendering, described in previous work [14]. The scenarios were run on a Core 2 Extreme X7900 processor at 2.8 GHz, with 4 GB of RAM and an Nvidia Quadro FX 3600M GPU with 512 MB of memory.

Active foot-water interaction (shallow pool). Our approach is particularly suited for foot-floor interaction, where the floor renders the vibrotactile feedback to the user's feet through appropriate vibrotactile transducers. We used a floor consisting of a square array of thirty-six 30.5×30.5 cm rigid vibrating tiles [115], rendering in the 20-750 Hz range. The virtual scene consisted of a virtual pool with a water depth of 20 cm filling the floor. The user's feet were modeled as parallelepiped rigid bodies and tracked through the floor pressure sensors. The user could walk about, splashing water as he stepped on the pool as seen in Figure 1.4 (left). Performance: 15,000 particles (1% bubbles), 152Hz.

Passive foot-water interaction (beach shore). Using the same hardware setup as the previous scenario, we designed a tidal action scene in which the user stands still and experiences waves washing up on a sandy beach, as shown in Figure 1.1. Performance: 15,000 particles (6% bubbles), 147Hz.

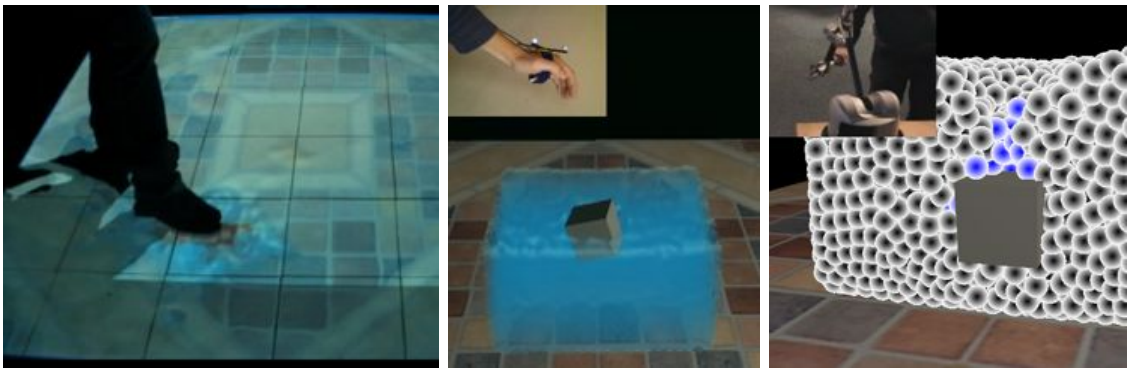

Figure 1.4: Interaction examples using different body parts and rendering devices: a foot-water pool scenario using actuated tiles (left), and a hand-water basin scenario using a small vibrator (center) or a 6DoF haptic device (right).

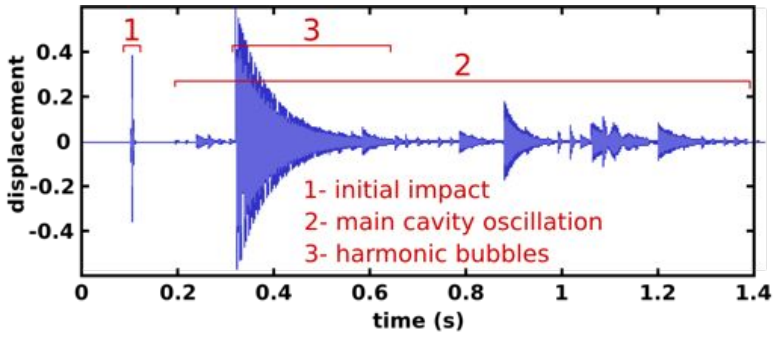

Figure 1.5: Vibrotactile signal generated with our model during a plunging movement, with its three distinct components: (1) the initial impact, (2) the small bubble harmonics, and (3) the main cavity oscillation.

Active hand-water interaction (water basin). The user can interact with fluids with his hands using a hand-held vibrotactile transducer. In this scenario, a small vibrator was attached to one of the user's hands. The hand was tracked by a motion capture system, and modeled in the virtual world as a parallelepiped rigid body. He could feel the water sensations by plunging his hand into a cubic volume of fluid, as seen in Figure 1.4 (right). Figure 1.5 shows the vibrotactile signal generated during a plunging movement. Performance: 7,000 particles (6% bubbles), 240Hz.

1.7 Extension to Other Modalities

The model can be combined with existing kinesthetic and auditory feedback techniques to achieve a truly multimodal interaction with fluids.

Kinesthetic Rendering. Kinesthetic feedback can be rendered through a suitable haptic device, such as a multiple degrees-of-freedom force-feedback manipulator. The approach is described in previous work [14], using the same SPH fluid and rigid body simulation model. This is illustrated in Figure 1.4 (right), which shows a pool of fluid with a rigid body coupled to a Virtuose 6DoF force-feedback device from Haption. The user can interact with the fluid through the rigid body while receiving kinesthetic feedback.

Acoustic Rendering. Since our vibrotactile model is built from sound generation mechanisms, we are able to produce acoustic feedback using the same model, by displaying the signal through a speaker and in the 12 Hz - 20 kHz range. Since the sound synthesis is coupled to a physically based fluid simulator, it enables richer interactions than previous real-time ones, where ad-hoc models (for the filling of a glass) [21] or shallow-water equations [60] were used.

1.8 User Feedback

The novelty of the interaction scenarios, for example, users walking on a volume of fluid and experiencing vibrotactile stimuli, motivated us to design the scenario of a pilot study to assess the perceived interaction qualitatively. The objective of the pilot study was to answer the following questions:

1. Can users recognize that they are stepping onto a simulation of water, despite the contradictory perceptual cues provided by the rigid floor surface?
2. Does the addition of vibrotactile or acoustic rendering improve the realism of the interaction compared to visual feedback alone?
3. How compelling is the overall experience?

1.8.1 Scenario

The scenario required subjects to walk on the *virtual* shallow pool of fluid described in section 1.6 and perform both a material identification and a subjective evaluation. On the initial presentation, only vibrotactile feedback was provided, with no graphics or audio. Users were asked to identify the simulated material onto which they were stepping. Since the equipment used generates auditory output in conjunction with vibrotactile feedback, we masked the audio by supplying sufficiently loud white noise through four speakers surrounding the floor and headphones worn by the users. The users were then presented three different feedback conditions: visual feedback alone (V), visual + audio feedback (V+A) and visual + vibrotactile feedback (V+Vi), counterbalanced across participants and with three repetitions (nine trials per participant). For each of the three feedback conditions, the users were asked to walk on the actuated floor for 20 seconds and complete a questionnaire, rating each condition on a seven-point Likert scale in terms of *believability* and *engagement* of the interaction.

We gathered the feedback of eight users, all naïve to the purpose of the simulation. An ANOVA on the *believability* question revealed a significant effect of the condition ($F(8, 2) = 3.5231$, p-value= 0.04555). A post-hoc analysis indicated that the V+Vi condition received significantly higher scores (p-value= 0.0417155) than the other conditions. An ANOVA on the *engagement* question did not show a significant effect of the condition ($F(8, 2) = 2.4816$, p-value= 0.1048.).

1.8.2 Discussion

The responses to our material identification question, with the presentation of only vibrotactile feedback, were highly encouraging. Six out of the eight subjects stated that they were interacting directly or indirectly with water. This user feedback suggests that our vibrotactile model, based on bubble vibrations, can efficiently convey the sensation of interacting with a fluid volume.

Further insights are available from the details of these responses. Of the six correct answers, two directly identified water, and the remaining four answered that they were walking on "plastic bottles with water inside", "a floor with wet shoes", "a water bed", and "a plank on top of water". In other words, most of the subjects who associated their experience with water felt that they were interacting with water through a solid material. We believe the lack of kinesthetic feedback of a sinking foot induced the perception of an intermediate rigid material between the foot and the water. We suggest two possible mechanisms for overcoming this sensory conflict: one involving shoes with vibrating soles to provide the vibrotactile rendering, so that stimuli can be presented prior to contact with the real floor; and a second, employing a vibromechanical actuator, as used in earlier prototypes [114], to deliver the necessary kinesthetic cues of compliance associated with foot-fluid interaction.

As expected, the subjective feedback provided by the users confirmed that the use of vibrotactile feedback greatly improved the interaction experience compared to having only visual rendering. More importantly, the results also suggest that the addition of vibrotactile feedback is more valuable for solid-fluid interaction than audio rendering in terms of the believability of the experience; it also improved engagement, although based on our limited subject pool, the result was not significant.

This shows that, despite the lack of kinesthetic feedback to compensate for the conflicting sensory cues provided by the contact of the foot with a rigid tiled floor surface, the effect was perceptually compelling.

1.9 Conclusion

In this paper, we introduced the use of vibrotactile feedback as a rendering modality for solid-fluid interaction. We proposed a novel vibrotactile model based on prior fluid sound rendering knowledge, leveraging the fact that acoustic and vibrotactile phenomena share a common physical source. The model is divided into three components: an initial impact with the fluid surface, a cavity oscillation created when the body enters the fluid, and a set of small bubble harmonics. We illustrated this approach with several fluid interaction scenarios, where users felt the fluid through vibrotactile transducers. User feedback regarding material identification solely based on vibrotactile cues suggested that the model effectively conveys the sensation of interacting with fluids, and highlighted the need for consistent kinesthetic cues.

Future work will focus on using our model with other vibrotactile devices, such as actuated shoes and vibrotactile gloves, as well as evaluating the user perception of their interaction with fluid in different application scenarios.

2 Navigation in virtual environments

2.1 Introduction

Navigation is one of the main interaction tasks in Virtual Environments (VE). Most of Virtual Reality (VR) applications need to provide the users with a way to move through the virtual world.

The most natural metaphore to achieve this goal is to simulate the natural human walking that is used in the real life for almost any navigations tasks. Moreover, the walk is a skill acquired very early in a child development [29], which can be done without even thinking.

In order to provide to the users the most realistic impression of walking, numerous types of interfaces have been developed. The navigation interfaces try to induce walking sensation within limited space by directly influencing on sense and/or body of the users.

To interact as smoothly as possible with the user, 3D user interfaces may take into account knowledge from all these fiels to provide interactions with the Virtual Environment (VE) as close as possible as the natural human walking.

In this chapter, we propose a bibliographical study of the *3D User Interfaces for interactive human walking in virtual environments*. First, we present the state of the art of the navigation techniques. Finally we focus on three new interaction techniques developed in the NIW project.

2.2 State of the Art

Many navigation techniques have been developed to simulate the walk in VE. In this section, we will make a quick overview of the state of the art techniques existing.

2.2.1 Navigation Hardware Interfaces

Natural walking is the ideal navigation interface. By mapping the user's real position to the user's virtual position, it provides the most direct and realistic way of controlling one's position. Moreover, if there is a 1:1 mapping between the real and the virtual motion, vestibular and proprioceptive feedback perfectly matches the visual perception of the virtual movement, producing an ideal multi-sensory perception of self-motion. Besides, natural walking is, after all, the locomotion interface that we use in our everyday life. However, in most simulations the virtual world is larger than the real workspace, hence 1:1 natural walking is not possible. The use of navigation interfaces is therefore required.

In this section, we will focus on devices that enable or augment natural walking.

Locomotion Interfaces

Hollerbach [31] defines a motion interface as "the mean by which a user travels through a virtual environment". It is the input device that controls the motion of the user in the virtual environment. Motion interfaces can be subdivided in two categories [22]:

- in passive interfaces, there is no user locomotive exertion. Typically, passive interfaces are related to a rate control: the user changes his motion rate in the virtual environment by controlling the displacement amount exerted on the device. He cannot directly control his position in the virtual world, since the interface is usually very small, which would require a very high number of repetitive motions. Clear examples are the joystick, where, usually, the greater the angle between the stick and the vertical the higher the speed, and a pedal, where speed is proportional to the displacement of the pedal.
- in active interfaces, or locomotion interfaces, the user is self-propulsed through a repetitive gait. Active interfaces directly control the position of the user in the virtual world. Therefore, for devices smaller than the virtual environment, a repetitive motion is needed. In situations where the limits of the device are reached, some interfaces can be switched from position control to rate control in order to navigate to farther locations.

In the locomotion interfaces category, most of the devices try to enable the use of natural walking. They usually work by compensating the user motion with an inverse motion produced by the device. For obvious reasons, we will not consider passive interfaces.

Hollerbach [31] classifies locomotion interfaces in 4 distinct categories: pedaling devices, walking in place, foot platforms and treadmills. However, we propose a slightly modified classification that we believe to be more appropriate, for two reasons. First, we do not consider walking in place techniques as a locomotion interface, but more as a 3D navigation technique, since they are not related to a specific device, and when devices are used it is only for tracking purposes. Second, many locomotion devices (mainly wearable devices) do not fit into any of the categories mentioned above.

We therefore categorize locomotion interfaces as pedaling devices, foot-based devices and recentering floors. Pedaling devices regroup mechanical interfaces where the locomotion gait is pedaling. Foot-based devices (subdivided in foot-wearables and foot-platforms) are those who compensate the motion of each foot separately, in opposition to the recentering floors that compensate the overall movement. Both classifications are based on how the devices work, not what they allow the user to do.

Pedaling Devices Pedaling devices such as unicycles or bicycles are straightforward examples of locomotion interfaces [31]. By pedaling, the user provides linear motion data to the system, and feedback can be returned through pedaling resistance. By steering, the user provides direction data.

Brogan et al. [10] used the bicycle as a locomotion device in their bicycle race environment, when studying dynamically simulated characters. The user was able to steer, brake and change gears. The device delivers the speed of the rear wheel and the turning angle of the front wheel. The bicycle was set up on a tilting platform, enabling a simulated slope of +/- 12 degrees.

The Sarcos Uniport, a unicycle, was set up on a turntable, and steering was possible through body pressure on the seat, thus freeing the hands. The possibility of turning in the real world when doing a virtual turn provided an accurate vestibular perception of the virtual rotational motion [31].

Foot-based Devices This class of locomotion interface regroups devices directly linked to the foot. It can be subdivided in foot-wearables, where the device is mainly a pair of special shoes standing on the ground, and foot platforms, where the device is a pair of moving platforms, acting as a support for the feet and allowing haptic rendering through the control of the platforms.

Foot-wearables Iwata's Virtual Perambulator [38] is a pair of roller skates and a hoop around the user's waist that limits his movement, as shown in Figure **??**. The user would slide in place while walking, with a body movement closer to a real-world movement than passive locomotion. The prototype was improved by using shoes with a low friction film and a brake pad instead of the roller skates. Extensive testing showed that 94% of the tested subjects could freely walk around the virtual space.

A more recent design from Iwata, The Powered Shoes [42], tries to overcome the need from the user to apply an unnatural extra force when sliding for recentering his position. The recentering is done automatically with roller skates actuated by two motors. The movement of the feet, measured by optical sensors, is compensated by a movement of the active roller skates on the opposite direction. However, the main limitation of this system is that the direction of motion is the same than the direction of the shoe, and therefore gaits like side-walking are not possible [43].

Iwata's String Walker [43] is a rather different implementation of a foot-wearable device, since the device is not made of shoes alone, but also of a surrounding turntable to which the shoes are attached through a set of 8 strings. Each string is actuated by a motor-pulley mechanism. The motion of the user is compensated by pulling the shoes in the opposite direction through the strings. The pull is applied to the foot in its stance phase. The device allows omni-directional walking thanks to the turntable that repositions the strings according to the user's orientation. Side-stepping is also possible, following the same mechanism than for regular walking.

Foot platforms Individually programmed foot platforms allow each feet to be positioned in the 3-dimensional space. The forward motion of the user is countered by a reverse centering motion, keeping the user in the center of the device. The platform can either follow the foot, or apply a force that limits its movement, from a recentering motion to a simulated ground.

The Sarcos Biport has 2 hydraulic 3 linear degrees-of-freedom (DOF) arms, where the user's feet are attached. The user can also pivot on the foot mounts. The device recenters the user's position translation-wise, but cannot compensate rotations, and the arms can occasionally interfere with each other [44]. It can simulate uneven terrain and smooth surfaces, but allows only moderate walking due to the structural stiffness of the device.

Iwata's Gait Master [44] is made of a pair of 3 linear DOF motion platforms mounted on top of a turntable. The feet, attached to the motion platforms, are recentered by the device, with the user's rotations compensated by the turntable. In the specifications, the Gait Master can be a 6 DOF locomotion device. However, a device of only 4 DOF was implemented, which is enough to simulate uneven terrains (a set of plain surfaces) but not inclined floors. The device was evaluated on its ability to render an uneven terrain by assessing the naturalness of the user's walk through the measure of the soles pressure on the platforms. The results showed that the force applied at the heel was close to the one on a real stair, and that differences compared to real case values at the toes were due to the use of a safety bar.

The HapticWalker [90] is another example of foot platform. It differs from the other devices in its capacity of allowing high-speed walking (120 steps/minutes) and complex movements (stumbling, sliding), and in its bigger size. It allows a 6 DOF motion, but as with the Gait Master a more restricted version was implemented, with 3 DOF in the saggital plane, allowing only linear walking. The device was successfully tested by many unexperienced subjects.

Recentering Floors

Treadmills A treadmill is a moving platform made of a conveyor belt and an electric motor. A user walking on a treadmill is able to reproduce the gait of walking while standing in place, since the conveyor belt repositions the user to the center of the platform. The main advantage of treadmill-like devices is the realism of vestibular and proprioceptive feedback, since the user is naturally walking forward without constraints.

1 DOF treadmills Linear treadmills, with only 1 linear DOF, are simple, straightforward, and yet satisfactory locomotion devices.

The Sarcos Treadport is a 4 by 8 feet treadmill with a mechanical tether attached to the back of the user's harness. The tether is used for tracking, safety and force rendering purposes. The second generation has a 6 by 10 feet surface and a fast tilting mechanism [34]. The belt speed is proportional to the user's distance to the center of the platform. The ability of applying forces on the user's torso through the tether is used to compensate missing inertial forces, and simulate obstacles in the virtual path. The device allows peak walking velocities of 12mph, 1g accelerations, and different body postures, like crawling and crouching [31]. However, turning capabilities are very limited.

The tether can also be used to simulate slopes. Hollerbach et al. [33] add a force on the opposite direction of motion when simulating walking uphill, with a magnitude equal to the horizontal component of the force in the real world case. Analogously, the force is applied on the same direction of motion when going downhill. Simulation of side slopes are also possible when applying lateral forces [32].

ATR's ATLAS [65] enhances linear treadmills by allowing the user to rotate, and by naturally representing slopes in any direction. The device is a 145 by 55cm linear treadmill mounted on an actuated spherical joint. The joint works like a turntable regarding the user's rotations, and allows the pitching and rolling of the platform in order to simulate slopes. However, rotations are limited, and the treadmill has to be recentered when the yaw reaches its limit. The user's feet are tracked through markers on the shoes and a video tracking system. The maximum belt speed is 4m/sec.

In order to simulate uneven terrain, ATR developed the Ground Surface Simulator (GSS) [66], made of a 150 by 60cm linear treadmill with a deformable belt. Six 25cm long platforms can locally push the belt up to 6cm, at a speed of 6cm/s. A tensioning system ensures the belt fits the changing geometry of the treadmill. The GSS can render a 5 degrees.

With the opposite approach, Bouguila et al. [5] exploit a rotational DOF instead of a linear one. Their device is a 70cm diameter turntable which recenters the user after he has performed a body rotation, so that he is always facing the same direction (the same display device) in the real environment, hence providing natural vestibular and proprioceptive cues for rotational motions. Although the device recenters the user with a smooth rotation in order to go unnoticed, large and fast rotations compensations might be perceived by the user, generating innapropriate kinesthetic cues. User rotations are tracked through cameras and two optical markers on the user's back. The forward motion is achived through walking in place. The device was specially conceived with semi-immersive environments (such as CAVEs) in mind, so that users will avoid looking at the missing parts of the display.

Omni-directional treadmills

The main problem of linear treadmills being the impossibility of performing natural rotations, some devices were developed with a 2D motion surface in mind.

The Omni-directional Treadmill [17] tackles the problem by using two orthogonal belts, one inside

the other. The outer belt is made of rollers that rotate around the axis perpendicular (in the horizontal plane) to the motion axis of the belt, which is the same axis of rotation than the inner belt. When the inner belt rotates, the rollers rotate with it, therefore creating a motion perpendicular to the motion of the outer belt. The user can walk in any direction, since his movement can be compensated in any direction by combining the two perpendicular rotation motions of the treadmill. The usable surface of the device is 1.3 by 1.3m, and the user can walk up to 2m/sec. The main limitations of the device are system lags and inaccurate tracking, which causes the user to stumble. Moreover, the walking gait has to be adapted in order to successfully walk on the device.

The Torus treadmill [44] is a set of 12 treadmills connected side-by-side, forming a belt that rotates in the opposite direction (in the horizontal plane) to the individual treadmills. The main drawbacks of the device are its usable area (1 by 1m), and the maximum walking speed (1.2m/s) a user can achieve, which forces the user to walk slowly and with short steps.

The Cyberwalk [92] uses a different approach. Inspired by the Ball Array Treadmill [104], it has one linear and one rotational DOF, instead of the two linear DOF of the previous devices. A small treadmill, generating the linear DOF, is mounted on a turntable, generating the rotational DOF. An array of balls, set up above the treadmill but independent from the turntable, sums up the 2 motion vectors, delivering a single motion that recenters the user when walking on top of the ball array. The usable surface has a 80cm diameter, and the user can walk up to 1.5 m/s in translation, and can rotate up to 2rad/s. The considerable friction between the balls and the shoes wears the soles, and the wear debris can degrade the device. Therefore, the user needs to wear shoes with a specific kind of sole material.

Tiles The CirculaFloor [39] is a set of tiles with omnidirectional motion capability. Together, the tiles form an infinite surface by constantly circulating under the user's feet. Through the tracking of the feet, the tile on which the user is standing moves in the opposite direction of the user's motion direction, therefore compensating the motion of the step. Meanwhile, the available free tiles are repositionned on the intended path of the user, in order to compensate his next step. Since the tiles can be positioned around the user in any direction, the user can freely change his motion direction while walking. Although in theory the system allows an omni-directional natural walking, implementations difficulties restrict the user to very slow walking motions. The device does not allow a normal speed locomotion.

Spheres With the Cybersphere [26] and the Virtusphere [50], it is the user that gives to the floor its recentering motion. As show in Figure **??**, the user is inside a sphere, and his movement creates a momentum which acts as a recentering motion.

Augmenting Interfaces

In order to augment a user's walking capabilities in a virtual environment, many devices have been used for the accurate tracking of different aspects of a user's footsteps. A complementing group of devices employ haptic techniques to convey information to the user about his surrounding virtual environment.

Input/Sensing devices There are many devices that have been developed to capture footsteps, either through contact, force or weight measuring. Most of them are used for artistic performances and entertainment, like augmented dance floors or the commercially available "Dance Dance Revolution" stage and "Wii Balance Board". They answer to different needs, since each application has a specific set of requirements: the need of pressure information together with the contact coordinates, the number of simultaneous users, the nature of the movements, the need to identify the whole foot or just part of it [58].

The Podoboard [47] is a 91 by 102 cm grid of 1 square inch aluminum tiles, which are used as conductors. It was designed for performing clackage (a type of step dancing) in any surface. The user

wears shoes with metal contacts and piezoelectric films at the toes and heel that close an electrical circuit when they touch the conductors grid. The contact coordinates can be quickly (10 nanoseconds) and easily retrieved. The piezoelectric films measures the impact force. However, the device has many drawbacks [30]. The user needs to use special shoes, and the shoes contact might interfere with the user's activity on the floor, such as dancing (other than step dancing). Moreover, light steps may be missed if good electrical contact is not ensured. Applications that require information about the part or the position of the foot that is in contact with the floor will not find the Podoboard suitable.

The MIDI Dance Floor [80] removes the need to use special shoes by designing a floor made of strips of 1 by 4 feet velocity and pressure sensitive regions, with a granularity of around 6 inches. Each region holds 32 Force Sensing Resistors (FSRs), and is assigned a MIDI note, so that the user can synthesize music while performing. However, the device does not deliver precise contact coordinates, since the application does not require it.

The Magic Carpet [73] is a 16 by 32 grid of piezoelectric wires, spaced by 10cm. When pressed or flexed, the wires produce voltage. Since the wires are crossed, the location of the contact can be determined by knowing which wires were activated in each dimension of the carpet. The voltage is converted into a 7bit pressure value. The temporal resolution of the carpet is around 2Hz. The small spatial and temporal resolution are main drawbacks [30] for a highly precise application, together with the fact two simultaneous contacts cannot be precisely located due to the crossed-wire technique. Two x coordinates and two y coordinates are generated, but which ones correspond to the same contact?

The ORL Active Floor [1] is a 4 by 4 array of 50 by 50cm tiles. Each tile has a compressive load sensor on each corner, and each load sensor can measure up to 500kg with a 50g resolution (10bits) and a temporal resolution of 500Hz. The ORL was designed for large spaces, such as offices, in order to track moving objects along the floor. Again, the small spatial resolution would be a major drawback for applications demanding precise coordinates of contact. In the same line than the ORL Active Floor, the Smart Floor [71] achieves footsteps recognition using a steel plate with load cells.

The LiteFoot [30] aims at precisely tracking and recording the footsteps of one or many simultaneous users. The LiteFoot is a 1.76 meter square matrix of optical proximity sensors, with a spatial resolution of 40mm. Each sensor is considered as a pixel, and changes of state (1 bit data per sensor: has changed or not) are transmitted with a temporal resolution of 100Hz. The floor can be used in two modes. In the reflective mode light is emitted from under the floor, and the footstep is detected when the foot reflects the light back to the sensors. This requires the foot to wear a reflective sole. In the shadow mode, the floor is lit from above, and it is the blocking of the light that triggers the sensors. Since this implementation could not measure the pressure of the feet on the floor, the use of an accelerometer was added to provide the impact force of the feet. However, data provided from a single accelerometer might not be as precise as what could be gathered with pressure sensors. Moreover, the spatial resolution provided by the floor might not be high enough for precise tracking applications [98].

The Z-Tiles [86] aim at having a modular floorspace, in size and shape. Each Z-Tile is made of 20 blocks, arranged in a particular shape that allows the tiles to be interlocked with each other. The Z-tiles have a spatial resolution of 40mm. Each block has one force sensitive resistor (FSR) for reading pressure values, from 30 to 900N (12 bits), at 100HZ. They do not improve the spatial resolution of previous devices.

A recent pressure sensing floor [98] overcomes many of the limitations described so far. The floor is made of 48.8 by 42.7cm mats with 2016 pressure sensors each, a spatial resolution of 10mm , and a scan rate of 30Hz. The device is modular, since mats can be interconnected without leaving unscanned regions. Coupled to video, audio and motion-based tracking, it aims at creating a robust and accurate tracking system.

With a different goal in mind, the Ada system [18] was designed for large-scale collective user interaction. Ada's 136 square meters floor is made of 360 hexagonal 66cm tiles. Each tile has an FSR on 3 of its vertices, and 3 colored neons across. The floor can sense and, to a certain level, track the users, which allows Ada to interact with a croud through dynamic visual effects.

With a completely different approach, in-shoe sensor devices have been developed for capturing footstep data.

The most complete example of such a device is the CyberShoe [75]. While other devices restrict themselves to the measure of a small set of parameters (speed [36] or pressure [13] related), the CyberShoe captures 16 different parameters per shoe. FSRs and a piezoelectric film gather pressure data, resistive bend sensors measure the bidirectional bend of the sole, a dual-axis accelerometer provides the pitch and roll values of the foot, a 3-axis accelerometer detects large impacts and fast kicks, a magnetic sensor gives the angular orientation, a gyroscope measures the angular speed about the vertical axis, a piezoceramic transducer coupled to sonar transmitters retrieve the foot's position in the space, and an electric-field-sensing electrode detects the foot's height when in proximity to conductive strips located on the floor. The data is gathered at 60Hz and broadcasted wirelessly through an FM radio transmitter. A 9 Volt battery powers each shoe, with the battery state also transmitted with the data. The hardware is set up on a regular sports shoe.

One particular device was developed specifically for virtual reality environments such as CAVEs. The CyberBoots [13] are boots with FSRs in their soles, at the heel, the inner and outer ball, and the toe tip, covering the main pressure areas of the foot/floor contact. The user wears the boots on top of his own shoes. The FSRs and the cables are inserted between two layers of vinyl, and the cables are drawn to the arch of the boot from where a terminal cable climbs up to the waist of the user. At his waist, the user wears a small interface box containing the electronics. The pressure data is used to determine if the user is walking naturally, walking in place, or leaning, and the corresponding gestures are used as navigation commands.

Output/Haptic devices Haptic devices are hardware interfaces that enable the communication from machine to user through the sense of touch. The device applies forces and/or vibrations to the user in order to simulate a virtual force, a virtual contact, or more generally to provide tactile sensory feedback.

Unlike with foot sensing devices, there are a very limited number of foot haptic devices. Some locomotion devices, as previously described, are considered haptic devices ([90]) since they provide kinesthetic feedback when simulating uneven terrain. Forces are applied, through the feet, to the entire body. However, devices where the haptic stimuli is located exclusively on the feet, and therefore mainly related to contact and tactile information, are scarce.

The Fantastic Phantom Slipper [49] is a pair of shoes with markers for optical tracking and vibrators on the sole for haptic rendering. The user can sense and interact with virtual objects projected onto the floor. However, with actuators located only on the sole, there are limited rendering possibilities. Moreover, the shoes are not meant for walking.

The FootIO prototype [88] is an insole with pager motors used as actuators for the foot sole. Each motor can deliver 15 different intensity levels of stimuli, in a wide frequency range. The device can provide localized stimuli or static and dynamic patterns throughout the sole, the main goal behind the design of the FootIO being the human-to-human communication over a network. A perception experiment showed that users were able to accurately (90-100%) recognize patterns in the longitudinal direction, but failed to accurately identify (45%) the direction of transversal patterns mainly due to the motors layout. A higher number of motors and a bigger spacing between them leads to a better recognition of the direction of the pattern. The FootIO is still at an early stage, and it is not yet usable as a foot based haptic interface. Walking is not yet possible.

The Haptic Shoes [28] use a simple haptic mechanism, similar to previously described haptic interfaces, to provide haptic cues. Each shoe holds 3 cell phone motors (used to produce the vibrating ring-tone in cell phones) in the sole, connected to a wireless receiver. Data is transmitted wirelessly from a distant computer. The device aims at haptically conveying stock market data through the feet. No perception study is provided.

ATR's ALive Floor (ALF) [106] aims at simulating irregular surfaces. The floor is made of a 1 by 2 meter matrix of 28 actuator units. Each actuator unit is made of 6 tilting triangle plates arranged in

an hexagonal configuration. The actuator can control the tilt of the triangles by raising or lowering a shaft on its center, with a maximum height of 10cm. The device can therefore display a patched irregular surface in real time.

Focusing on the haptic rendering of textures, Visell et al. [116] developed the EcoTile, a vibrotactile device than can simulate different types of ground materials. The user, walking on top of the actuated floor tile, can sense the type of material he is virtually standing on. Physically based models of contact interaction have been used for sound synthesis, and were adapted for vibrotactile haptic rendering. The vibrating device is made of a single motor actuator, with a frequency response ranging from 10Hz to 18kHz, while the tile itself is a thin 12 by 12 inch polycarbonate layer, placed on top of the motor and held by a frame. FSRs are placed between the tile and the frame in order to detect the user's feet pressure. The device generates a vibrational as well as an auditory stimuli according to the user's footsteps. So far, the snow material has been successfully rendered. A sensing shoe with FSRs, acoustic vibration transducers, an acoustic microphone and an accelerometer is currently being developed for the measuring of the interaction between the feet and ground materials. It will provide data for the accurate rendering of the different ground materials.

2.2.2 3D Navigation techniques

Static Navigation

Several relatively well-known VR locomotion metaphors

- Teleportation. This is an instantaneous change of position, once the new location has been chosen. This seems desirable, but it has been found that it causes disorientation and simulator sickness.
- Worlds In Miniature (WIM). In this model, the user has a copy of the whole virtual world in their hand(s). They can choose a location in the copy and be smoothly taken to the location in the "real" VW. Some drawbacks include difficulty in using the model without haptic feedback.
- Scene In Hand. Here the virtual world is supposed to be attached to the user's hand, so that simple hand movements manipulate the world and therefore change the user's position.
- Eye In Hand. Similar to WIM, except the user is told to "imagine" the miniature copy of the world.
- Flying Vehicle. In this model the environment is not manipulated. The illusion is that the user can move through the world, either by using a mock-up or a wand or other device.
- Leaning. In this model the user moves by leaning in the direction they wish to travel.
- Voodoo Dolls. Similar to WIM, except that the user can create miniature "dolls house" of any part of the world, and then move through the house, and manipulate the "dolls".

These metaphors are implemented within a virtual world. This means they must take the confines of the system into account. Many of these "metaphors" do not map well-understood source domains. This may reduce their usefulness as metaphors.

Walking in place

The Walking in Place technique [95] allows the user to navigate in a limitless virtual environment without moving in the real environment, but with an increased degree of presence compared to other static navigation techniques. By "walking in place", the user simulates the physical act of walking but without forward motion of the body, and a virtual forward motion is introduced. Therefore, the visual optical flow of navigation that should be matched by proprioceptive information from the natural walking gait is matched by proprioceptive information from a gait close to natural walking. The user's sense of presence is greatly increased compared to static navigation techniques, as shown in [112]. However, other sensory data that completes the illusion of walking, mainly vestibular cues, are not perceived by the user.

The Walking in Place technique can be combined to natural walking. The user moves freely in the working space, but whenever he has to go beyond the limits of the physical space he can switch to a Walking in Place navigation. The direction of motion is naturally defined by the user's gaze direction.

Compared to previous techniques where navigation was controlled through the use of the hands (hand-held input devices, gestures), this technique has a major advantage since navigation is entirely devoted to the region of the body that performs that task in the real world. Therefore, what could be done in the real world while walking, can also be done in the virtual one: the hands are free to be used in any other task, such as the manipulation of virtual objects.

In [95], the technique does not require additional hardware, the existing head tracking system is enough. Walking in place is detected through pattern analysis using a neural network, based only on the head position data over time. However, there is a training period of around 5-10 minutes. If walking in place is detected, the user is translated in the virtual world giving the illusion of real motion. The system yielded a 91% accuracy rate in average.

Additional tracking devices are used in the Gaiter system[102]. Six degrees of freedom sensors are attached to the knees, transmitting translation and rotation values. The Walking in Place function is triggered when the system detects an excess of motion of the legs compared to the normal gait. This can happen when actually walking in place, or when adding extra motion to a real walking motion (hybrid steps). Unlike the previous technique, the direction of motion is defined by the legs motion direction, and not the user's gaze direction, which corresponds to real world locomotion. Hence, the system allows forward, backward, and side-stepping motion.

Natural Walking

Several techniques try to adapt natural walking to restricted size workspaces.

The Step Wim [51] allows a user to invoke at his feet a miniature version of the world. Inspired from the World in Miniature navigation (WIM) technique [76], the miniature world enables the user to travel distances greater than his real working environment. Upon invocation, the WIM appears on the ground with the position of the user in the virtual world corresponding to the position of the user's feet in the real world. The user can then walk on the WIM to a new position and trigger a rescaling command that will scale up the WIM until it reaches the virtual world's size. The scaling animation was chosen instead of translation techniques since it was shown [76] that it helps the user's sense of location. The invocation and scaling commands are triggered by bringing together the heels or the toes of the user's shoes. Conductive cloth patches located at the heels and toes of the shoes ensure the detection of the contact between the two shoes.

Since the Step Wim is implemented on a CAVE environment, the proposed technique solves the limited workspace problem translation wise, but not rotation wise. As in any semi-immersive environment, for a given position and orientation of the user there is a portion of the scene that cannot be displayed. In the case of the CAVE, the main problem are the missing walls. In order to minimize the missing wall problem, the authors of the Step WIM also developed an auto-rotation technique that would ensure that the user would never face a missing wall. The technique is based on a mapping of the 360 degrees of the scene to the system's display field. A linear mapping proved to be uncomfortable since it quickly produced cybersickness. Therefore, the mapping depends on the user's orientation and position inside the CAVE, and is matched by applying a scaling coefficient to the user's head rotation. The closer the user is to the front wall, the larger his field of view, hence the rotation scaling is not as large as if the user were at the edge of the CAVE, where the field of view is the narrowest. The technique is noticed by the user, who needs a small adaptation time in order to adjust to the new mapping, and does not entirely solve the missing wall problem.

Working with Head-Mounted Displays (HMDs) and large virtual environments, Williams et al. [121] explored different techniques to overcome the limited working environment problem when using natural walking for navigation. They proposed and compared 3 different techniques that would reset the user's position in the real world when he reaches the workspace limits, without breaking his spatial awareness of the virtual world. The main idea is to warn the user when he is about to reach the limits,

so he can execute the resetting motion and keep traveling in the same direction he has been traveling in the virtual world. In the freeze-backup technique, the virtual world is frozen and the user takes steps backwards to recenter his real world position inside the workspace. In the freeze-turn technique, the orientation of the user is frozen while he physically pivots 180 degrees, while in the 2:1-turn technique a 360 degrees virtual rotation is mapped to a 180 degrees real world rotation, and the user also physically pivots 180 degrees For each technique, the spatial awareness was evaluated according to 2 criteria: the angle error when asked to face, eyes closed, a particular object after a reset, and the time the user took to face the object. Results showed that the freeze-backup technique had the lowest angle error, making it better suited for applications where spatial orientation is important, while the 2:1-turn had the lowest time, more appropriate for applications where time is a constraint. Both were superior to the freeze-turn technique.

The Seven League Boots [37] are another technique to naturally walk in virtual worlds larger than the real one. It relies on the scaling of the user's speed only along his intended direction of travel, unlike with a uniform gain where both components of the forward motion are scaled. In order to detect the intended direction of travel, an approach considering both the user's gaze direction and the direction of previous displacement over the last reasonable period of time is proposed. Both criteria are weighted according to the amount of displacement over the last period of time: if the user stood still or did not travel much, the gaze criteria will be weighted higher, but if the user showed a will to move, the displacement will be weighted higher. However, the experiments conducted in order to evaluate the Seven League Boots used a simpler version of the technique, since the intended direction of travel was pointed by the user through a wand. The goal of the experiments was to show how comfortable it was for the users to have their speed scaled only on the intended direction of travel. Therefore, the simpler version was compared to natural walking, natural walking with uniform gain, and wand flying. Results showed that, overall, the Seven League Boots were more appreciated than the other techniques. No quantitative evaluation measuring spacial awareness has been reported.

Redirected Walking The Redirected Walking [82] technique brings a solution to the limited workspace problem, very close to the ideal. When experiencing Redirected Walking, a user is tricked into walking in a curved path in the real world when walking in a straight line in the virtual world. In a sufficiently large workspace, and with a straight virtual path [101], the user can naturally walk endlessly without reaching the limits of the real workspace, and without noticing that he is actually walking in circles. The technique relies on the interactive and imperceptible rotation of the scene around the user, so that he is always walking towards the farthest side of the real workspace. The "redirection" takes into account the user's position, orientation and speed. However, in small workspaces such as CAVE environments it is not possible to keep the user away from the walls by injecting only unnoticeable amounts of rotation distortion: there is a trade-off between workspace size and user perception of the rotations. Such a trade-off can be avoided by using controlled events that force the user to look around him, hence allowing the injection of additional rotation distortion.

In a pilot user study, Razzaque et al. [82] showed that the technique worked well in a 4 by 10 meters workspace with a virtual room twice the size and a predefined path. Users had to follow that virtual path and press a button at each turning point. They were told to look at the next button once the had pushed the one next to them. The systematic rotation of the user to face the new button was the occasion to inject the rotation distortion. Then, if needed, more could be injected during the walking phases. The tested subjects could complete the test without noticing the extra amount of rotation or running into walls. However, the path was predefined: the user could not freely wander about the room.

Since Redirected Walking relies on the unnoticeability of the redirection, Steinicke et al. [101] statistically quantified the maximum amount of rotation distortion that could be injected without being noticed by the user. Moreover, they added the possibility to scale translational movements, and quantified the maximum scaling that could be introduced without altering the user's perception of his virtual self motion. After a user study that tested several users with different rotation and translation

gains, results showed that, statistically, rotations scaled down or up to 30% and translation scaled down to 15% and up to 45% are perceived in less than 20% of all cases. Therefore, a user walking on a straight virtual path can be redirected into a 3.3 meter radius circle. However, these results are subjective, varying from a user to another.

Peck et al. [78] introduced the use of distractors in order to force the user to rotate his head and allow the injection of rotation distortion. They compared Redirected Walking coupled to a distractor technique (a flying moving sphere) with other reorientation techniques such as direct scene rotation, mapping of the 360 degrees of the virtual world to 180 degrees in the real world and Redirected Walking with audio instruction to force the rotation of the head. The virtual scene was a 200 meter straight path with checkpoints every 5 meters, while the real workspace was a 5 meter straight path. At every checkpoint, the users were presented one of the four reorientation techniques, which made them rotate around themselves in the real world, but kept them in the same direction of motion in the virtual world. Results showed that the distractor and the instructed head rotation techniques were rated significantly higher than the others, in terms of presence, preference and naturalness. In a second experiment with different subjects, a butterfly distractor was compared to the sphere distractor and the instructed head rotation, showing a preference for the more "natural" distractor technique.

Redirected Walking has been used in applications beyond its original goal of natural walking navigation. Razzaque et al. [84] explored the use of Redirected Walking with a walking in place technique in a CAVE environment. They wanted to avoid looking at the missing wall of the CAVE when exploring a virtual environment. By comparing the technique with a hand-held navigation technique, they showed that the frequency of looking at the missing wall was not reduced, although the variance was. In a different context, Kohli et al. [48] showed how to take advantage of the Redirected Walking technique to use passive haptics (simulating the haptic feedback of a virtual object through the use of a real object) on multiple identical virtual objects with only one real object. They use a distractor (a moving target that had to be neutralized) triggered at specific synchronization zones in order to make the real object coincide with the virtual object faced by the user. As with the previously described Redirected Walking applications, the path in the virtual world has to be predefined and carefully chosen.

Motion Compression Along the same idea than Redirected Walking, Motion Compression [64] rotates the virtual environment, keeping the user away from the limits of the real environment and allowing him to explore a virtual world larger than his workspace. However, Motion Compression differs from Redirected Walking in many aspects. First, rotations and translations are mapped with a 1:1 ratio: no rotation (or translation) is injected, only the environment is rotated. Hence, distractors cannot be used. Second, the unnoticeability of the rotations is not a constraint: the user can perceive the manipulation, specially when close to the limits, where it is simply not possible to bend the user's path by unnoticeably rotating the virtual world. And third, the user's path in the virtual world does not need to be predefined: without the unnoticeability constraint, paths can be greatly curved, giving to the user a higher freedom of movement. Instead, the path is predicted in real time.

There are three steps in the Motion Compression algorithm: path prediction (finding the intended virtual path of the user), path transformation (mapping the intended virtual path to a real path inside the boundaries of the workspace, with minimum curvature and preserving distances and angles) and user guidance (applying the rotations as the user walks the path). In Nitzsche et al. implementation, path prediction is based on the user's gaze direction and a set of manually selected potential targets. A potential target is activated as the actual target if it stayed the longest, among the other potential targets, inside the user's field of view. The more a user looks at a potential target, the higher the probability the target is the actual path target. Path prediction is performed at every update of the tracking system. However, it requires the availability of a model of the virtual world. Path transformation is done by minimizing the total curvature deviation.

In a first experiment, Nitzsche et al. tested different curvature deviations in order to find the perceptual threshold under which subjects could not tell the direction of the curvature. The threshold

was found equal to $0.1m^{-1}$. In a second experiment they simulated a 30 by 20 meter virtual museum hall, with a 3 by 3 meter real workspace. With a curvature deviation usually greater than $0.3m^{-1}$ due to the small workspace, the virtual visit was a positive experience for most of the users regarding comfort and intuitiveness.

Engel et al. [23] proposed a different technique for path transformation by including a perceptual component. Instead of minimizing the curvature deviation, they look for a path that both maximizes the distance to the boundaries of the workspace and minimizes the discomfort of the user. A perceptual study allowed the construction of a psychometric function that gives, for each rotation gain, the probability of noticing the gain. The function is then translated into a cost function representing the level of discomfort of the user for each rotational gain.

2.3 JoyMan: Navigating in Virtual Reality Using Equilibrioception as a Human Joystick

In order to provide VR users with realistic sensations of locomotion while keeping them in their limited workspace and without any break in the immersion, numerous types of VR interfaces have been proposed so far. The most straightforward solution probably consists in walking naturally in the real world, using for instance Head Mounted Display (HMD) together with a tracking system which measures head position for updating virtual camera position and orientation in the virtual environment. However, this solution can rarely be used due to the limited physical workspace in the real world, as well as limitations of the tracking devices.

To overcome this problem, numerous "locomotion interfaces" have been designed. The locomotion interfaces keep the user within the limits of the physical workspace while allowing to navigate infinitely in the virtual world. A classical hardware approach consists in using a treadmill which can be either unidirectional (such as the Sarcos Treadport [99]) or omni-directional (such as the Torus Treadmill [45]). Other locomotion interfaces can be based on feet-platforms [41], walking inside a rotating sphere [25] or a set of moving tiles [40], standing on a Cybercarpet [91]. These locomotion interfaces have the advantage of providing realistic walking sensations, involving proprioceptive information. However, the current locomotion interfaces are often too expensive, cumbersome, or complex for being integrated easily in both VR applications.

Besides, classical input devices such as mice, keyboards and joysticks can also be used within interaction metaphors to navigate from one virtual place to another [6]. However, such interaction techniques, relying exclusively on manipulated input devices, do not provide any proprioceptive or vestibular feedback of walking and lead to a poor sensation of locomotion in the VE.

Thus, numerous interaction techniques have been developed to provide software-based navigation capabilities in VR. Software solutions have been proposed in order to keep the user inside the real workspace limits while he is really wallking, such as the World in Miniature technique [77], the Magic Barrier Tape [15] or the Freeze-Backup technique [120]. However these techniques induce breaks in the feeling of immersion during the navigation in a virtual environment. Redirected Walking Technique [83] or Motion Compression [63] preserve the immersion by forcing the user into walking in a curved path in the real world while walking in a straight line in the virtual world. However, these techniques require large workspaces, can be confusing when doing unpredictable or quick changes of direction, and may require distracting events.

Other solutions have been proposed to enhance the sensation of walking in virtual worlds without modifying the user position in the real environment. An efficient solution consists in playing with visual feedback and adding artificial camera motions [52]. The Walking-In-Place technique has been introduced by Slater et al. [96] to enable a real physical walking movement and an efficient navigation technique in 3D virtual environments. In this interaction paradigm, the user is walking in place in the real world, allowing proprioceptive feedback while keeping him in the real workspace. However, the absence of real global displacement neglects equilibrioception.

In this section, we would like to explore the possibility of an interface that aims at preserving equilibrioception in place of proprioception. The ChairIO project [3] also implies equilibrioception but the user is seated while he is navigating in the virtual world, decreasing the immersion, especially for locomotion tasks. Our new interface is based on the metaphor of a human-scale joystick. Some devices have already been built [105] but they are not designed for enhancing immersion in VE, especially as they are not combined with a control law allowing to reproduce locomotion tasks. The main principles and the design of our new interface, called Joyman, are detailed in the next section.

2.3.1 Description of the Joyman Interface

The Joyman is a new interface to navigate into virtual worlds. One main objective of this interface is to go towards realistic locomotion trajectories in the virtual environment. The Joyman meets this objective by combining two components: a peripheral device allowing users to indicate the desired

direction of locomotion and a control law which transforms the device state into a virtual velocity vector. The mechanical design of the device is based on the metaphor of a human-scale joystick and mainly consists of a board upon which the user stands. The inclination of this board can be changed by the user. Thus, the user indicates the desired locomotion direction by tilting the platform in the corresponding direction. The mechanical design of the platform allows the user to change the platform inclination by leaning and prevents him from falling, whereas repealing forces tend to maintain the platform horizontal. The control law transforms the device state (pitch and roll angles) into a virtual locomotion velocity vector. The proposed law ensures that humanly feasible velocities are achieved during virtual locomotion. Particularly, the tangential velocity is bounded and depends on the desired angular velocity. These two components are detailed below.

Device Mechanical Design

The proposed mechanical design of the interface is proposed in Figure 2.1.

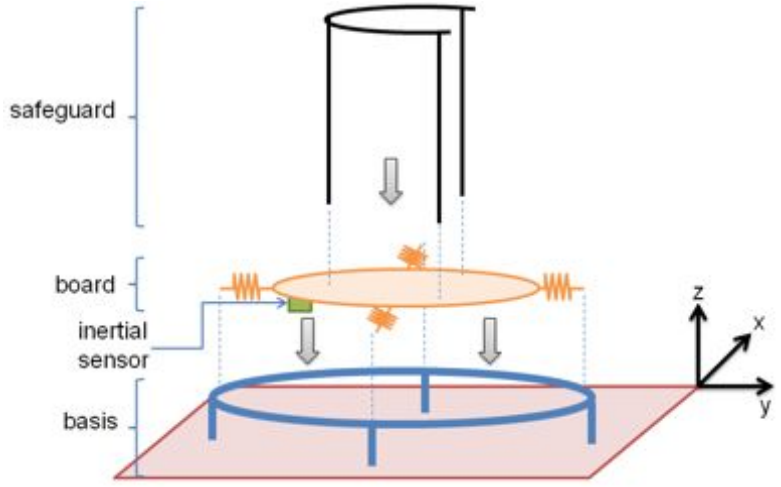

Figure 2.1: Schematic representation of the main elements composing our peripheral device. The board and the basis are articulated using springs. Springs generate repealing forces that tend to maintain the board in the horizontal plane, however, the board can be freely oriented around the two horizontal axis (pitch and roll) with limited range. Users act on the device by standing on it, and apply forces in order to tilt the board by leaning. The board inclination is measured using an inertial sensor. The safeguard prevents users from falling when leaning.

The key principle of the Joyman is to let the user indicate the desired direction of his virtual locomotion by leaning in the corresponding direction. In other words, our objective is to let a user control his locomotion from his equilibrioception. This main functionality stressed at first the mechanical design of the Joyman device.
The main properties of our design are :

- the user should be able to stand on the platform, which is the natural position in the task of locomotion;
- the user should feel perfectly safe when using the platform which, as a result, should prevent him from falling;
- the user should be able to lean over the limit of his own balance in order to increase the amplitude of vestibular sensations;
- the device should tend to bring the user back in vertical position to avoid highly tiring manœuvers during its use.

We also carefully considered other criteria during the device design process: the most important was to build an affordable interface. Exploring the equilibrioception in place of proprioception to control

oneself locomotion in the virtual world is a promising direction to lower these costs. However, we did not want to lower the costs by degrading the feeling of immersion provided by the interface. We aim at building a new interface that does not involve any active mechanical device in its movement and does not imply the use of sophisticated and complex tracking devices. We meet the objectives listed above based on a simple mechanical device as illustrated in Figure 2.1.
The device is made of 4 components:

The basis supports the whole device. It is made of a flat square platform lying on the ground and of a circle shaped steel frame linked to the board by legs. The legs and the platform are fixed together and the board is large enough to prevent the device from tipping over.

The board is the element supporting users. It is circle shaped and linked to the frame by springs. As a result, the board is free relatively to the frame within limited range of motion. We can approximately but reasonably consider that the board is articulated to the frame by two degrees of freedom which are rotations around the x and y axis as illustrated in Figure 2.1.

The safeguard is rigidly fixed to the board and prevent users from falling.

The inertial sensor is rigidly attached to the board and acquires the current orientation of the board.

The basis is made of a square woodcut platform. The steel frame is the one of a trampoline where the springs have also been retrieved (1 m diameter). The basis is a machined plastic plate (0.55 m diameter) equipped with 18 hooks on its rim. A rigid rope is tied up between the hooks and the springs attach the board to the frame. Finally, the safeguard is made of welded iron tubes. Its height can be adapted to each user (height between 0.8 m and 1.2 m). Building costs did not exceed 500$ in spite of the fact that this device is a unique prototype excluding the inertial sensor.

The user can stand straight on the platform and has to lean in any direction to start the locomotion in the virtual world. The control law that allows to navigate in a VE with the new interface is described in the next paragraph.

Virtual Locomotion Control

The main steps of the virtual locomotion control are summarized in Figure 2.2. The control is composed of two elements detailed in the following paragraph: a locomotion model and a control law.

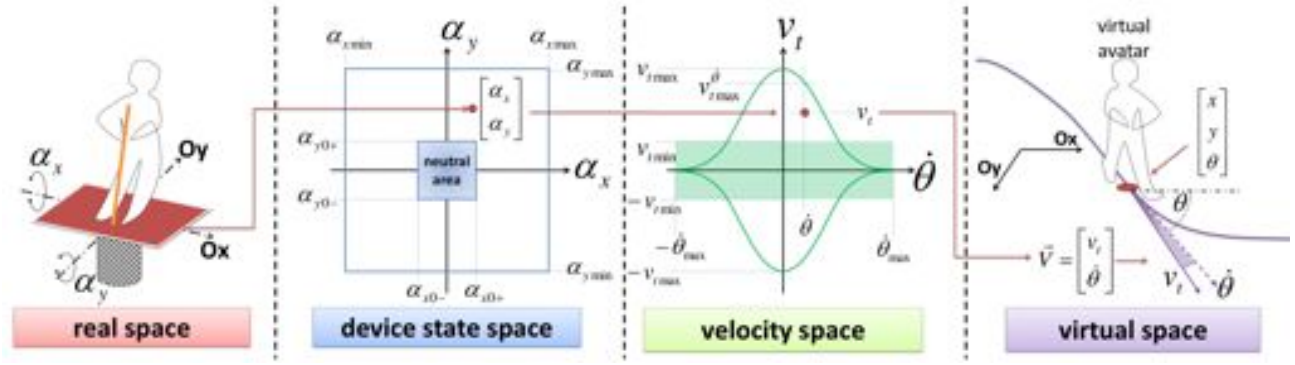

Figure 2.2: Summary of the different steps for the virtual locomotion control. From the position and orientation of the Joyman, we can compute the virtual velocity vector by using our locomotion model and control law.

Locomotion model The locomotion model allows to translate the position and orientation of the platform into a virtual motion in the VE. We exploited experimental observations of the walking human trajectory to build our model [2].

We model the position P of the user in the virtual space as an oriented point moving in a horizontal plane (see Figure 2.2):

$$P = \begin{bmatrix} x \\ y \\ \theta \end{bmatrix} \tag{2.1}$$

The virtual motion is velocity controlled. We assume the virtual trajectory is non-holonomic, which means that the velocity vector orientation and the body orientation are always identical. The non-holonomic nature of walking human trajectory has been experimentally observed in [2].

Non-holonomy constraint allows us to decompose the velocity vector V as follows:

$$V = v_t . \begin{bmatrix} cos(\theta) \\ sin(\theta) \end{bmatrix} \tag{2.2}$$

Such a decomposition allows us to independently control the tangential speed v_t and the orientation θ. Human tangential velocity is limited. We denote $v_{t_{max}}$ the maximum tangential velocity bound. By default, we set this bound at $v_{t_{max}} = 1.4\ m.s^{-1}$. It has equally been observed that very slow walking velocities are never reached: such velocities are humanly feasible but not used in practice. We thus define $v_{t_{min}} = 0.6\ m.s^{-1}$. It has been however observed that changing orientation θ affects the amplitude of v_t during human locomotion [123]: humans decelerate when turning. In order to take into account such an observation, we define $v^{\dot{\theta}}_{t_{max}}$ the bound of the tangential velocity knowing the current turning velocity $\dot{\theta}$. $v^{\dot{\theta}}_{t_{max}}$ is defined as follows:

$$\begin{aligned} v^{\dot{\theta}}_{t_{max}} &= a.v_{t_{max}}.e^{-b.\frac{\dot{\theta}}{c}} \\ v^{\dot{\theta}}_{t_{min}} &= a.v_{t_{min}}.e^{-b.\frac{\dot{\theta}}{c}} \end{aligned} \tag{2.3}$$

where a, b and c are parameters. As a result, we model the reachable tangential velocity $v^{\dot{\theta}}_{t_{max}}$ as a Gaussian function of the current turning velocity $\dot{\theta}$. The higher $\dot{\theta}$, the lower the tangential velocity bound. By default, we arbitrarily choose: $a = 1.07$, $b = 0.5$ and $c = 0.7$. Such values match the experimental observations provided in [123]. Finally, the absolute value of the angular velocity is also bounded to $\dot{\theta}_{max}$. We arbitrarily choose $\dot{\theta}_{max} = 1\ rad.s^{-1}$.

Control Law The control law allows users to modify the virtual velocity vector V by standing on the device and leaning. The modification of the platform orientation affects the state of the device s_d. s_d is defined by the orientation of the board α_x and α_y relatively to the two horizontal axis $\vec{Ox}$ and $\vec{Oy}$, as measured by the inertial sensor:

$$s_d = \begin{bmatrix} \alpha_x \\ \alpha_y \end{bmatrix} \tag{2.4}$$

We neglect the orientation around the Oz axis. During calibration stage, we ask the user to stand on the platform and to successively firmly lean towards all the cardinal directions. We estimate the reachable bounds of the board orientation by averaging the reached orientation over a short period of time, $\alpha_{x_{min}}$, $\alpha_{x_{max}}$, $\alpha_{y_{min}}$ and $\alpha_{y_{max}}$. We also ask the user to stand straight on the platform and define a neutral area bounded by $\alpha_{x_{0+}}$, $\alpha_{x_{0-}}$, $\alpha_{y_{0+}}$ and $\alpha_{y_{0-}}$.

We want the user to control his tangential velocity v_t by leaning in forward or backward direction, i.e., by playing on α_x, whilst angular velocity $\dot{\theta}$ is controlled by leaning on the sides, i.e., by playing on the α_y value. The angular velocity is controlled as follows:

$$\begin{cases} \dot{\theta} = \dot{\theta}_{max}.\frac{\alpha_{y_{max}} - \alpha_y}{\alpha_{y_{max}} - \alpha_{y_{0+}}} & \text{if } \alpha_y > \alpha_{y_{0+}} \\ \dot{\theta} = \dot{\theta}_{max}.\frac{\alpha_y - \alpha_{y_{min}}}{\alpha_{y_{0-}} - \alpha_{y_{min}}} & \text{if } \alpha_y < \alpha_{y_{0-}} \\ \dot{\theta} = 0 & \text{otherwise} \end{cases} \tag{2.5}$$

which allows to deduce $v^{\dot{\theta}}_{t_{max}}$ and $v^{\dot{\theta}}_{t_{min}}$ according to equation 2.3 and finally v_t:

$$\begin{cases} v_t = v^{\dot{\theta}}_{t_{min}} + (v^{\dot{\theta}}_{t_{max}} - v^{\dot{\theta}}_{t_{min}}).\frac{\alpha_{x_{max}} - \alpha_x}{\alpha_{x_{max}} - \alpha_{x_{0+}}} & \text{if } \alpha_x > \alpha_{x_{0+}} \\ v_t = -v^{\dot{\theta}}_{t_{min}} + (v^{\dot{\theta}}_{t_{max}} - v^{\dot{\theta}}_{t_{min}}).\frac{\alpha_x - \alpha_{x_{min}}}{\alpha_{x_{0-}} - \alpha_{x_{min}}} & \text{if } \alpha_x < \alpha_{x_{0-}} \\ v_t = 0 & \text{otherwise} \end{cases} \tag{2.6}$$

At each time step n, the virtual velocity vector $V_n = [v_{t_n}, \dot{\theta}_n]$ is thus deduced from s_d, the state of the device. Before updating the simulation accordingly, we check that no unbelievable acceleration is performed. Thus, given the previous velocity vector V_{n-1}, we finally compute the effective velocity

vector V_f:

$$V_f = V_{n-1} + \lfloor \frac{V_n - V_{n-1}}{\Delta t} \rceil \tag{2.7}$$

where Δt the simulation time step, and $\lfloor \rceil$ a function that truncates the velocity vector variation so that the absolute tangential acceleration does not exceed 1 $m.s^{-2}$ and the angular one does not exceed 1 $rad.s^{-2}$.

Summary of Joyman Interface

The Joyman interface is designed for immersive virtual locomotion into virtual worlds. Our main aim was to explore the possibility of an interface that tends to preserve the equilibrioception in place of proprioception, in contrary to many current interfaces. The mechanical design of our platform is relatively simple. We showed how the proposed device allows a user to safely perform exaggerated leaning motion, over the limit of balance, in order to indicate his virtual navigation wills. We here remind our long-term objective to be able to perform realistic locomotion trajectories into the virtual world. What can we expect from the Joyman interface?

Immersion The device does not limit technical choices about visual or audio feedback. Typically, the device can be directly used in any immersive environment. More important, the initiation of the locomotion is made by leaning ahead. This corresponds to the real motion to initiate walking, which is considered to be a constant loose of balance and succession of falling ahead. We expect this interaction mode, where vestibular sensory system stimulation is preserved, to be greatly immersive.

Realism We expect that the dynamics of the required manœuvers to operate the Joyman are close to the ones of real walking motions. As an example, using the Joyman, users perform a left turn following a right turn by changing the platform orientation toward its opposite using their whole body. Involved inertias and frictions prevent from performing this change of state immediately: this reproduces the walking behavior during which the inclination of the body is naturally adapted to face centrifuge force. We clearly do not expect users to achieve efficient navigation (in terms of task completion time) using the Joyman.

2.3.2 Preliminary Evaluation

As a preliminary evaluation, we propose to compare our interface to the joystick, a classical peripheral often used in VR applications. The joystick can be considered as one of the most performant interface to achieve efficient navigation in terms of task completion time. Thus, the first objective of our evaluation was to quantify the loss of performances of our interface compared to the joystick (our interface involves namely the whole body compared to the joystick where only hands and arms are used). The experiments were conducted using 3D VE displayed either on a screen or on a Head-Mounted Display (HMD). We investigated the effectiveness of our interface to travel complex paths composed of different gates placed in the VE. The second objective of our evaluation was to verify if we can obtain a more immersive navigation with our interface compared to classical peripherals already used in VE, ie. the joystick. Thus, a subjective questionnaire was proposed to the participants to evaluate their subjective preferences in terms of quality of the VE navigation.

Performances: Task Completion Time and Accuracy

For each participant, the task completion time of each trial was measured for the different experimental conditions. A repeated two-way ANOVA was performed on the two different interfaces and the two visual conditions. The ANOVA accounting for the visual conditions and the task completion time revealed a significant effect ($F(1, 526) = 11.63$, p-value< 0.001). A significant effect was also found for the interfaces and the task completion time ($F(1, 526) = 54.47$, p-value< 0.001). As expected, the

results reveal that the joystick was better than the Joyman in terms of speed of the navigation. The mean value for the completion time of a path was $187s$ (standard deviation=$13s$) for the joystick and $321s$ (standard deviation=$89s$) for the Joyman.

A specific analysis was developed to study the learning effect of the two configurations (Joyman and Joystick). A linear model where all conditions are mixed was fitted to explain the relation between the task completion times and the trial number. For the Joyman configuration, it revealed that the slope of the linear regression was significantly lower than zero (p-value < 0.000001), reflecting a significant decrease in the task completion time as the number of trials increases. The same analysis, where the first trial was removed, showed that the slope was not significantly different from zero anymore (for both visual conditions). Figure 2.3 illustrates the task completion time of the different trials for the Joyman configuration.

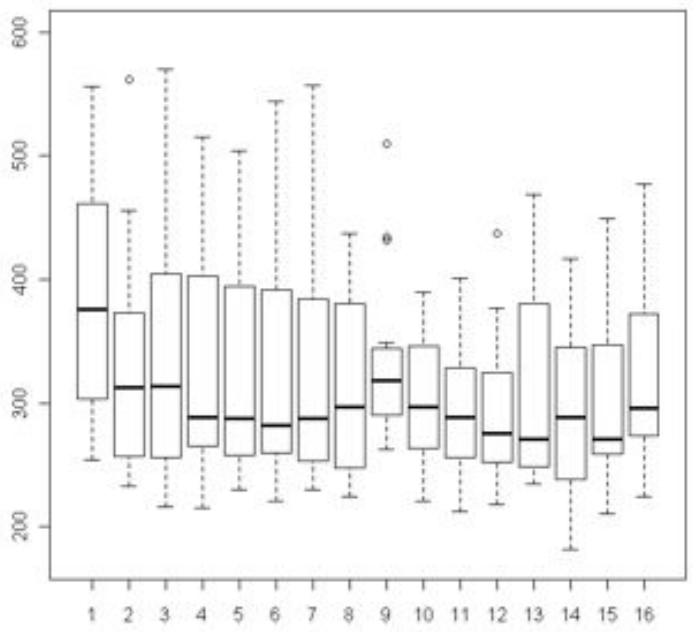

Figure 2.3: Task completion time (in seconds) for the Joyman configuration (16 trials). The first trial was left to illustrate the learning effect. Each boxplot is delimited by the quartile (25% quantile and 75% quantile) of the distribution of the condition over the individuals. The median is also represented for each trial.

For each participant and for each trial, the percentage of errors for the different paths was measured. There was no error at the end, for both configurations (Joyman and Joystick) and both visual conditions.

Subjective Questionnaire

A preference questionnaire was proposed in which participants had to grade from 1 (low appreciation) to 7 (high appreciation) the two configurations (Joyman and Joystick) according to 8 subjective criteria : (a) Fun, (b) Intuitive, (c) Acccuracy, (d) Presence, (e) Rotation realism, (f) Fatigue, (g) Cybersickness and (i) Global appreciation. The grade 7 for Fatigue and Cybersickness respectively means that the interface does not induce any fatigue and does not imply any cybersickness feeling.

After performing an ANOVA on the two different conditions, we found a significant effect for 6

criteria: Fun ($F(1,30) = 17.77$, $p < 0.001$), Intuitive ($F(1,30) = 21.25$, p-value< 0.001), Accuracy ($F(1,30) = 23.52$, p-value< 0.001), Presence ($F(1,30) = 13.35$, p-value< 0.001), Rotation realism ($F(1,30) = 6.63$, p-value$= 0.015$) and Fatigue ($F(1,30) = 87.51$, p-value< 0.001). In particular, our new platform was better ranked for Fun, Presence and Rotation realism. No significant effect was found for Cybersickness ($F(1,30) = 4.01$, p-value$=0.054$) and Global appreciation ($F(1,30) = 0.69$, p-value$= 0.411$).

2.3.3 Discussion and Perspectives

As a global conclusion of the preliminary evaluation of the Joyman, we can state that the feeling of immersion in the virtual world is significantly improved - in comparison with traditional joystick-based techniques - at the cost of some easiness of use. The Joyman is still at a early stage of development, however, first results are promising and open a large set of possible directions to improve usability and the level of realism of virtual navigation. This section discuss as exhaustively as possible these directions as well as future work to meet our objectives: an easy, intuitive, immersive interface allowing realistic locomotion in virtual worlds.

Interface Calibration

The navigation with the Joyman implies the whole body and, as expected, the task completion times were higher for our interface compared to the joystick where only the arms and hands are involved. The use of the whole body, implying more movements, is also a reason for the lower rating for the Accuracy criterion in the subjective questionnaire. However, the results concerning the evaluation of our interface in terms of performances are encouraging as the participants always succeeded to complete the navigation task during the evaluation. One way of improvement for the performances of our interface could be the interface calibration. The control law directly transforms the angle of the interface platform into virtual walking velocities. 3 types of parameters were proposed to design our control law: those define the active angles area into the device state space ($\alpha_{x_{min}}$, $\alpha_{x_{max}}$, $\alpha_{y_{min}}$ and $\alpha_{y_{max}}$, $\alpha_{x_{0+}}$, $\alpha_{x_{0-}}$, $\alpha_{y_{0+}}$ and $\alpha_{y_{0-}}$), those controlling the reachable virtual velocities ($\dot{\theta}_{max}$, $v_{t_{min}}$ and $v_{t_{max}}$) and finally those controlling the dynamics of the relation between tangential and angular velocities (a, b, and c, see Equation 2.3). Ideally, the first type of parameters should be calibrated for each user. Indeed, given his size, weight and strength, each user provides relatively different efforts to reach a same given platform inclination angle (i.e., a given walking velocity). The slight fatigue reported by participants in the results of our experiments confirms the need for individual calibration. Currently, calibration process consists in recording leaning motions as well as neutral positions to define the bounds of the effective zone for the platform. We made such a calibration before experiments, but kept a unique one all along our evaluation in order to have the same behavior of our interface for all participants. For VR applications, we could envisage to explore various calibration techniques as well, such as by displaying a virtual environment with a moving point-of-view and to ask users to apply the effort on the platform they believe to be corresponding to the motion.

Mechanical design

We presented in this paper the first prototype of our new interface. The results of the subjective questionaire suggest that the participants enjoyed the navigation with the Joyman. They namely gave a higher rating for the Fun criterion but also for Presence criterion and Rotation Realism criterion. Specifically, the higher rate given to this last criterion confirms that our mechanical design was appreciated for locomotion tasks in the virtual environment. A way of improvement for our mechanical design could concern the linkage between the platform and the basis. The proposed prototype implements this linkage using an inextensible rope and a set of springs (see Figure **??**) as we wanted to have a simple mechanical design. However, we could envisage to modify the linkage to increase the possibility of movements of our interface. Ideally (mechanically speaking), the linkage could be a 2 rotational degrees of freedom, one with a restoring force proportional to the platform inclination.

Control Law

Experimental evaluation reveals that the proposed control law is intuitively grasped by users: establishing a relation between linear and angular velocities seems to be naturally accepted by users, and is consistent with observations of human locomotion trajectories. Future work will deal with potential modifications of the control law, for example by using a law that can lower vestibular and visual sensory conflicts.

The device state vector is currently two dimensional: the two angles that describe the platform inclination. With some experience, it appears that the device can be fully and accurately controlled by involving the lower body only, it is not even required to hang on to the barrier (except for moving backward in the current state of the device). Such a property opens interesting perspectives and makes possible to increase the dimension of the device state space. Possible extensions are numerous, and immediately within reach if the VR system is equipped with tracking abilities: hands remain free to achieve secondary actions in the virtual worlds (grasping, touching, pointing tasks, etc.), view direction and locomotion control can be decomposed, etc. Nevertheless, we observed during experiments that most of the participants intuitively attempts to control their locomotion also by moving their upper body, in spite of the inefficiency of such motions to significantly increase the inclination of the platform. We however could use this input (i.e., the orientation of the spine relatively to the hips) to control navigation based on an holonomic locomotion model: as opposed to the non-holonomic one, lateral velocities are allowed in addition to tangential and angular ones by removing the constraint imposed by equation 2.2 (lateral velocities here correspond to side-steps).

Immersion and sensory feedback

Evaluation showed that the Joyman provides a satisfying level of immersion into the virtual world. In addition, the Presence criterion was better ranked for our interface, compared to the joystick. As futur work, we could envisage to improve the level of immersion by adding other interaction techniques. Thus, the visual perception of motion could be improved by adding camera motions like in [103] for example. Thus, having oscillating view point could reinforce the accuracy of the perception of the traveled distance by reproducing the natural oscillations of the head during human locomotion. This could also lower the feeling that locomotion is too slow as reported by participants in the subjective questionnaire.

2.3.4 Conclusion

In this section, we presented a novel interface to control locomotion into a virtual world while remaining globally static in the real space. The Joyman interface is composed of a new peripheral device and a dedicated control law which transforms the device state into a virtual locomotion velocity vector. Our main contributions are (1) to explore the users ability to exploit equilibrioception to control their virtual locomotion and (2) to maintain a high level of immersion compared to handheld devices (e.g., joysticks).

After a preliminary evaluation, we obtained promising results as the users enjoyed the navigation with our new interface. Presence and realism in the virtual rotations were also underlined. The evaluation also opened several future work directions to improve and extend our interface. Various VR but also real applications could be envisaged, when navigating in a 3D world. Our interface could be used for example for videogames, rehabilitation, training tasks or virtual visits.

A patent for the interface presented in this paper has been filed on November 19, 2010 under the number FR10/595551.

2.4 Shake-Your-Head: Walking-in-Place Using Head Movements

Navigation is one of the fundamental tasks needed for 3D interaction with Virtual Environments (VE) [8]. The possibility to walk inside the VE is necessary in many applications of Virtual Reality (VR) such as for virtual visits (review of architectural and urban projects) or virtual training (technical procedures, military infantry).

In order to provide VR users with realistic sensations of walking while keeping them in their limited workspace, numerous types of VR interfaces have been proposed so far [8]. Locomotion interfaces, such as unidirectional or omnidirectional treadmills [100][46], are hardware solutions that enable the user to walk infinitely. But they are often costly, complex and cumbersome. Thus, many software interaction techniques have been developed based on classical 3D input devices such as joysticks or gamepads [8]. But they induce very limited physical movements, and sometimes generate a sensation of "flying" in the virtual scene rather than walking [53].

The Walking-In-Place (WIP) technique has been introduced by Slater et al. [97] to enable a real physical walking movement and an efficient navigation technique in 3D virtual environments. The user has to consciously walk in place while motions of his body are tracked and analyzed. The tracked Walking-In-Place motion is used as input for the locomotion simulation inside the VE. First implementations of WIP were all based on the processing of head positions using a neural network. More recent models track the positions of the heels or knees of the user to compute the resulting virtual locomotion [24][119]. However, all existing WIP techniques require the user to stand up, and they focus on immersive VR applications based on sophisticated tracking devices and head-mounted-displays or CAVE setups.

In this section, we propose to revisit the whole pipeline of the WIP technique. Our intention is to extend it to match a larger set of configurations, by notably applying it to desktop setups. As a result, we propose to improve the WIP technique on three main aspects, corresponding to the three following main innovations:

- **Novel user interface capacities**. We introduce (1) the possibility for the user to stand or sit (using classical interaction paradigms of desktop VR with mouse), and (2) the possibility to use screens (with limited field of view) and low-cost tracking (webcam) by using the use of head movements as main input of WIP control laws.
- **Novel locomotion simulation**. We extend the range of possible motions with WIP to: walking, turning, jumping and crawling, using heuristics and control laws based exclusively on users' head movements.
- **Novel visual feedback of walking**. We introduce the possibility to combine our approach with more realistic visual feedback of walking, i.e. camera motions, that are known to improve the sensation of walking [53].

2.4.1 Shake-Your-Head: Revisiting WIP for Desktop VR

We propose to revisit the whole pipeline of the Walking-In-Place technique to match a larger set of configurations and apply it to the context of desktop Virtual Reality. Our approach is schematized in Figure 2.4. The Figure highlights the main differences between our approach and the classical and existing WIP techniques.

With the Shake-Your-Head technique, the user can be standing or sitting (such as in traditional video games or desktop VR configurations). The user interacts with the system by means of head movements. These head movements can be captured using different tracking interfaces, but we insist on the use of low-cost optical tracking with standard webcams. The locomotion simulation proposes not only the computation of a virtual walking motion but also turning, jumping, and crawling possibilities. As a result, the user can perceive the locomotion in the virtual world by means of integrated virtual camera motions on the three axes of motion, to further enhance the sensation of walking.

In the following section, we will successively describe the different parts of our approach, namely: (1) the 3D user interface input/output, (2) the interaction techniques developed for the computation

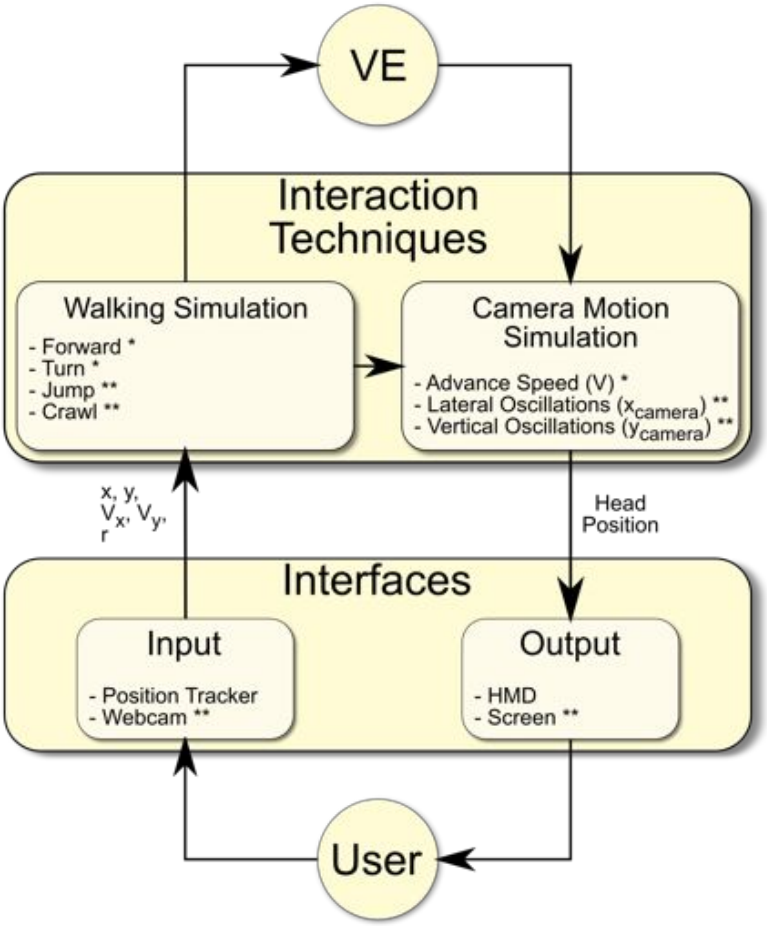

Figure 2.4: Overview of our novel approach for Walking-In-Place (one star stresses improvement of existing component, two stars stress additional components).

of the virtual locomotion, and (3) the visual feedback relying on camera motions.

Input/Output Interfaces

Our method proposes new features in terms of interfaces in order to extend the set of configurations where WIP can be applied, especially for Desktop VR. Thus, we propose to incorporate new devices for both input and output user interfaces in the VE.

Input: Tracking Based on Head Motions The input interface of our method is only based on head motions. In our implementation, the input interface is reduced to a webcam, allowing the use of our method for Desktop configuration without any additional peripheral. However, our method can also be implemented with other classical VR tracking systems.

Use of Head Movements. The main concept of our method is to exploit the head oscillations as a transposition of the one observed during natural walking. While walking, the head of the user oscillates along the lateral, vertical and forward axes [53]. The oscillations are strongly correlated to gait events and foot steps. Moreover, these oscillations also occur while walking in place and can be measured.

The head motions are classically retrieved in the existing WIP techniques thanks to the use of regular position trackers [97]. More generally, any tracking device can be used, as long as its accuracy is within the range of $1cm$. Moreover, the acquisition process of the required position does not have any influence, as long as the real time constraint is maintained. In our method, we propose the use of

a video camera system to handle the tracking of the user head. Thus, our interaction technique can be deployed on a large scale at low cost for training purpose or video games for example.

Extracted Data. In our method, we propose to use 3 Degrees of Freedom (DOF) that can be easily accessed in the image frame provided by the webcam:

- The lateral position x (and the computed speed V_x);
- The vertical position y (and the computed speed V_y);
- The rotation of the head in the frontal plan r.

Implementation. The use of the webcam to track the 3D position of the user head without using markers requires the implementation of real time constraints for the algorithms, i.e. more than 25 frames per seconds. Our implementation is based on the Camshift (Continuously Adaptive Mean Shift) algorithm [9] implemented in the OpenCV library. This algorithm is based on color tracking and is well-suited for real-time tracking of features of a given color, such as the face of the user.

While the user is standing in front of the webcam, our algorithm recognizes him as an ellipsoid. The position of the head (x, y) can be deduced from the center of the ellipse and the orientation angle of the head r is given by the angle with the vertical of the ellipse. However, x and y depend on the resolution of the used webcam. Thus, we compute a normalized position (x_n, y_n) on [-1; 1] on both axes. From the normalized positions of the head (x_n, y_n), we compute the instantaneous speeds V_x and V_y, and we use Kalman filters on all values to reduce the noise produced by the algorithm.

Output: Immersive and Desktop Visual Displays Our method can be used with both immersive and regular screens as output interface with the VE. A requirement of desktop VR applications was to propose a technique that is usable with a limited Field of View (FOV). In regular WIP techniques, the provided FOV is always 360 degrees, except in [85] where they used the redirected walking technique to simulate a 360 degrees FOV in a 4-walls CAVE. In this paper, our technique is evaluated with both LCD laptop screens to fulfill the Desktop VR pre-requisite or video projected output. Our method can also be used with the classical output used with WIP, i.e. with HMDs or CAVEs.

Walking Simulation

Walking States The main goal of our interaction technique is to translate the inputs of the user, i.e. head motions, into virtual motions in the VE. The user should be able to perform various motions while navigating in the VE. We implemented different locomotion states: walking, turning, jumping and crawling. To manage these different states, we added a state automaton to our algorithm. The state transitions are governed by the user head motions and the main inputs are the lateral velocity V_x and the vertical velocity V_y.

Forward State The forward movements in the VE are governed by the lateral oscillations as main input. Our technique is designed to emphasis the idea of walking with a variating advance speed depending on the user head motions. The advance velocity V_a oscillates regularly, accordingly to the lateral head motions. One head oscillation period corresponds to one step. The footstep events are simulated by a null advance speed and correspond to a modification of the lateral velocity sign of the user's head. Thus, when the user's head reaches the maximal amplitude of the oscillations, the oscillating speed is null, as well as the advance speed, simulating a foot step.

For more realistic movements, we introduced two thresholds T_{min} and T_{max}. T_{min} allows to stop the forward movements when the lateral head motions are too small. T_{max} allows to avoid unrealistic high speed walks. The advance velocity V_a is computed in two steps in order to test these thresholds.

The equations of V_a are:

$$\begin{aligned} V_{n_1} &= \frac{min(abs(V_x), T_{max})}{T_{max}} \\ V_{n_2} &= \begin{cases} 0 & if\ V_{n_1} < T_{min} \\ V_{n_1} & otherwise \end{cases} \\ V_a &= V_{n_2} * V_{max} \end{aligned} \tag{2.8}$$

Finally, the advance speed V of the camera inside the VE is adapted depending on the current locomotion state and is given by:

$$V = \begin{cases} V_a & if\ state = walk \\ 0.4 * V_a & if\ state = crawl \\ V_{max} & if\ state = jump \end{cases} \tag{2.9}$$

We chose normalized thresholds with the following values: $T_{min} = 0.05$ and $T_{max} = 0.5$. We also set V_{max} to 3.5 meters per second, corresponding to the maximal speed that can be achieved.

Jump and Crawl States Comparing to existing WIP techniques, we chose to add 2 new states to the navigation possibilities in the VE: Jump and Crawl motions. The jump and crawl states are governed by the vertical oscillations of the user's head. If the vertical velocity exceeds normalized thresholds T_{jump} in upward direction and T_{crawl} in downward direction, the user can jump and crawl respectively in the VE. In practice, it means that the user will need to slightly jump or bend forward if he is seated, or jump or crouch down if he is standing. The user has to stand-up to stop crawling.

When a jump is detected, the vertical position of the camera is set to follow a classical parabolic trajectory defined by:

$$y_{camera} = \frac{1}{2} * g * t^2 + V * t \tag{2.10}$$

with g the gravity acceleration and t the time. The jumping state is left automatically while landing, i.e. when the camera reaches again its normal height H (known as the reference state when the algorithm starts). After preliminary testings, we set $T_{jump} = 0.3$ and $T_{crawl} = 0.4$. While the crawling state is activated, the vertical position of the virtual camera y_{camera} is lowered by 1 m.

Turn State In parallel to Forward, Jump and Crawl states, the user has the possibility to turn inside the VE in order to modify his navigation direction. During a turn in a normal walk, the human body leans slightly in direction of the center of the turn to compensate the centrifugal force [16]. This phenomenon is often reproduced by video games players which can tend to lean in the direction of the turn even if it does not have any influence on their in-game trajectory. Thus, we choose to use this property to implement turns in our system as a control law based on the head orientation on the roll axis. To turn in the VE, the user has to lean his head in the left or right side respectively to turn left or right in the VE. The rotation speed V_r of the virtual camera is given by:

$$V_r = \begin{cases} V_{r_{max}} & if\ r > r_{max} \\ -V_{r_{max}} & if\ r < -r_{max} \\ 0 & otherwise \end{cases} \tag{2.11}$$

where r_{max} is the minimum angle of head inclination to start the rotation and $V_{r_{max}}$ is the maximal angular speed of the rotations. In our experiment, we set $r_{max} = 15 degrees$ and $V_{r_{max}} = 45 degrees/s$.

Visual Feedback Based on Camera Motions

To further emphasize the perception of walking in the VE, we extended the visual rendering of the WIP using camera motions driven by the user's head oscillations. There are existing models in the literature that make the virtual camera oscillating along the three axes ([53] for example). However, the oscillations are totally independent from the user interactions.

We introduce a new model of camera motions adapted to the user's head motions. The camera oscillations along the different axes must follow the user in real time to maintain the coherency of the system. Thus we have implemented a novel visual feedback with camera motions along the vertical, lateral and advance axes.

Advance Oscillations The advance speed V of the view point already oscillates. The camera motions are indeed intrinsically linked to the advance velocity of the control law presented in paragraph 2.4.1. As a result, extra camera motion is not necessary along this axis and the advance camera velocity corresponds exactly to V.

Lateral Oscillations In order to move in the VE, the user has to make his head oscillating from left to right. Thus, as the user moves in front of the screen, his view point of the scene is modified to follow the head oscillations.

The lateral oscillations of the camera are computed as a function of the user's position. If d is the distance of the user to the screen and α and β the opening angles of the webcam, the real world position of the user in front of the screen depends on the normalized coordinates x_n and y_n. The real world position of the user's head is given by the following coordinates:

$$\begin{cases} x_{real} &= x_n * d * tan(\alpha/2) \\ y_{real} &= y_n * d * tan(\beta/2) \end{cases} \tag{2.12}$$

Finally, the virtual camera is moved along the lateral axis by a distance x_{camera} equals to: $x_{camera} = A_x * x_{real}$. We set the scale factor $A_x = 1$ to match the user's head displacement and thus generate the illusion that the screen is a window through which the user can observe directly the VE. However, other values can be used to amplify the camera motions for example. The webcam used during the experiment was such as $\alpha = 60 degrees$ and $\beta = 45 degrees$.

Vertical Oscillations The vertical oscillations of the camera can not be computed with the same algorithm as for lateral oscillations. In a desktop VR context the user can be seated and not be able to produce high vertical oscillations.

In our method, we propose to generate pseudo-sinusoidal vertical camera oscillations based on the current phase of the virtual gait cycle. Similarly to advance speed control law, the vertical amplitude y_{camera} of the camera oscillations is given by:

$$y_{camera} = V_{n_2} * y_{camera}^{max} \tag{2.13}$$

where y_{camera}^{max} is the amplitude of the vertical oscillations for velocities greater or equal to the T_{max} threshold. For smaller speeds, the amplitude of the oscillations is proportional to this maximum, thus increasing the perception of the variations in advance speeds. Using the same factor between the camera motions and the advance velocity forces the synchronization, resulting in a smooth final visual rendering. In our implementation, we set $y_{camera}^{max} = 15cm$.

Discussion

To summarize, our approach is composed of (1) an input interface based on the sole user's head movements, (2) a locomotion simulation in the VE composed of various possibilities such as jumping, crawling, turning, and (3) a visual feedback of walking relying on oscillating camera motions. The head motions are tracked along 3 Degrees of Freedom (DOF): lateral, vertical and roll axis (Figure **??**). These different physical motions are transposed in virtual movements thanks to a locomotion automaton (Figure **??**). We then added oscillating camera motions (Equations 2.9, 2.12 and 2.13) to the visual feedback to enhance the walking sensation. The different control laws were parametrized after preliminary testings. But of course some parameters can be modified in order to amplify/decrease

some effects during the locomotion simulation. Besides, other movement possibilities could also be envisaged and added to our automaton such as running state or backward movement.

2.4.2 Evaluation

The evaluation of the proposed technique was performed using a comparison with classical techniques in Desktop VR. We chose keyboard and joystick peripherals for seated and standing positions respectively as they are often used in Desktop VR context. The experiments were conducted using 3D VE displayed on a screen and we investigated the effectiveness of our technique to travel complex paths composed of different gates placed in the VE.

In this paper, we choose to not compare our technique to existing WIP techniques, and instead used more common interfaces that followed our low cost requirement. We conducted the evaluation of the proposed technique in both immersive Standing Up (SU) position and Desktop Sitting Down (SD) position. The keyboard and joystick were chosen respectively as the control conditions in the SD position and the SU position. Both keyboard and joystick peripherals are referred as Control techniques (Ctrl) in the following paragraphs. Our technique is referred with the "WIP" suffix.

Results

Task Completion Time For each participant, the task completion time of each trial was measured for the different experimental conditions. An exploratory analysis was first performed. A Principle Component Analysis revealed the presence of one spurious individual who has been taken out from the analysis. A specific analysis was developed to study the learning effect of the two conditions (Joystick/Keyboard and WIP techniques). A linear model where all conditions are mixed was fitted to explain the relation between the task completion times and the trial number. It revealed that the slope of the linear regression was significantly lower than zero (p-value < 0.000001), reflecting a significant decrease in the task completion time as the number of trials increases. The same analysis, where the first trial was removed, showed that the slope was not significantly different from zero anymore. In the following paragraph, the first trial was removed from the analysis as it corresponds to a learning effect.

A two-way ANOVA was performed on the 2 different conditions (Joystick/Keyboard and WIP techniques) and the 2 different positions (Sit-Down and Stand-Up). A post-hoc analysis using Tukey's procedure was then performed. The ANOVA was achieved separately for the two different types of paths (normal/steeple).

Concerning the normal path, the two-way ANOVA accounting for the conditions and the positions revealed a significant dependency between the position and the task completion time ($F(1,11) = 31.9981$, $p < 0.0001$) and between the condition and the task completion time ($F(1,11) = 6.0449$, $p = 0.0143$). Interaction between condition and position was also considered as a significant factor to discriminate task completion time ($F(1,11) = 27.1921$, $p < 0.0001$). Post-hoc analysis showed that the task completion time in the SD-WIP configuration ($M = 55.67s$) was significantly lower than in the SD-Ctrl configuration ($M = 61.72s$), adjusted p-value < 0.0001, in the SU-Ctrl configuration ($M = 62.06s$), adjusted p-value < 0.0001, and in the SU-WIP configuration ($M = 64.23s$), adjusted p-value < 0.0001. The other pairs of effects did not give any significant adjusted p-values.

The results concerning the different conditions are represented in Figure 2.5 for Sit-Down and Stand-Up experiments respectively, for the normal path only. The results are ordered in function of the trials. The first trial was kept to illustrate the learning effect.

For the steeple path, the pre-analysis suggested the presence of a spurious individual and the existence of one learning trial. The two-way ANOVA accounting for the conditions and positions revealed a significant dependency between the position and the task completion time ($F(1,11) = 8.5665$, $p < 0.005$), and the condition and the task completion time ($F(1,11) = 11.8925$, $p < 0.001$. Interaction between condition and position was also considered as a significant factor to discriminate task completion time ($F(1,11) = 8.7647$, $p < 0.005$). Post-hoc analysis showed than the task completion

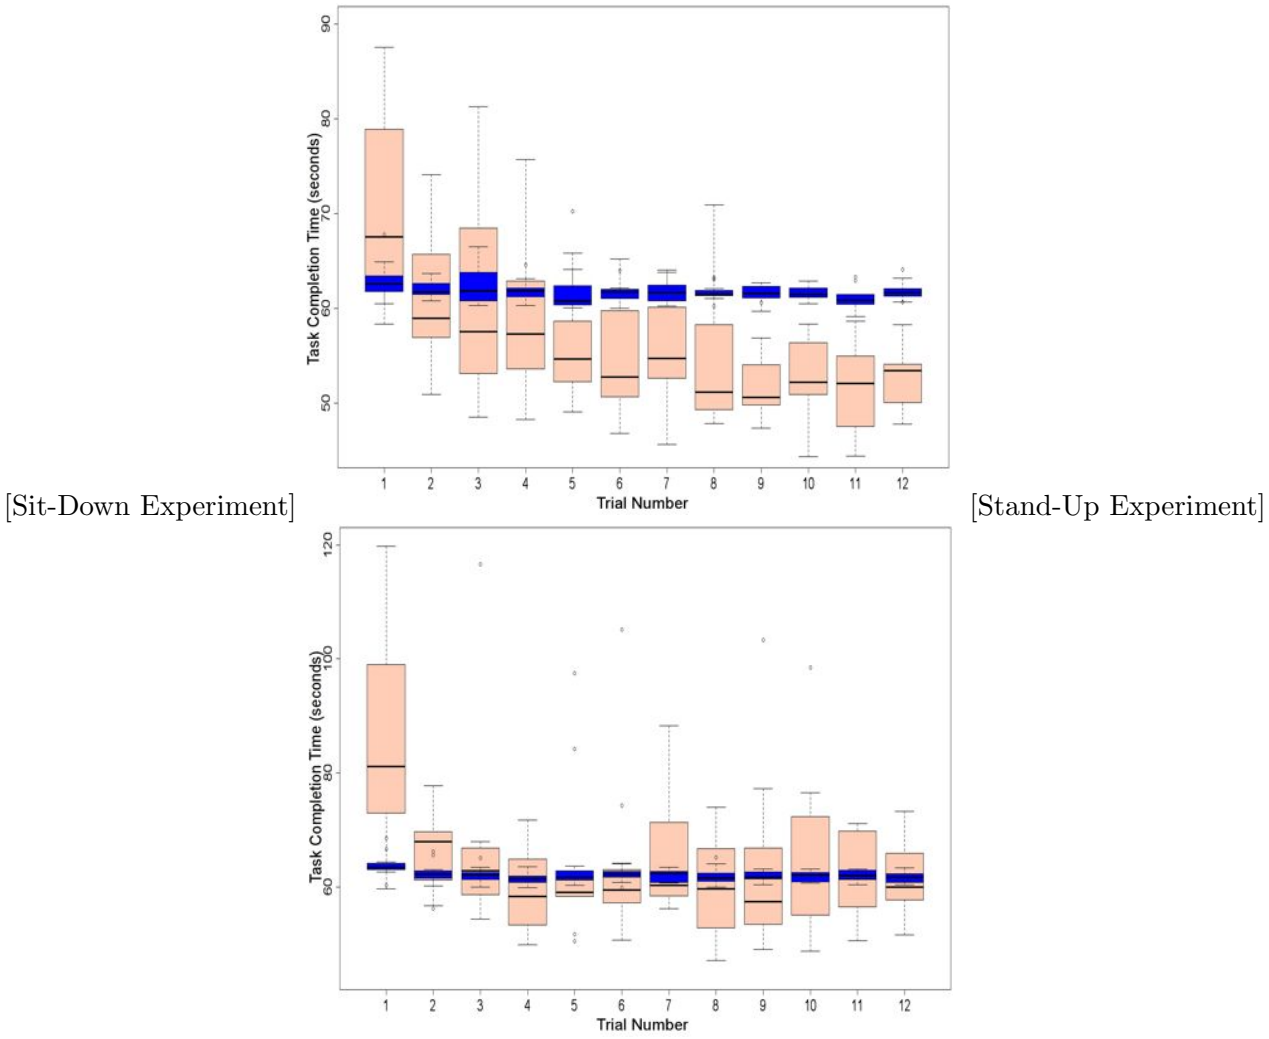

Figure 2.5: Task completion time for the two different techniques on normal paths for (a) SD experiments and (b) SU experiments. The blue and red light colors correspond to the Ctrl and WIP conditions respectively. Each boxplot is delimited by the quartile (25% quantile and 75% quantile) of the distribution of the condition over the individuals. The median is also represented for each condition.

time in the SU-WIP configuration ($M = 86.07s$) was significantly higher than in the SD-Ctrl configuration ($M = 74.78s$), adjusted p-value < 0.0001, in the SU-Ctrl configuration ($M = 74.72s$), adjusted p-values < 0.0001 and in the SD-WIP configuration ($M = 75.64s$), adjusted p-values < 0.0001. The other pairs of effects did not give any significant adjusted p-values.

Accuracy For each participant and for each trial, the percentage of errors for the different paths was measured. The two types of path (Normal and Steeple) are separated. The resulting percentages are for the normal path: 0.61% of error for SD-WIP configuration, 0% of error for SD-Ctrl configuration, 1.39% of error for SU-WIP configuration, 0.09% of error for SU-Ctrl configuration.

The resulting percentages are for the steeple path: 20.49% of error for SD-WIP configuration, 7.81% of error for SD-Ctrl configuration, 27.28% of error for SU-WIP configuration, 13.19% of error for SU-Ctrl configuration. We found a significant effect between Ctrl and WIP techniques for the steeple path.

Subjective Questionnaire After Sit-Down and Stand-Up configurations, a preference questionnaire was proposed in which participants had to grade from 1 (low appreciation) to 7 (high appreciation) the

four different conditions (SD-Ctrl, SD-WIP, SU-Ctrl, SU-WIP) according to 9 subjective criteria: (a) Fun, (b) Easiness of Use, (c) Intuitive, (d) Accuracy, (e) Presence, (f) Walking realism, (g) Fatigue, (h) Cybersickness and (i) Global appreciation. The results concerning the grades (Likert-scale) were obtained for the two different techniques for each of the subjective criteria, for the two experimental conditions (SD and SU). The grade 7 for Fatigue and Cybersickness respectively means that the technique does not induce any fatigue and does not imply any cybersickness feeling.

Concerning SD configuration, no significant effect was found for the following criteria: Intuitive ($p = 0.052$) and Cybersickness ($p = 0.12$). Concerning SU configuration, no significant effect was found for the following criteria: Intuitive ($p = 0.3$), Walking realism ($p = 0.19$) and Cybersickness ($p = 0.21$). We found a significant effect for all other criteria. In particular, our technique was better ranked for Fun, Presence and Global Appreciation, for both configurations.

2.4.3 Discussion

The results suggest that the technique can allow efficient navigation even compared with standard and well-known input devices such as keyboards and gamepads. The participants could sometimes go even faster with WIP, without any strong loss in precision. The WIP-based interaction seems also fast to learn, after only a couple of trials. The technique is well appreciated and perceived as more immersive and more fun than classical configurations.

The quick learning of our technique could be explained by the fact that interfaces based on webcam are generally intuitive and simple to learn [81]. After the learning phase, WIP tends to become faster than the keyboard in sitting condition. One explanation could be that with the technique (but also with the joystick in the standing condition) we could observe that participants tended to turn without stopping their advance motion. On the contrary, with the keyboard condition, participants tended to walk and turn sequentially, which might have globally increased the task completion time. Another explanation could be that, in the implementation of WIP, the advance speed is influenced by the speed of lateral movements. The seated position allows the user to make faster oscillations than the standing position, and thus to accelerate the walking motion by making fast lateral oscillations. Interestingly, these faster motions did not impair the precision of users.

The longer task completion time observed for WIP in the steeple paths (involving jumping and crawling motions) could be due to unexpected behaviors which induced incorrect transitions in our locomotion automaton. Indeed, some participants acted as if they were "anticipating" the jumps and bent down prior jumping. The problem could easily be fixed in the future by using additional conditions in the automaton based on both speed and position.

Results from the questionnaire are very consistent with previous subjective evaluations of WIP [111]. In the study, the WIP is more appreciated, and is perceived as more fun, and improving presence. As expected, more standard techniques (i.e. joystick and keyboard) are found easier to use, more precise, and less tiring (as they induce less physical movements). Interestingly, impression of cybersickness is not increased by WIP. This could be due to our desktop (and thus less immersive) configurations. Last, realism of walking in the VE was significantly improved only in the sitting condition. The perception of walking with WIP is actually quite complex, as participants wrote: "we have the impression to be a video game character", "the motions are exaggerated", or "we really have the sensation of walking, and not running". In the standing condition, some participants found that the physical motions were closer to "skiing" or "skating", as they noticed that they did not lift their feet from the ground but only oscillated their body. For these people we could further stress in the future that our implementation of WIP still works very well when lifting the feet and walking in place, as the oscillations of the head can be captured the same way in both situations (lifting the feet or not).

Taken together, the results suggest that the Shake-Your-Head technique could be used in a wide range of applications, when navigating in a 3D world, in sitting or standing configurations. It seems to be both a low-cost and an efficient paradigm that can match a lot of walking motions. It could thus be used for training in VR with more physical engagement (military infantry, vocational procedures), or more realistic virtual visits such as for project review in architecture or urban planning.

2.4.4 Conclusion

We presented the Shake-Your-Head, a technique designed to revisit the whole pipeline of the Walking-In-Place technique to match a larger set of configurations and apply it notably to the context of Desktop Virtual Reality. The Shake-Your-Head technique uses as sole input the head movements of the user. It can be used in a desktop configuration with the possibility for the user to sit down and to navigate in the VE through small screens and standard input devices such as a basic webcam for tracking. Various motions have been implemented such as turning, jumping and crawling in the locomotion simulation. Additional visual feedback based on camera motions to enhance the walking sensation were also introduced.

An experiment was conducted to evaluate the technique compared to standard techniques such as keyboard and joystick. In this experiment, participants had to walk through a series a gates forming a slalom path. The evaluation was performed both in an immersive and desktop configurations. It was notably found that WIP technique only requires a small learning time to allow faster navigation in seated position compared to the keyboard. Moreover, the technique was more appreciated and considered as more fun and inducing more presence than the other classical techniques.

2.5 Magic-Barrier-Tape: Walking in Large Virtual Environments With a Small Physical Workspace

There has been a large amount of previous work in the field of locomotion interfaces [31]. These devices allow users to be self-propulsed through a repetitive gait while staying in place by canceling the user's motion. Thus, they can provide kinesthetic feedback to the user. Some locomotion interfaces provide linear walking capabilities, such as treadmills and pedaling devices. Although the proprioceptive feedback matches real world movements, they are limited to a 1 degree of freedom (DOF) motion. Omni-directional walking is possible with 3DOF treadmills and foot platforms [31], together with some foot-wearable devices [43]. For the reasons mentioned in the introduction, namely their size and weight, their cost or their accuracy, they have not yet been widely adopted outside the laboratory.

Passive locomotion techniques were surveyed by Bowman et al. [7]. These travel techniques allow the user to navigate inside a virtual world without moving from his real-world position. Among rate control techniques, the most common and widely used technique is the flying vehicle metaphor [118], often coupled to an input device such as a wand through which the user controls his speed and orientation inside the virtual world. Using position control, the eyeball-in-hand and the scene-in-hand [7] techniques map the camera and the scene respectively to the user's hand. A clutching mechanism is required to navigate beyond one's arm length. The World-in-Miniature technique [76] provides a hand-held miniature of the scene, through which the user can select the location he wants to navigate to, and then be taken to that location in the real size virtual environment. Although usually intuitive and efficient, these navigation techniques are often inadequate for real world simulation scenarios, since their metaphors do not match real world navigation.

Active locomotion techniques, oriented towards real walking, provide natural metaphors by adapting real walking to restricted size workspaces. Moreover, the vestibular and proprioceptive feedbacks produced by self-propulsion increase the realism of the techniques and therefore the degree of immersion of the simulation. The most basic example is the Walking-in-Place technique [95], where the user simulates the physical act of walking by stepping in place but without forward motion of the body, making a gesture that is interpreted as a virtual forward, backward or side-step motion. However, although the technique has the advantage of not having to deal with workspace limits, it falls short regarding kinesthetic feedback. By scaling the user's speed along his intended direction of travel, the Seven League Boots [37] technique implements a scaling technique for real walking in order to reach virtual places beyond the real world workspace. The technique does not solve the limited workspace problem, and precise navigation can become very difficult. The Step WIM [51] takes a different approach on scaling by adapting the World-In-Miniature technique to real walking. The miniature world is drawn on the floor, and the user walks to the new destination on the miniature, instead of using his hands. Although the technique involves real walking, the metaphor might not be considered as natural. Resetting techniques such as the Freeze-backup [121] use a clutching approach by freezing the scene while the user recenters his position in the real world once he has reached the limits of the workspace. Other resetting techniques such as the 2:1-Turn [121] map a 360° virtual rotation to a 180° real rotation to keep walking on the same direction in the virtual world while taking the opposite direction in the real one. These resetting techniques are performed consciously by the user following a warning signal, which implies a break of immersion. Moreover, the resetting itself might feel unnatural: there is a sudden change in locomotion direction and orientation that does not correspond to the natural movement.

Redirected Walking [82] and Motion Compression [64] techniques solve many of the aforementioned problems by forcing the user into walking in a curved path in the real world when walking in a straight line in the virtual world through the progressive rotation of the scene around him. In a sufficiently large workspace, and with a straight virtual path, the user can walk endlessly without reaching the limits of the real workspace. These techniques are natural and in some cases imperceptible. However, they require large workspaces, can be confusing when doing unpredictable or quick changes of direction, and may require distracting events. In practice, they are more suited for Head Mounted Displays and

wide area tracking systems.

2.6 The Magic Barrier Tape

We propose a novel interaction metaphor, the Magic Barrier Tape, that brings a solution to immersive infinite walking in a restricted workspace through a natural and efficient metaphor.

Walking workspaces of virtual reality systems are often bounded by the tracking area, the display devices or by the walls of the immersive room. Hence, the Magic Barrier Tape has two fundamental objectives. The first one is to inform and display the limits of the workspace in a natural way, without break of immersion, in order to avoid the collision with physical objects outside the workspace boundaries or leaving the tracking area. The second one is to provide an integrated navigation technique to reach any location in the virtual scene, beyond the walking workspace.

To overcome the mismatch between the restricted size workspace and the potentially infinite size of the virtual scene, we followed the concept of hybrid position/rate control [20], used in a different context for object manipulation, where position control is used inside the available workspace for fine positioning, while rate control is used at the boundaries for coarse positioning. This concept can be found in common desktop applications and games where the mouse switches to rate control when it reaches the edge of the screen: in a file manager when doing multiple selection, or in top-view strategy games such as Starcraft when panning on the map. In our context, we applied the concept to navigation, with the available workspace being the walking workspace. The boundaries of the workspace are represented by a virtual barrier at mid body height textured with slanted black and yellow stripes, evoking the use of barrier tape and its implicit message: "do not cross".

The real workspace, delimited by the physical boundaries, is mapped to a virtual workspace inside the scene, delimited by the virtual barrier tape. Inside the workspace, we use position control: the user can freely walk, and objects inside the virtual workspace can be reached and manipulated through real walking and real life movements. When reaching the boundaries of the workspace, we switch to rate control: the user can move farther in the scene by "pushing" on the virtual barrier tape, hence translating the virtual workspace in the scene. He can then perform a task inside the virtual workspace at the new location.

The Magic Barrier Tape concept is not subject to a specific technology. It can be implemented in many different virtual reality systems. Any object or body part can be used as an actuator for the virtual barrier tape, depending on the application, and the rate control law can be fitted to specific behavioral needs. In the remaining of this section, we detail the Magic Barrier Tape concept. We take as implementation example our own virtual reality environment, consisting of a Head Mounted Display (HMD) with a 1.5m radius cylindrical tracking space, and one of the user's hands as actuating object.

2.6.1 Display of the Workspace Limits

The boundaries of the workspace are displayed through 3 complementary visual cues: the main virtual barrier tape, the warning virtual barrier tape, and their grey shadow on the floor.

The main virtual barrier tape is presented as a band that matches the shape of the workspace boundaries, such as a square for a CAVE or a circle for a cylindrical tracking system. It is positioned at a safe distance ahead of them, high enough from the virtual floor so that the user does not need to look down to see the barrier tape, and low enough so that it does not occlude the user's forward vision. The boundaries of the workspace are therefore clearly and continuously visible. The tape is made slightly translucent so what would have been normally hidden by the tape is still discernible.

The warning virtual barrier tape appears when the user's body is close to the main tape, as a warning signal. This second tape has the same shape and origin than the main one, and has a red glow to capture the user's attention. For the same reason, it is positioned at the user's eyes height. The tape is fully transparent when the user is at a reasonably safe distance from the main tape, and becomes progressively opaque as the user gets closer, therefore making the warning signal also progressive,

from dim to strong. The warning virtual barrier tape is complementary to the main tape, since it is triggered as a safety measure, and it gives an idea of when to stop walking and start "pushing".

The tapes shadow is drawn on the floor as if the barrier tapes were lit from above, in order to have a visual cue about the limits of the workspace when the user looks down. Hence, at least one of the 3 visual components is always visible at almost any viewing direction, which is particularly helpful with an HMD setup where there is usually a narrow field of view. Figure 2.6 shows the three components of the Magic Barrier Tape: the main barrier tape, the warning barrier tape (here visible) and the tapes' shadow.

In our virtual reality environment implementation, the main virtual barrier tape is 30 cm high and at 30 cm from the boundaries. It is shaped as a ring with a 1.2m radius and the center of the tracking area as origin. It is positioned at 1.3m from the virtual floor. The warning tape is activated when the user is at 30 cm from the main tape.

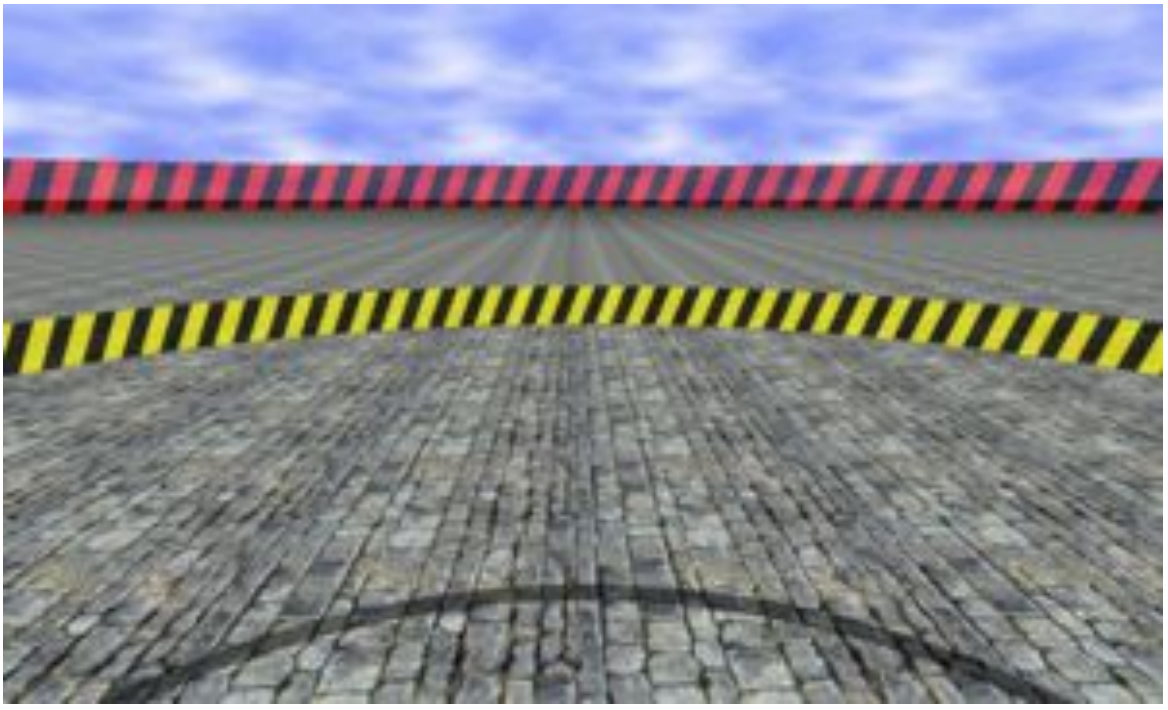

Figure 2.6: The three Magic Barrier Tape visual cues to show the workspace boundaries: the main virtual barrier tape (middle), the warning tape (top) and the tapes shadow (bottom).

2.6.2 Navigation Through Rate Control

The Magic Barrier Tape allows the use of position control inside the workspace, and rate control at the boundaries. The user is switched from position control to rate control whenever his hand (or any other tracked body part) penetrates the boundaries represented by the virtual barrier tape. The speed of the resulting translation in the virtual scene is a function of the hand penetration distance. When the user's hand is pulled back inside the workspace, the user is switched back to position control.

The virtual barrier tapes (main and warning) are deformed when the user's body (preferentially, the hand) penetrates the boundaries. This elastic behavior allows the user to see how deep he is "pushing", and therefore to evaluate how fast he will move in the virtual scene. A visual feedback on the rate control is also important so the user can know where the neutral position is located [20].

The deformation follows the shape of a centered Gaussian curve D, of equation:

$$D(p) = p\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{x^2}{2\sigma^2}}$$

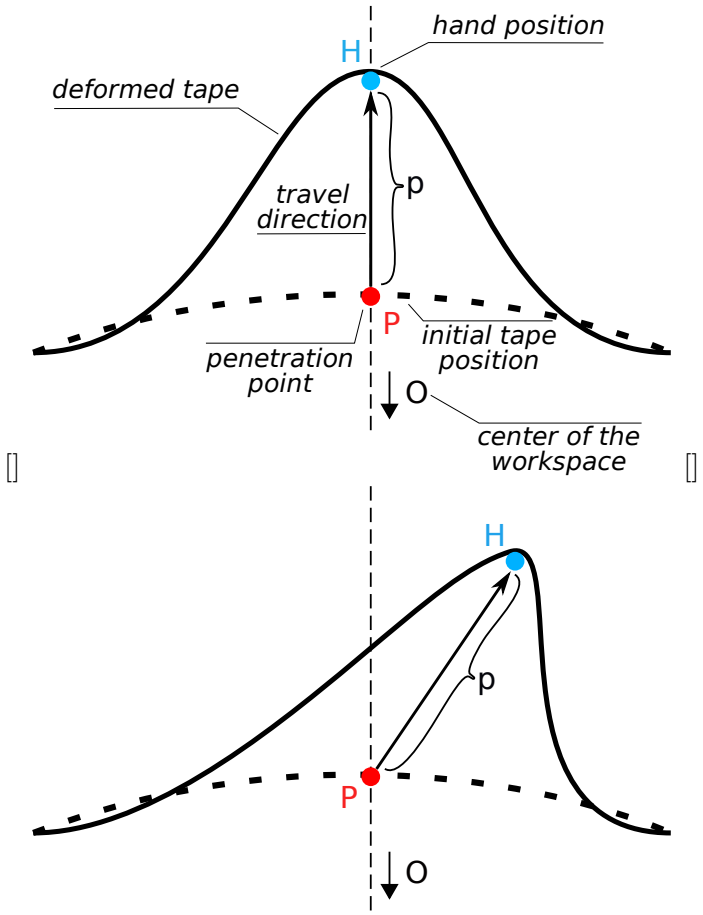

Figure 2.7: The Gaussian deformation of the Magic Barrier Tape in top-view (a) and its shifted version (b) to follow the hand position.

where p is the penetration length (in meters), and σ the standard deviation, which controls the "width" of the deformation. The virtual barrier tape is rotated so that the center of the Gaussian curve matches the penetration point P, the collision point between the hand and the virtual barrier tape. Therefore, as shown in Figure 2.7, the Gaussian deformation is centered around the penetration point, and its symmetry axis is given by the $\overrightarrow{OP}$ direction, where O is the center of the virtual barrier tape. Since the deformation follows the user's hand, the Gaussian curve has to be shifted to take into account the lateral deviation of the hand position H with respect to the Gaussian axis, as shown in Figure 2.7. The final result gives the impression of having an elastic region around the penetration point than can be deformed in any direction. This deformation direction, $\overrightarrow{PH}$, gives the travel direction of the virtual workspace (Figure 2.7).

The velocity V, a function of p, gives the speed of travel. It has the following equation:

$$V(p) = k * p^n$$

where k and n are constants. We use a polynomial function in order to have both slow speed when the user is close to the boundaries for small distances, and high speeds to move fast for distant targets. In our implementation, after preliminary testing, we used $\sigma = 0.15$, $k = 1.4$ and $n = 3$.

Our Magic Barrier Tape implementation provides both a safe walking environment and a natural and efficient navigation technique.

2.6.3 Extending Resetting Techniques for Omni-Directional Walking

In order to conduct the evaluation of the Magic Barrier Tape, we chose similar purpose techniques among existing navigation techniques. Among the surveyed active navigation techniques, based on real walking, only the resetting techniques developped by Williams et al. [121] provide collision free and infinite navigation capabilities. However, the resetting techniques were originally designed for straight paths and right angle turns, where in most virtual reality applications the user is allowed to freely explore his surrounding virtual environment, taking arbitrary paths and freely rotating around him. For fair comparison throughout the evaluation, since our Magic Barrier Tape technique enables such a navigation, we propose to add visual cues to these techniques in order to make them well suited for omni-directional navigation.

Extended Freeze-Backup Technique.

In the original Freeze-backup technique, in order to reset his position the user has to walk backwards in a straight line, until he reaches the resetting position. Since he is not guided while walking backwards, paths can only be straight. Otherwise, he could reach the workspace boundaries prematurely and find himself "locked" in a very short path resetting loop.

In the extended Freeze-backup technique, backups now need to take the user to the center of the real workspace. Before the reset, the user can be at any position in the real workspace, and with any orientation. Hence, we propose to add visual cues to guide the user through his resetting motion, which is divided in two steps. First, the body needs to be oriented towards the resetting position. An horizontal segment is drawn on the screen representing the user's orientation with respect to the resetting position, like his shoulder line seen from above in the real workspace reference frame. The user has to change his orientation until the segment becomes parallel to his body. Then, as a second step, the user has to walk to the resetting position by following an arrow direction. The arrow becomes smaller as the user gets closer to the resetting position, indicating how far he is from his target. Through this mechanism, the user can reach the center of the real workspace from anywhere in the real workspace. Figure **??** shows the segment and the arrow drawn at the top of the screen.

Extended 2:1-Turn Technique.

In the original 2:1-Turn technique, a 180° real rotation of the user is mapped to a 360° virtual turn, and the user stays on the same real path but on the opposite direction. Since real turns are always of 180°, walking paths need to be straight with eventually right angle turns to avoid the same "locking" problems mentioned above.

In the extended 2:1-Turn technique, real turns can no longer be of only 180°. The resetting angle is given by the non oriented angle between the viewing direction and the body position - resetting position vector. The virtual angle remains the same, 360°. For any orientation before resetting, two turning directions are possible: to the left and to the right. To each direction corresponds an angle, with usually one greater than the other. The direction with the largest angle is chosen, so that the rotation gain when mapping the turn to a 360° virtual turn is lower, and the illusion is therefore less perceivable. An arrow drawn at the top of the screen indicates the turning direction to the user.

2.6.4 Evaluation

In order to demonstrate its suitability for infinite navigation within a restricted workspace, we evaluated the Magic Barrier Tape by comparing it to two other existing navigation techniques that enable collision free infinite walking within a restricted workspace, namely the Freeze-Backup and the 2:1-Turn resetting techniques [121] with our extensions for omni-directional walking. We conducted two experiments, a pointing task and a path following task.

Experiment #1

In Experiment #1, our goal was to compare the 3 techniques over a pointing task where the user had to move from a central initial location to a new location, indicated by a target, as fast as possible. We a priori assumed that the Magic Barrier Tape will be faster, since rate control allows speeds greater than the average walking speed.

For the different comparison analyses, a correction for experiment-wise error was realized by using Bonferroni-adjusted alpha level ($p = 0.05$ divided by the number of tests). Thus, in order to compare the Barrier Tape technique to the two other techniques (Freeze-backup and 2:1-Turn) the alpha level was adjusted to $p = 0.025$.

Completion Time. Using the completion time data collected during the experiment, we conducted a statistical analysis. For each participant, statistics (mean M, standard deviation SD) were computed on the 18 trials in each condition. A one-way within subject design ANOVA (Techniques: Barrier Tape, Freeze-backup, 2:1-Turn) on the mean completion time (in seconds) revealed a significant main effect of the technique ($F(2, 22) = 183.22$, $p < 0.001$). Follow up t tests revealed that completion time in the Barrier Tape technique ($M = 6.37$ sec, $SD = 1.30$ sec) was significantly shorter than in the Freeze-backup technique ($M = 21.49$ sec, $SD = 3.11$ sec, $t(11) = -19.15$, $p < 0.001$). Similarly, completion time in the Barrier Tape technique was significantly shorter than in the 2:1-Turn technique ($M = 14.54$ sec, $SD = 2.41$ sec, $t(11) = -14.61$, $p < 0.001$).

Amplitude of Walking in the Real World. An ANOVA on the mean amplitude of walking in the real world (in meters) revealed a significant main effect of the technique ($F(2, 22) = 434.75$, $p < 0.001$). Follow up t tests revealed that the amplitude of walking in the real world in the Barrier Tape technique ($M = 1.46$ m, $SD = 0.16$ m) was significantly shorter than in the Freeze-backup technique ($M = 4.42$ m, $SD = 0.30$ m, $t(11) = -30.13$, $p < 0.001$). Similarly, the amplitude of walking in the real world in the Barrier Tape technique was significantly shorter than in the 2:1-Turn technique ($M = 3.37$ m, $SD = 0.23$ m, $t(11) = -20.80$, $p < 0.001$).

Experiment #2

In the second experiment, our goal was to compare the 3 techniques over a path following task where the user had to follow a path delimited by two virtual walls, as fast as possible and as accurately as possible by trying to stay right between the two walls. We a priori assumed that the Magic Barrier Tape will be faster, as in Experiment #1, but less precise due to the controlability of rate control [122].

For the different comparison analysis, a correction for experiment-wise error was realized by using Bonferroni-adjusted alpha level ($p = 0.05$ divided by the number of tests). Thus, in order to compare the Barrier Tape technique to the two other techniques (Freeze-backup and 2:1-Turn) the alpha level was adjusted to $p = 0.025$.

Completion Time. An ANOVA on the mean completion time (in seconds) revealed a significant main effect of the technique ($F(2, 22) = 84.01$, $p < 0.001$). Follow up t tests revealed that completion time in the Barrier Tape technique ($M = 31.62$ sec, $SD = 9.71$ sec) was significantly shorter than in the Freeze-backup technique ($M = 99.54$ sec, $SD = 21.63$ sec, $t(11) = -12.06$, $p < 0.001$). Similarly, completion time in the Barrier Tape technique was significantly shorter than in the 2:1-Turn technique ($M = 52.33$ sec, $SD = 6.59$ sec, $t(11) = -6.48$, $p < 0.001$).

Path Deviation. An ANOVA on the mean path deviation (in square meters) revealed a significant main effect of the technique $F(2, 22) = 4.77$, $p = 0.019$ (Figure **??**). Follow up t tests revealed that the path deviation in the Barrier Tape technique ($M = 3.46$ m^2, $SD = 1.76$ m^2) was not significantly different from the path deviation in the Freeze-backup technique ($M = 2.45$ m^2, $SD = 1.04$ m^2,

$t(11) = 1.72$, $p = 0.1143$.). By contrast, the analysis indicated that the path deviation in the 2:1-Turn technique ($M = 1.93$ m^2, $SD = 0.54$ m^2) was significantly lower than in the Barrier Tape technique, $t(11) = 2.81$, $p = 0.017$.

Amplitude of Walking in the Real World. An ANOVA on the mean amplitude of walking in the real world (in meters) revealed a significant main effect of the technique ($F(2, 22) = 379.81$, $p < 0.001$). Follow up t tests revealed that the amplitude of walking in the real world in the Barrier Tape technique ($M = 6.81$ m, $SD = 1.33$ m) was significantly shorter than in the Freeze-backup technique ($M = 19.03$ m, $SD = 1.25$ m, $t(11) = -32.63$, $p < 0.001$). Similarly, the amplitude of walking in the real world in the Barrier Tape technique was significantly shorter than in the 2:1-Turn technique ($M = 13.61$ m, $SD = 1.54$ m, $t(11) = -13.17$, $p < 0.001$).

Subjective Questionnaire

After both experiments, a preference questionnaire was proposed in which participants had to grade from 1 to 7 the 3 techniques according to 6 subjective criteria: *easiness of use*, *fatigue*, *navigation speed*, *navigation precision*, *general appreciation* and *naturalness*. Figure 2.8 shows the means and standard deviations of the 3 techniques for each of the subjective criteria.

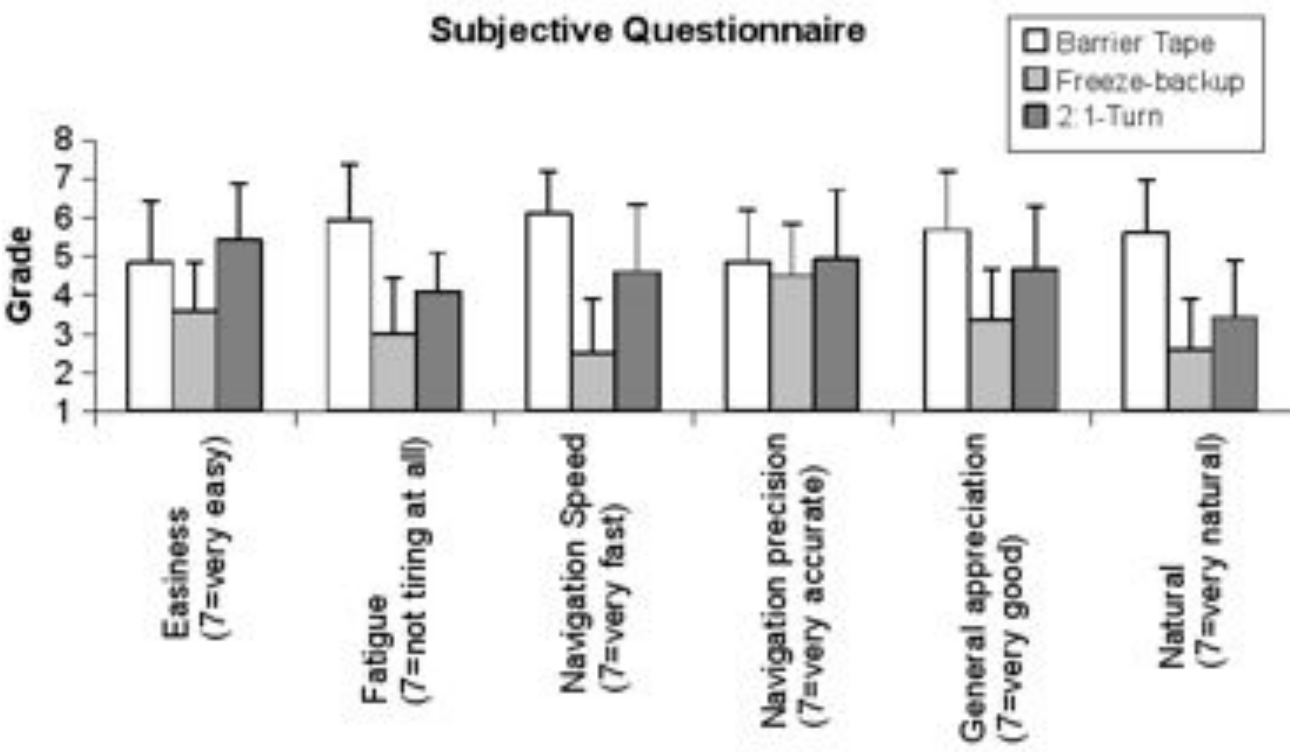

Figure 2.8: Mean and standard deviation of subjective ratings about the different criteria for the three techniques.

Wilcoxon signed rank tests with Bonferroni correction showed significant differences: for the *fatigue*, between the Barrier Tape and the Freeze-backup techniques ($z = 2.69$, $p = 0.007$) and between the Barrier Tape and the 2:1-Turn techniques ($z = 2.41$, $p = 0.016$); for the *naturalness*, between the Barrier Tape and the Freeze-backup techniques ($z = 2.77$, $p = 0.006$) and between the Barrier Tape and the 2:1-Turn techniques ($z = 2.53$, $p = 0.011$); for the *navigation speed*, only between the Barrier Tape and the Freeze-backup techniques ($z = 2.82$, $p = 0.005$); and for the *general appreciation*, only between the Barrier Tape and the Freeze-backup techniques ($z = 2.65$, $p = 0.008$).

2.6.5 General Discussion

Both experiments showed that the Magic Barrier Tape is faster compared to the other techniques. Indeed, results show that Experiment #1 trials were completed more than 3 times faster with the Magic Barrier Tape than with the Freeze-backup technique, and more than 2 times faster than with the 2:1-Turn technique. In Experiment #2, completion time using the Magic Barrier Tape was also roughly 3 and 2 times faster respectively. This result is consistent with the user's impression from the questionnaire regarding the navigation speed of the different techniques (Figure 2.8). It is mainly due to the fact that there is no time lost in the resetting of the position when using the Magic Barrier Tape, and that the control law allows navigation speeds greater than the average walking speed. Completion times could be further reduced by tuning the control law for greater speeds, although controlling the Magic Barrier Tape could become increasingly difficult, as testified by 3 users which complained about an acceleration behavior that was sometimes hard to control.

The experiments also showed that users walked less when using the Magic Barrier Tape than with the other 2 techniques, which was expected due to the use of rate control at the boundaries of the workspace. However, an interesting observation can be made when considering that trials were completed significantly faster with the Magic Barrier Tape. If we do a ratio between the amplitude of walking in the real world and completion time, in a per user basis, we obtain similar values for the Magic Barrier Tape ($M = 0.24$, $SD = 0.04$), the Freeze-backup ($M = 0.21$, $SD = 0.04$) and the 2:1-Turn ($M = 0.24$, $SD = 0.04$) techniques in Experiment #1, as well as in Experiment #2 with ($M = 0.22$, $SD = 0.036$), ($M = 0.20$, $SD = 0.037$) and ($M = 0.26$, $SD = 0.052$) respectively. Hence, the amount of "useful walking", contributing to moving forward in the virtual world, relative to time is as large with the Magic Barrier Tape as with the other techniques. If we consider that walking speeds are the same for the 3 techniques, users spend roughly the same percentage of the total time doing useful walking with the Magic Barrier Tape technique than with the other 2 techniques.

Experiment #2 showed that the Magic Barrier Tape was less precise when following a given path, with a higher path deviation when compared to the 2:1-Turn technique (roughly 2 times less precise). We cannot conclude on the comparison with the Freeze-backup technique, since results were not significantly different. Again, these results were expected. By nature and design, the use of the Magic Barrier Tape is meant for coarse positioning. The user gets close enough to the navigation target in order to have it inside his workspace, and can then reach it by fine positioning navigation through real walking. As explained by one of the subjects of the Experiment #2, when asked about the strategies he used: "I sent the barrier tape as far as possible without going into the walls in order to take advantage of the workspace". However, path deviation could be improved by allowing users to cutomize their control law, like when they choose the mouse speed in desktop computers. Moreover, Zhai [122] observed that, with sufficient training, rate control and position control can achieve similar performances. Hence, further user training on the Magic Barrier Tape rate control might improve its mean path deviation.

Overall, users graded the Magic Barrier Tape higher in all criteria of the questionnaire where comparisons were significantly different. We highlight that 6 subjects complained about having cybersickness when using the 2:1-Turn technique, while 2 said it made them loose balance. Many subjects found the Freeze-backup technique exhausting and frustrating. It is also important to note that 2 subjects had a very hard time using the Magic Barrier Tape. They used an inadequate strategy, and complained about the control law. They might have needed a longer training time, or more guidance on the strategy to adopt. They consistently graded it lower than the other techniques in every criteria of the questionnaire.

In a nutshell, the Magic Barrier Tape is faster than the Freeze-backup and the 2:1-Turn techniques, but is less precise when using it in rate control. The 2:1-Turn technique is the most precise, but seems to induce cybersickness to users, as well as stability issues. There is a general dissatisfaction with the Freeze-backup technique, mainly due to its physical exertion and slow speed, leading to a frustrating experience. People generally prefer the Magic Barrier Tape, and find it more natural and less tiring.

2.6.6 Perspectives

Through user feedback on the experiments and our own observations, we found some ways of potentially improving the Magic Barrier Tape.

The user could use any part of his body in order to "push" on the virtual barrier tape. Figure **??** (right) shows the virtual barrier tape being triggered by using the elbow when the user's hands are busy carrying an object. One could think about using the shoulders, the pelvis or the feet, since we often naturally use these body parts when we are unable to use our hands.

In their "Bubble" technique [20], a hybrid position/rate control haptic interaction technique for devices with restricted workspace, Dominjon et al. successfully used haptic feedback to represent the workspace boundaries and their virtual elasticity. In the RubberEdge technique [11], Casiez et al. used a passive haptic feedback through an elastic ring on top of a tracking surface such as a touchpad to allow the user to switch from position to rate control when reaching the elastic boundaries. Similarly, the Magic Barrier Tape could be augmented with haptic feedback when "pushing" on the tape. A possibility could be the use of passive haptics through tangible objects such as queue barriers with retractable belts as one could find in airports and queue-up places. The queue barriers could follow the workspace boundaries, and the virtual barrier tape would match the queue barriers position. Since retractable belts are elastic, the haptic feedback of the virtual barrier tape elastic deformation would be straightforward. Many users complained about the translation speed when using the Magic Barrier Tape, since the acceleration could be hard to control. A solution to this problem could be the use of a discrete approach. The control law, according to the hand penetration, would deliver one of three discrete velocities, corresponding to a human walking, jogging or running. Side-stepping human velocities could be used when moving in a direction orthogonal to the body orientation. The translation speed would therefore be more predictable, although capped by the running speed.

Last, although the barrier tape is made semi-transparent to reduce occlusion, visibility might be reduced in environments were the dominant color is close to the tape color. A way to enhance visibility in such cases would be to use different tape textures using the complementaries of the dominant colors of the surrounding environment, in order to emphasize the presence of the Magic Barrier Tape while increasing the visibility of the scene.

2.6.7 Conclusion

This paper introduces the Magic Barrier Tape, a new interaction metaphor for navigating in a potentially infinite virtual scene while confined to a restricted walking workspace. Using the barrier tape metaphor and its "do not cross" implicit message, the walking workspace is surrounded with virtual barrier tape in the virtual scene. The technique uses a hybrid position/rate control mechanism: real walking is used inside the workspace, while rate control navigation is used to move beyond the boundaries by "pushing" on the virtual barrier tape. Moreover, the technique naturally informs the user about the boundaries of his walking workspace, providing a walking environment safe from collisions and tracking problems.

We conducted two experiments in order to evaluate the Magic Barrier Tape by comparing it to other state-of-the-art navigation techniques previously extended for omni-directional navigation. In Experiment #1 participants had to walk to a target, while in Experiment #2 they had to navigate inside a scene following a path. Results showed that the Magic Barrier Tape was faster than the other techniques. Experiment #2 results confirmed that, by design, navigation through rate control with the Magic Barrier Tape is not meant for precise path following, but rather for coarse positioning between fine positioning tasks. Overall, the Magic Barrier Tape was more appreciated, while being more natural and less tiring.

Future work will focus on exploring the different perspectives highlighted in this paper, namely the use of haptics for a more compelling and immersive experience, and the use of discrete velocities to produce a more predictable and realistic motion in rate control.

3 A multimodal architecture for natural interactive walking

3.1 Introduction

In the academic community, foot-based interactions have mostly been concerned with the engineering of locomotion interfaces for virtual environments [79].

A notable exception is the work of Paradiso and coworkers, who pioneered the development of shoes enhanced with sensors, able to capture 16 different parameters such as pressure, orientation, acceleration [74]. Such shoes were used for entertainment purpose as well as for rehabilitation studies [4]. The company Nike has also developed an accelerometer which can be attached to running shoes and connected to an iPod, in such a way that, when a person runs, the iPod tracks and reports different information.

In this paper we mostly focus on enhancing our awareness of auditory and haptic feedback in foot based devices, topic which is still rather unexplored. Virtual augmented footwear has interesting applications in different fields related to virtual reality. As an example, auditory and haptic feedback in foot-based interaction can assist rehabilitation. Moreover, foot-based interfaces has started to appear in the entertainment industry, in the form of platforms such as the Wii fit from Nintendo, which is connected to the Wii console[1]. Furthermore, having the possibility to provide auditory and haptic feedback has the potential of providing interesting applications in the field of navigation, especially for visually impaired people. When exploring a place by walking, two main categories of sounds can be identified: the person's own footsteps and the surrounding soundscapes. In the movie industry, footsteps sounds represent important elements. Chion writes of footstep sounds as being rich in what he refers to as materializing sound indices – those features that can lend concreteness and materiality to what is on-screen, or contrarily, make it seem abstracted and unreal [12]. We believe that footsteps sounds represent an important element also in interactive entertainment, and novel foot-based interactions present new possibilities in this area.

We describe a multimodal interactive space which has been developed with the goal of creating audio-haptic-visual simulations of walking-based interactions. The system requires users to walk around a space wearing a pair of shoes enhanced with sensors and actuators. The position of such shoes is tracked by a motion capture system, and the shoes drive a audio-visual-haptic synthesis engine based on physical models.

We have used this architecture to perform several psychophysical experiments in order to understand the contribution of the auditory and haptic modalities when interacting with different simulated surfaces using the feet [109]. We have also investigated the role of the different modalities when providing feedback in balancing tasks, as well as the possibility of recreating sense of presence in virtual environments [67].

Possible applications of the architecture are envisaged in the field of navigation in real and virtual environments, architecture, rehabilitation and entertainment. As an example, a better understanding of the role of the different modalities in helping balance control, can advance the field of virtual reality for rehabilitation purpose. Moreover, the possibility to have faithful reproduction of real places, both indoors and outdoors, allows to advance the field of virtual reality for architecture, as well as the ability to visit a physical place virtually. Moreover, in the entertainment industry, several interfaces such as the Wii Fit by Nintendo and the Kinect by Microsoft are starting to explore the possibilities

[1] www.nintendo.com

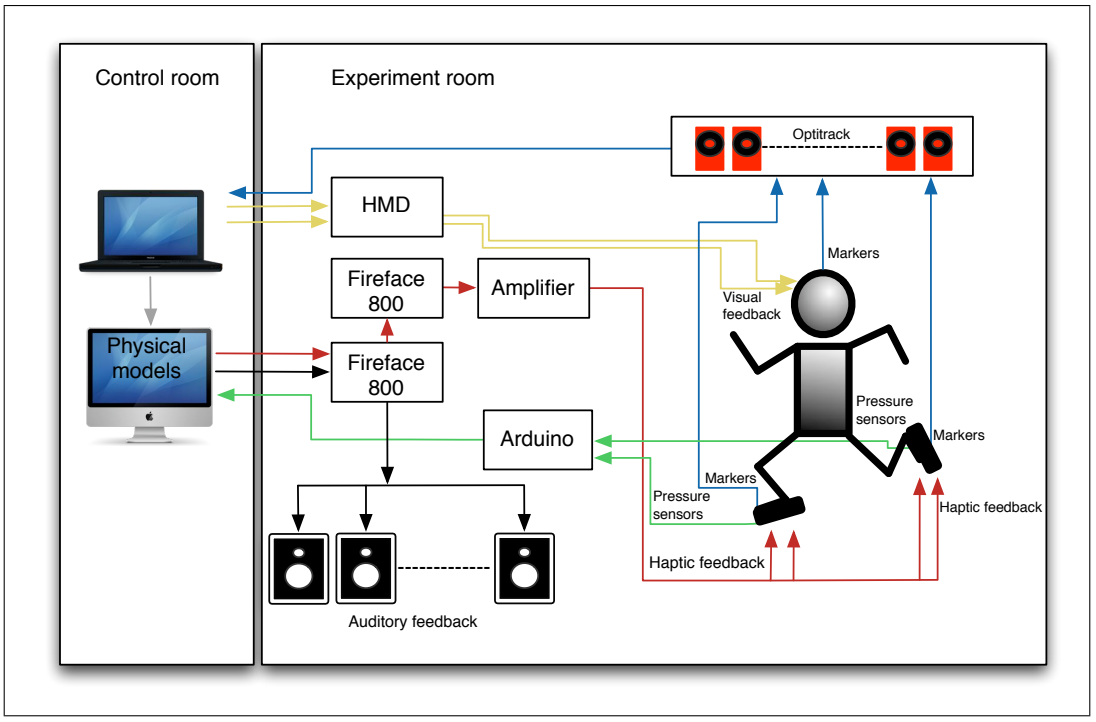

Figure 3.1: A schematic representation of the developed architecture to simulate natural interactive walking.

offered by feet-based and full-body interactions. Amusement parks are also exploring the possibilities offered by virtual reality and multimodal interaction in order to provide several illusions such as vection, i.e., the illusion of self-movement in space. The results presented in this paper are part of the Natural Interactive Walking (NIW) FET-Open EU project (http://www.niwproject.eu/), whose goal is to provide closed-loop interaction paradigms enabling the transfer of skills that have been previously learned in everyday tasks associated to walking. In the NIW project, several walking scenarios are simulated in a multimodal context, where especially audition and haptic play an important role.

The overall architecture

Figure 3.1 shows a schematic representation of the overall architecture developed. The architecture consists of a motion capture system (MoCap), two soundcards, twelve loudspeakers, two amplifiers, two haptic shoes, a head-mounted display (HMD), and two computers. Such system is placed in an acoustically isolated laboratory which consists of a control room and a bigger room where the setup is installed and where the experiments are performed. The control room is 5.45 m large, 2 m long, and 2.85 m high, and it is used by the experimenters providing the stimuli and collecting the experimental results. It hosts 2 desktop computers.
The first computer runs the motion capture software (Tracking Tools 2.0) and the visual engine, while the second runs the audio-haptic synthesis engine. The two computers are connected through an ethernet cable and communicate by means of the OSC protocol (http://opensoundcontrol.org/). The data relative to the motion capture system are sent from the first to the second computer which processes them in order to control the sound engine.

A transparent glass divides the two rooms, so it is possible for the experimenters to see the users performing the assigned task.
The two rooms are connected by means of a standard talkback system such as the ones used in recording studios.

The experiment room is 5.45 m large, 5.55 m long, and 2.85 m high, and the walking area available

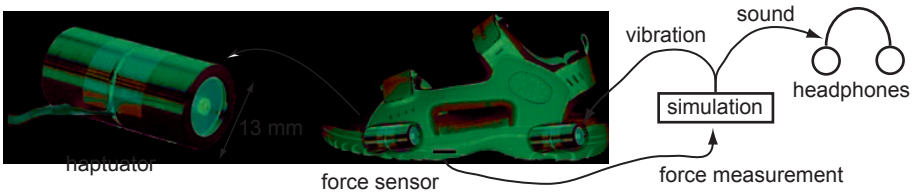

Figure 3.2: The developed haptic shoes used as part of the multimodal architecture

to the users is about $24m^2$.

Figure 3.3: A picture of one pressure sensor and two actuators embedded in the shoes.

SIMULATION HARDWARE

Tracking the user

The user locomotion is tracked by an Optitrack motion capture system (http://naturalpoint.com/optitrack/) composed by 16 infrared cameras (OptiTrack FLEX:V100R2). The cameras are placed in a configuration optimized for the tracking of the feet and head position simultaneously. In order to achieve this goal, markers are placed on the top of each shoe worn by the subjects as well as on top of the head (using the HMD for this purpose, or, in its absence, a cycling helmet).

Users are also tracked by using the pressure sensors embedded in a pair of sandals. Specifically, a pair of light-weight sandals was used (Model Arpenaz-50, Decathlon, Villeneuve d'Ascq, France).

The sole has two FSR pressure sensors (I.E.E. SS-U-N-S-00039) whose aim is to detect the pressure force of the feet during the locomotion of a subject wearing the shoes. The two sensors were placed in correspondence to the heel and toe respectively in each shoe. The analogue values of each of these sensors were digitalized by means of an Arduino Diecimila board (http://arduino.cc/) and were used to drive the audio and haptic synthesis.

Haptic feedback

In order to provide haptic feedback during the act of walking, a pair of custom made shoes with sensors and actuators has been recently developed [107]. The particular model of shoes chosen has light, stiff foam soles that are easy to gouge and fashion. Four cavities were made in the tickness of the sole to accommodate four vibrotactile actuators (Haptuator, Tactile Labs Inc., Deux-Montagnes, Qc, Canada). These electromagnetic recoil-type actuators have an operational, linear bandwidth of 50–500 Hz and can provide up to 3 G of acceleration when connected to light loads. As indicated in Fig. 3.2 and Fig. 3.3, two actuators were placed under the heel of the wearer and the other two under the ball of the foot. These were bonded in place to ensure good transmission of the vibrations

Figure 3.4: A screenshot of the audio-haptic synthesis engine used in the architecture.

inside the soles. When activated, vibrations propagated far in the light, stiff foam. In the present configuration, the four actuators were driven by the same signal but could be activated separately to emphasize, for instance, the front or back activation, to strick a balance, or to realize other effects such as modulating different, back-front signals during heel-toe movements.

A cable exits from each shoe, with the function of transporting the signals of the pressure sensors and for the actuators. Such cables were about 5 meters long, and they were connected, through DB9 connectors, to two 4TP (twisted pair) cables: one 4TP cable carries the sensor signals to a breakout board[2], which then interfaces to an Arduino board; the other 4TP cable carries the actuator signals from a pair of Pyle Pro PCA1 (http://www.pyleaudio.com/manuals/PCA1.pdf) mini 2X15 W stereo amplifiers, driven by outputs from a FireFace 800 soundcard. (http://www.rme-audio.com/english/firewire/ff800.htm).

Each stereo amplifier handles 4 actuators found on a single shoe, each output channel of the amplifier driving two actuators connected in parallel. The PC handles the Arduino through a USB connection, and the FireFace soundcard through a FireWire connection.

Auditory feedback

In our virtual environment the sound can be delivered by means of headphones (Sennheiser HD 650, http://www.sennheiser.com) or a set of loudspeakers (Dynaudio BM5A speakers, http://www.dynaudioacoustics.co
In detail, the loudspeakers configuration is illustrated in Figure 3.5. In the current setup we use 12 or 16 loudspeakers depending on whether the haptic feedback is involved or not. Indeed for the delivery of both the haptic and audio feedback we use two FireFace 800 soundcard connected through a firewire 800 cable (see figure 3.1). Since there are 8 output channels available on each soundcard, and the handling of the haptic feedback requires 4 output channels, we use the remaining 12 for the auditory feedback (loudspeakers 1-12 in figure 3.5). Conversely, when the haptic feedback is not involved all the 16 channels are available for the auditory feedback (loudspeakers 1-16 in figure 3.5). In future we

[2]containing trimmers, that form voltage dividers with the FSRs

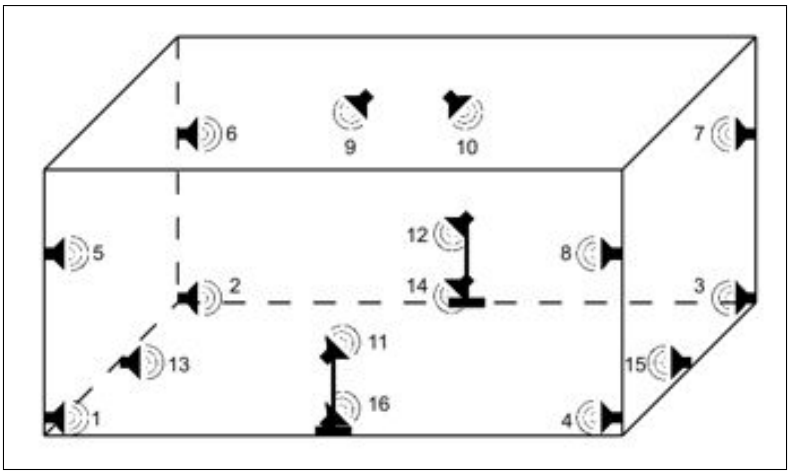

Figure 3.5: A schematic representation of the sound diffusion system.

plan to extend this configuration adding a third soundcard dedicated exclusively to the handling of the haptic feedback.

Visual feedback

The visual feedback is provided by an head-mounted display (HMD) nVisor SX from nVis (www.nvis.com). The HMD is connected to the PC by using a Matrox TripleHead2Go Digital Edition graphics expansion module. As previously mentioned, three markers are placed on top of the HMD, in order to track orientation and position of the head.

The goal of the visual feedback is to render, through the use of a commercial game engine, the visual sensation of exploring different landscapes. In particular, in our simulation the Unity3D game engine has been used (http://unity3d.com/). This engine was used for its ability to render realistic visual environments without being skilled visual designers. This choice was ideal for us, since our main interest is a physically based audio-haptic engine, so the visual feedback is used only for supporting it, without being the main goal.

SIMULATION SOFTWARE

We developed a multimodal synthesis engine able to reproduce auditory and haptic feedback As concerns the auditory feedback, we developed a sound synthesis engine based on a footsteps sounds synthesis engine and on a soundscapes synthesis engine.
As regards the haptic feedback, it is provided by means of the haptic shoes previously described. The haptic synthesis is driven by the same engine used for the synthesis of footsteps sounds, and is able to simulate the haptic sensation of walking on different surfaces, as illustrated in [107].
The engine for footsteps sounds, based on physical models, is able to render the sounds of footsteps both on solid and aggregate surfaces. Several different materials have been simulated, in particular wood, creaking wood, and metal as concerns the solid surfaces, and gravel, snow, sand, dirt, forest underbrush, dry leaves, and high grass as regards the aggregate surfaces. A complete description of such engine in terms of sound design, implementation and control systems is presented in [110].
Using such engine, we implemented a comprehensive collection of footstep sounds. As solid surfaces, we implemented metal, wood, and creaking wood. In these materials, the impact model was used to simulate the act of walking, while the friction model was used to simulate the creaking sounds typical of creaking wood floors. As aggregate surfaces, we implemented gravel, sand, snow, forest underbrush, dry leaves, pebbles and high grass. The simulated metal, wood and creaking wood surfaces were furthermore enhanced by using some reverberation.

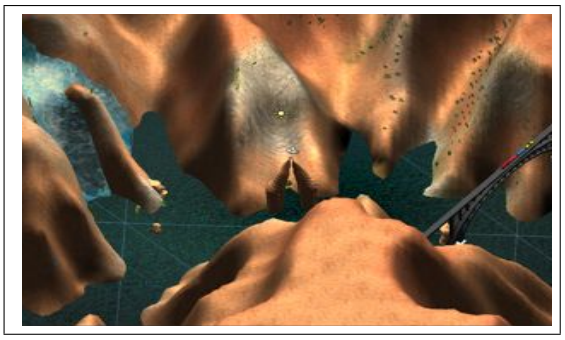

Figure 3.6: A screenshot of the visual feedback representing the Gran Canyon.

To control the audio-haptic footsteps synthesizer in our virtual environment, we use the haptic shoes: the audio-haptic synthesis is driven interactively during the locomotion of the subject wearing the shoes. The description of the control algorithms based on the analysis of the values of the pressure sensors coming from the haptic shoes can be found [107].
Such engine has been extensively tested by means of several audio and audio-haptic experiments and results can be found in [70] [68] [93] [108].

Soundscape rendering

We implemented a soundscape engine able to provide various typologies of soundscapes: static soundscapes, dynamic soundscapes and interactive soundscapes. Static soundscapes are those composed without rendering the appropriate spatial position of the sound sources, nor their tridimensional movements in the space. Conversely, in the dynamic soundscapes the spatial position of each sound source is taken into account, as well as their eventual movements along tridimensional trajectories. Finally, the interactive soundscapes are based on the dynamic ones where in addition the user can interact with the simulated environment generating an auditory feedback as result of his/her actions. The position and the movements of the user are tracked by means of the MoCap system and are used as input for the designed interactive soundscapes. As an example of sound interaction, one can imagine the situation in which the virtual environment simulates a forest, and when the user walks enough close to a bush where there are some animals the sounds of the movements of the animals fleeing are triggered. Furthermore, also the footsteps sounds interactively generated during the locomotion of the user can be conveyed to the user taking into account the position of the user feet in the simulated space, that is the footsteps sounds follow the user position.

The sound synthesis engine has been implemented in Max/MSP. To achieve the dynamism in the soundscapes we use the ambisonic tools for Max/MSP (Available at http://www.icst.net/research/projects/ambiso
tools/) which allows to move virtual sound sources along trajectories defined on a tridimensional space [89]. In particular the sound synthesis engine is currently set with 16 independent virtual sound sources, one for the footsteps sound, and fifteen for the sound sources present in the soundscape.

As regards to the visual feedback, several outdoor scenarios have been developed using the Unity3D engine. The goal of such outdoor scenarios is to provide a visual representation of the physically simulated surfaces provided in the audio-haptic engine. As an example, a forest, a beach and a skyslope were visually rendered, to match the physically simulated sand, forest underbrush and snow. The user interacts with the visual engine by means of the markers placed on the top of the HMD, and by means of the pressure sensors embedded in the shoes. Moreover, as shown in Figure 3.6 and 3.7, the Gran Canyon was simulated.

Figure 3.8 shows a subject interacting with the virtual environment. The subject is wearing the HMD and the shoes enhanced with actuators, pressure sensors and markers. In the background it is

Figure 3.7: A screenshot of the visual feedback representing a waterfall in the Gran Canyon.

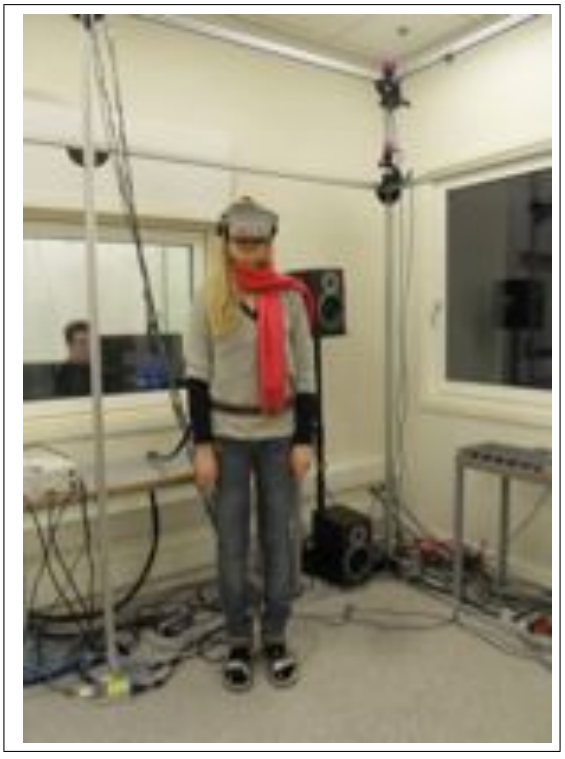

Figure 3.8: A subject interacting with the architecture

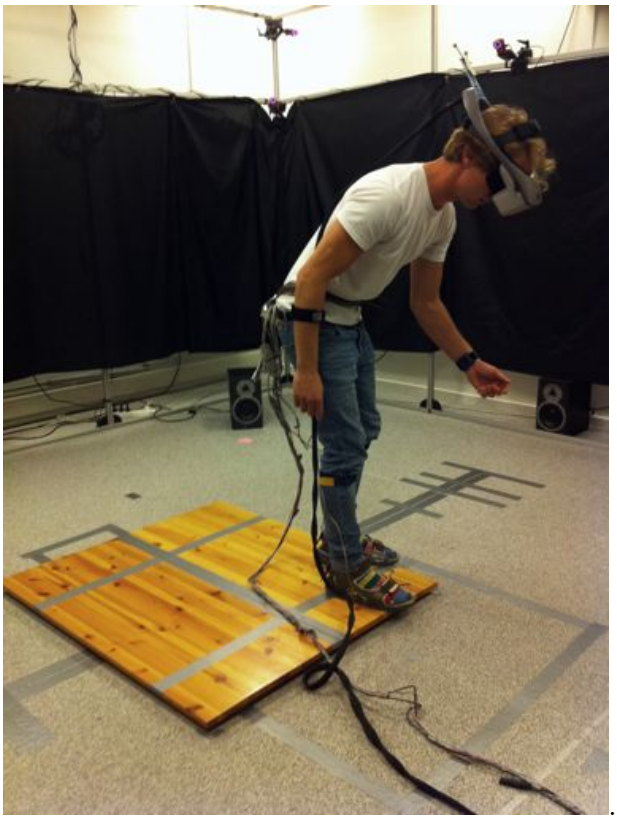

Figure 3.9: A subject performing the experiment.

possible to notice the motion capture cameras as well as the surround sound system.

3.2 Experiment design

We designed two experiments whose goal was to investigate the role of auditory and haptic feedback in enhancing presence and realism in the simulated virtual environment. As can be seen in Figure 3.9, in the first experiment subjects were asked to stand on a physical wooden plank while experiencing the environment. Such plank was not present in the second experiment. The reason was to investigate whether passive haptic had an effect in the results.

Both experiments were designed as within-subjects experiments, where half of the subjects experienced the condition without audio-haptic feedback (named NF in the following) first and the one with audio-haptic feedback (named F in the following) afterwards, while the other half experienced the condition with audio-haptic feedback first and the one without audio-haptic feedback afterwards.

3.2.1 Equipment and task

Before starting the experiment, each participant was asked to wear the HMD and haptic shoes previously described, together with a wireless Q sensor device developed by Affectiva (www.affectiva.com), which, placed around the wrist, allows to measure skin conductance and temperature. As can be seen in Figure 3.10, subjects were also asked to wear a wireless device able to measure heart beat (Scosche mytrek wireless pulse monitor).

After being ready to start the experiment, subjects were taken on the wooden platform, for those subjects exposed to the condition with passive feedback. For about a minute, subjects were allowed to freely explore the visual environment. In addition to allowing the participants to become familiar with the equipment, the hope was that this would minimize the effects of the orienting effect, that is,

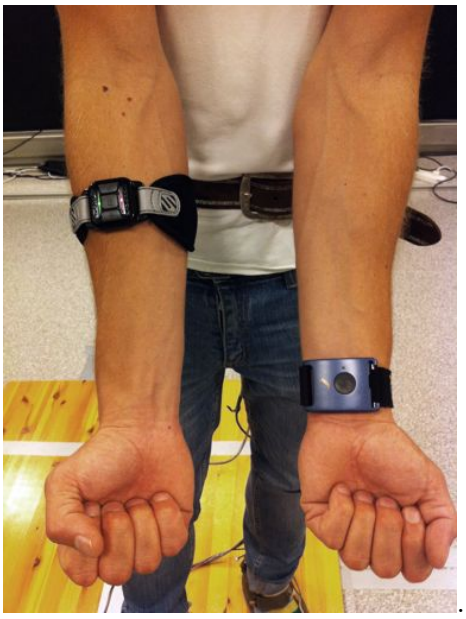

Figure 3.10: A view of the wireless devices used to measure physiological data.

individuals usually elicit a stronger physiological response the first time they are exposed to a given stimulus event [56].

Before starting the actual experiment, subjects were instructed that their task was to find three objects in the environment: a backpack, a camera and a hat.

The objects where sufficiently hard to see in such a way to encourage subjects to explore the environment.

After two minutes, subjects were asked to stop their search and to complete a presence questionnaire described later. Once subjects were done with the questionnaire, they were asked to repeat the experiment with the other condition. After two minutes, subjects were again asked to stop the experience of visiting the environment and asked to fill the same presence questionnaire. At the end of the experiment, subjects were asked questions to assess their ability of recognizing the feedback provided.

3.2.2 Participants

Forty participants were divided in two groups (n=20) to perform the two experiments. The two groups were composed respectively of 15 men and 5 women, aged between 20 and 34 (mean=23.05, standard deviation=3.13), and of 15 men and 5 women, aged between 20 and 32 (mean=23.5, standard deviation=3.17). Participants were primarily recruited from the campus of the Media Technology Department of the Aalborg University Copenhagen; however no restrictions on background were imposed. All participants reported normal, or corrected to normal, hearing. Participants were primarily recruited from the campus of the Media Technology Department of the Aalborg University Copenhagen; however no restrictions on background were imposed.

Participants were not awarded after the completion of the study. They were provided an informed consent form discussing the possible effects of participation in the study. Additionally, participants were informed that they could stop at any time during the experiment.

3.3 Results

In this section we present the results of both experiments, discussing the importance of feedback by examining both the case with passive haptics and the case without passive haptics.

NF-F		F-NF	
NF	F	NF	F
29.46 ± 0.80	30.27 ± 1.09	30,51 ± 0,95	29,92 ± 0,79

Table 3.1: Mean skin temperature (degrees celsius) for the two condition-orders NF-F and F-NF, experiment with passive haptics.

NF-F		F-NF	
F	NF	F	NF
31.19 ± 0.63	30.71 ± 0.62	30.42 ± 0.70	31.01 ± 0.69

Table 3.2: Mean skin temperature (degrees celsius) for the two condition-orders NF-F and F-NF, experiment without passive haptics.

3.3.1 Measuring physiological data

Table 3.1 and Table 3.2 show the results of the measured temperature with and without passive haptics respectively.

Table 3.3 and Table 3.4 show the results pertaining to the measures of skin conductance and skin temperature.

The skin temperature and skin conductance measures used during the experiment including passive haptic feedback did generally not suggest an increase in presence as a consequence of the added feedback. It is possible that the skin temperature measure have been corrupted by the orienting effect, individuals usually elicit a stronger physiological response the first time they are exposed to a given stimulus event. The results suggest that the participants in average experienced an increase in skin temperature between the first and the second trial, regardless of what condition was experienced first. Note that these differences were statistically significant ($p(19) = -5{,}4930$, $p \leq 0.05$). Similarly, significant difference was found between the averages pertaining to the skin temperature in case of both condition orders (NF-F: $t(9) = 7945$, $p = \leq 0.05$ and F-NF: $t(9) = -4.1416$, $p \leq 0.05$). It is possible to offer at least two explanations for this set of results, one being that the participants generally found the first exposure to the VE the scariest and therefore had a lower skin temperature during the first trial. A second explanation is that the high temperature within the laboratory caused their temperature to rise gradually for the duration of the experiment. Notably, the one explanation does not preclude the other. The results obtained from the skin conductance measure are inconclusive at best since no meaningful tendencies are present. It is expected that the skin conductance increases during stressful situations since one would sweat more. When the participants experienced the o feedbackondition prior to the feedback condition there was an increase in the average skin conductance across the two trials. However, when the order of the conditions was reversed there was a decrease in the average across the two conditions. It would in other words appear that the difference in averages cannot be ascribed to the orienting effect. Indeed it seems likely that the participants have not have experienced an increase in presence due to the added feedback. Bluntly put, one may argue that there is no logical explanation why one group should feel a stronger sensation of presence than the other. In both conditions the participants found themselves on a wooden platform above the canyon and auditory and haptic stimuli experienced in both conditions supported the illusion of being in this place. In

NF-F		F-NF	
NF	F	NF	F
1,69 ± 2,15	1,71 ± 1,68	1,54 ± 1,87	1,29 ± 1,37

Table 3.3: Mean skin conductance (microSiemens) for the two condition-orders NF-F and F-NF, experiment with passive haptics.

NF-F		F-NF	
F	NF	F	NF
5.79 ± 8.18	4.80 ± 8.41	2.07 ± 2.43	2.30 ± 2.35

Table 3.4: Mean skin conductance (microSiemens) for the two condition-orders NF-F and F-NF, experiment without passive haptics.

the o feedbackondition the participants experienced standing on a sturdy wooden platform, since they heard and felt their real world footsteps on the wooden platform in the lab. Contrarily, in the eedbackondition, the participants experienced standing on an unstable platform since the auditory and haptic feedback gave the impression that the wooden planks were giving in as the participants stepped on them. While it is possible that an unstable platform may be experienced as scarier, there is no reason to believe that that a person should feel more present.

The measures of skin temperature and skin conductance applied during the experiment where passive haptic feedback yielded similar results. It is worth noting that there is a significant difference between the averages pertaining to the skin temperature in case of both condition orders for the experiment without passive haptic feedback. (NF-F: t(9) = 4.6577, p = $\leq$ 0.05 and F-NF: t(9) = -5,0466, p $\leq$ 0.05). As was the case with the experiment including passive haptic feedback, it seems possible that the participants have experienced less stress or fear during the second condition or simply have gotten gradually warmer due to the high temperature in the lab. Moreover, there was a significant difference between the mean skin conductance during condition order NF-F (t(9) = 2,5008, p $\leq$0.05). However, the corresponding means related to condition order F-NF did not differ significantly. With this being said, it should be stressed that a significant difference would not have changed the fact that these results do not suggest that the participants experienced a higher level of presence during the condition with added feedback. Regardless of the condition order the participants seem to have experienced an increase in skin conduction from the first to the second condition, which suggests a higher degree of skin perspiration. Notably a comparison by means of paired sample t-test revealed that there is a significant difference between the first and second trial, both in case of skin temperature (p(19) = -6,946, p $\leq$ 0.05) and skin conductance (p(19) = -2,4511, $\leq$0.05).

Table 3.5 and Table 3.6 show the results for the heart-rate measurements in the two experiments. As it is possible to notice, the average values in both the conditions and in both the experiments are always higher for the typology of stimuli presented first. In particular the statistical analysis conducted by means of a paired t-test revealed that such differences are significant for the heart-rate mean and maximum of the condition F-NF in the first experiment (t(9) = 4.7555, $p \leq 0.5$ and t(9) = 3.0251, $p \leq 0.5$ respectively), and for the heart-rate mean and maximum of the condition NF-F in the second experiment (t(9) = 3.4804, $p \leq 0.5$ and t(9) = 5.0558, $p \leq 0.5$ respectively).

	Trials NF-F			**Trials F-NF**		
	Mean	**Max.**	**Min.**	**Mean**	**Max.**	**Min.**
NF	90±13.96	99.2±14.92	81.9±14.45	89.2±15	97.2±14.28	83.7±15.83
F	87.5±11.17	96.9±8.68	78.6±11.89	93.3±14.17	101.6±14.71	86.2±15.12

Table 3.5: Heart-rate results of the experiment with passive haptics. Legenda: NF-F: trials in which the no feedback condition was presented first and the feedback condition afterwards; F-NF: trials in which the feedback condition was presented first and the no feedback condition afterwards.

Note there is a significant difference between the averages pertaining to the skin temperature in case of both condition orders. This can be interpreted in two different, albeit not necessary mutually exclusive, ways. First, the participants may have experienced less stress or fear during the second condition. Secondly, it is possible that the participants got gradually warmer as due to the hight

	Trials NF-F			Trials F-NF		
	Mean	**Max.**	**Min.**	**Mean**	**Max.**	**Min.**
NF	92.2±9.54	103.3±10.97	84.1±9.38	86.1±10.21	95.8±11.79	77.6±8.07
F	88.1±8.19	94.4±10.04	82.2±9.24	89.6±14.45	99.3±19.04	79.5±10.69

Table 3.6: Heart-rate results of the experiment without passive haptics. Legenda: NF-F: trials in which the no feedback condition was presented first and the feedback condition afterwards; F-NF: trials in which the feedback condition was presented first and the no feedback condition afterwards.

temperature in the lab. Moreover, there was a significant difference between the mean skin conductance during condition order NF-F. The corresponding means related to condition order F-NF were not significantly different ($p = 0.2$). Even if it had been it wouldn't have provided us with particularly valuable information, considering how the averages differ. Regardless of the condition order the participants seem to have experienced an increase in skin temperature, which suggests that they became less afraid or that they were simply sweating. The following two tables detail the means of the skin conductance and temperature. The condition order is in other words not taken into consideration.

It should be noted that neither of the means presented in the two tables above significantly differ from one another. In addition to comparing these averages I also performed a comparison of the means obtained from each trial, that is, I compared the mean skin conductance and temperature obtained from all first trials with the ones pertaining to the second trial. The two tables below detail the corresponding data.

Notably a comparison by means of paired sample t-test revealed that there is a significant difference between the first and second trial, both in case of skin temperature and skin conductance.

3.3.2 Measuring presence

The participantsxperience of presence was assessed by means of the Slater-Usoh-Steed (SUS) questionnaire [94], [113], which is a questionnaire intended to evaluate the experience after exposures to a virtual environment (VE). The SUS questionnaire contains six items that evaluate the experience of presence in terms of, the participant sense of being in the VE, the extent to which the participant experienced the VE as the dominant reality, and the extent to which the VE is remembered as a place. All items are answered on scales ranging from 1 to 7 where the highest scores would be indicative of presence [113]:

Q1: Please rate your sense of being in the virtual environment, on a scale of 1 to 7, where 7 represents your normal experience of being in a place.

Q2: To what extent were there times during the experience when the virtual environment was the reality for you?

Q3: When you think back to the experience, do you think of the virtual environment more as images that you saw or more as somewhere that you visited?

Q4: During the time of the experience, which was the strongest on the whole, your sense of being in the virtual environment or of being elsewhere?

Q5: Consider your memory of being in the virtual environment. How similar in terms of the structure of the memory is this to the structure of the memory of other places you have been today?

Q6: During the time of your experience, did you often think to yourself that you were actually in the virtual environment?

The general level of presence experienced by the participants may be determined by summarizing the data obtained from the all of the questionnaire items in two ways. First, one may present the central tendency as the mean of all ratings to all items and the variability may thus be presented as the corresponding standard deviation. Secondly, it is possible to present the general experience of presence across participants (SUS count), as the mean of the individual presence scores. The presence score is taken as the sum of scores of 6 and 7 out of the number of questions posed [113].

Tables 3.7 and 3.8 illustrate the questionnaire evaluations for the first and second experiment respectively.

	Trials NF-F		**Trials F-NF**	
	NF	**F**	**NF**	**F**
Q1	5.3±1.49	6±1.15	5.63±1.2	5.45±1.12
Q2	5.5±1.08	5.6±1.26	5.45±1.36	5.45±1.12
Q3	3.9±1.79	5.1±1.59	5.09±1.57	5.81±0.98
Q4	5.9±0.99	6.1±0.56	5.18±1.53	6.18±0.87
Q5	3.5±1.5	4.7±1.63	4.72±1.79	4.54±0.93
Q6	4.9±1.72	5.3±1.7	5.09±2.02	6.18±1.16
SUS count	0.38±2.13	0.65±2.16	0.6±0.89	0.6±2.52

Table 3.7: Questionnaire results of the experiment with passive haptics.

	Trials NF-F		**Trials F-NF**	
	NF	**F**	**NF**	**F**
Q1	4.7±1.15	5.4±1.26	5.3±1.56	5.6±0.96
Q2	4.±1.15	5.±1.33	4.8±1.22	5.3±0.94
Q3	4.5±1.77	4.7±1.56	4.8±1.68	4.8±1.68
Q4	5.4±1.07	5.3±1.15	5.5±0.97	5.3±1.15
Q5	3.9±2.55	4.3±2.49	4.1±1.79	4.2±2.14
Q6	3.6±1.57	5.2±1.68	5.7±1.25	4.9±1.37
SUS count	0.28±1.17	0.5±1.41	0.46±1.75	0.41±0.75

Table 3.8: Questionnaire results of the experiment without passive haptics.

As outlined in [113], to check if the differences found in the questionnaire results for the two typologies of stimuli F and NF are statistically significant, one should not compare the means of the questionnaire items results, but rather the number of answers having a score of 6 or 7. Following this approach we found statistical significance in both the experiments (with and without passive haptics) for the trials in which the no feedback condition was presented first and the feedback condition afterwards ($\chi^2(1) = 5.0364$, p-value = 0.02482 and $\chi^2(1) = 7.5083$, p-value = 0.006141 respectively). Conversely, no significance was found in none of the two experiments for the trials in which the feedback condition was presented first and the no feedback condition afterwards.

It is interesting to notice that the mean presence score pertaining to the feedback condition is significantly higher when this condition was presented first while there was no significant difference between the scores when the no-feedback condition was presented first , despite this average being the higher. One may argue that this lends some credence to the claim that the addition of the feedback did increase the participant sensation of presence. This claim does, however, not come without reservations considering the inconclusive results obtained from the physiological measures, which did not indicate a difference in presence. While it is possible that this data has be tainted by the orienting effect as well as factors such as the temperature in the room, we do not feel that it is possible to disregard them

altogether. It does therefore remain an open question whether the added feedback did in fact increase the sensation of presence on behalf of the participants. With this being said, it is worth noting that results obtained from the questionnaire at least in part correspond with the statements made by the participants who generally thought that the feedback added to the sense of realism and in some cases intensified the experience of vertigo.

Moreover, while the choice of the SUS-presence questionnaire was motivated by the fact that it is extensively validated and used in the VR community, it can be questioned whether it is the most suitable for examining the relationship between feedback and presence.

3.3.3 Realism and audio-haptic feedback

As a final analysis of the experiments' results, it is interesting to discuss the observations provided by the subjects when the experiments were completed. Specifically, we asked subjects if they had noticed any difference on the two conditions and, in affermative case, if they could elaborate on the differences noticed and how they affected their experience.

During the first experiment, when asked whether they had noticed a difference between the two trials, 13 of the participants mentioned that they had noticed the change in the haptic and/or auditory feedback provided by the shoes. Precisely, 5 subjects noticed a difference in both auditory and haptic feedback, 7 only noticed the difference in auditory feedback, while 1 only noticed the difference in haptic feedback. All of the participants who noticed the difference expressed a preference towards the added feedback. When asked to elaborate, 11 of the 13 stated that it added realism, 5 felt that it made the experience more scary or intensified the sensation of vertigo, while 1 explicitly stated that it increased the sensation of presence in the virtual environment.

During the second experiment, out of the 20 participants, 16 noticed the additional feedback, 5 participants noticed both the auditory and haptic feedback while 7 just noticed the sound and 4 only noticed the haptic feedback. With one exception, all of the participants who noticed the difference preferred the additional feedback. The one participant who did not, described that he did like the haptic feedback, but he had found it too intense. Out of the 16 who noticed the feedback 13 thought that it added realism, 2 described that it made it more scary and 2 explicitly stated that it intensified the sensation of eing there

Such observations show that subjects indeed were able to notice and appreciate the provided feedback in both experimentals' conditions. The lack of the same evidence while analyzing physiological data or presence questionnaire can be due to the fact that the provided feedback does not necessarily elicit a higher physiological response or sense of presence.

3.4 Conclusion and future work

In this chapter, we have described two experiments whose goal was to assess the role of auditory and haptic feedback delivered at feet level to enhance sense of presence and realism in a multimodal virtual environments. The first experiment was performed with passive haptics, while subjects were experiencing the environment with and without auditory and haptic feedback. The second experiment was performed without passive haptics. While quantitative results obtained while measuring physiological data and performing a post-experimental presence questionnaire do not show significant differences among the two conditions, subjects were actually able to perceive the differences among the experiences. As discussed in the paper, indeed several subjects noticed the auditory and haptic feedback and reported on appreciating it and experiencing it as a way to simulate realism and sense of "being there".

In the future, we are interested in further investigating the role of auditory and haptic feedback provided at feet-level, also as a mean to provide useful information such as indications for navigation in virtual environments or feedback for actions in computer games.

References

[1] M. Addlesee, A. Jones, F. Livesey, and F. Samaria. The ORL active floor [sensor system]. *Personal Communications, IEEE [see also IEEE Wireless Communications]*, 4(5):35–41, 1997.

[2] G. Arechavaleta, J.-P. Laumond, H. Hicheur, and A. Berthoz. The nonholonomic nature of human locomotion: a modeling study. In *Proceedings of IEEE International Conference on Biomedical Robotics and Biomechatronics*, pages 158–163, 2006.

[3] S. Beckhaus, K. Blom, and M. Haringer. *ChairIO - the Chair-Based Interface*. Ed. Magerkurth und Rötzler, Shaker Verlag, 2007.

[4] A. Benbasat, S. Morris, and J. Paradiso. A wireless modular sensor architecture and its application in on-shoe gait analysis. In *Sensors, 2003. Proceedings of IEEE*, volume 2, 2003.

[5] L. Bouguila, M. Ishii, and M. Sato. Realizing a new step-in-place locomotion interface for virtual environment with large display system. In *Proceedings of the workshop on Virtual environments 2002*, pages 197–207, Barcelona, Spain, 2002. Eurographics Association.

[6] D. Bowman, E. Kruijff, J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Boston (MA), 2005.

[7] D. A. Bowman, E. Kruijff, J. J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, 1 edition, Aug. 2004.

[8] D. A. Bowman, E. Kruijff, J. J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Boston (MA), 2005.

[9] G. R. Bradski. Computer vision face tracking for use in a perceptual user interface. *Intel Technology Journal*, 2(2):12–21, 1998.

[10] D. C. Brogan, R. A. Metoyer, and J. K. Hodgins. Dynamically simulated characters in virtual environments. In *ACM SIGGRAPH 97 Visual Proceedings: The art and interdisciplinary programs of SIGGRAPH '97*, page 216, Los Angeles, California, United States, 1997. ACM.

[11] G. Casiez, D. Vogel, Q. Pan, and C. Chaillou. RubberEdge: reducing clutching by combining position and rate control with elastic feedback. In *Proceedings of the ACM symposium on User interface software and technology*, pages 129–138, Newport, Rhode Island, USA, 2007. ACM.

[12] M. Chion, C. Gorbman, and W. Murch. *Audio-vision: sound on screen*. Columbia Univ Pr, 1994.

[13] I. Choi and C. Ricci. Foot-mounted gesture detection and its application in virtual environments. In *Systems, Man, and Cybernetics, 1997. 'Computational Cybernetics and Simulation'., 1997 IEEE International Conference on*, volume 5, pages 4248–4253 vol.5, 1997.

[14] G. Cirio, M. Marchal, S. Hillaire, and A. Lécuyer. Six degrees-of-freedom haptic interaction with fluids. *IEEE Transactions on Visualization and Computer Graphics*, 2010.

[15] G. Cirio, M. Marchal, T. Regia-Corte, and A. Lécuyer. The magic barrier tape: a novel metaphor for infinite navigation in virtual worlds with a restricted walking workspace. In *Proceedings of the ACM Symposium on VRST*, pages 155–162, 2009.

[16] G. Courtine and M. Schieppati. Human walking along a curved path. i. body trajectory, segment orientation and the effect of vision. *Eu. J. of Neuroscience*, 18(1):177–190, 2003.

[17] R. P. Darken, W. R. Cockayne, and D. Carmein. The omni-directional treadmill: a locomotion device for virtual worlds. In *Proceedings of the 10th annual ACM symposium on User interface software and technology*, pages 213–221, Banff, Alberta, Canada, 1997. ACM.

[18] T. Delbrück, A. M. Whatley, R. Douglas, K. Eng, K. Hepp, and P. F. Verschure. A tactile luminous floor for an interactive autonomous space. *Robotics and Autonomous Systems*, 55(6):433–443, June 2007.

[19] K. Doel. Physically based models for liquid sounds. *ACM Transactions on Applied Perception (TAP)*, 2(4):534–546, 2005.

[20] L. Dominjon, A. Lecuyer, J. Burkhardt, G. Andrade-Barroso, and S. Richir. The "Bubble" technique: interacting with large virtual environments using haptic devices with limited workspace. In *Proceedings of World Haptics conference*, pages 639–640, 2005.

[21] C. Drioli and D. Rocchesso. Acoustic rendering of particle-based simulation of liquids in motion. *Journal of Multimodal User Interfaces*, 2011.

[22] N. I. Durlach, A. S. Mavor, N. R. C. U. C. on Virtual Reality Research, and Development. *Virtual Reality: Scientific and technological challenges*. 1995.

[23] D. Engel, C. Curio, L. Tcheang, B. Mohler, and H. H. Bülthoff. A psychophysically calibrated controller for navigating through large environments in a limited free-walking space. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*, pages 157–164, Bordeaux, France, 2008. ACM.

[24] J. Feasel, M. C. Whitton, and J. Wendt. Llcm-wip: Low-latency, continuous-motion walking-in-place. In *IEEE Symposium on 3D User Interfaces*, pages 97 –104, 2008.

[25] K. Fernandes, V. Raja, and J. Eyre. Cybersphere: the fully immersive spherical projection system. *Communications of the ACM*, 46(9):146, 2003.

[26] K. J. Fernandes, V. Raja, and J. Eyre. Cybersphere: the fully immersive spherical projection system. *Commun. ACM*, 46(9):141–146, 2003.

[27] G. Franz. Splashes as sources of sound in liquids. *The Journal of the Acoustical Society of America*, 31:1080, 1959.

[28] X. Fu and D. Li. Haptic shoes: representing information by vibration. In *proceedings of the 2005 Asia-Pacific symposium on Information visualisation - Volume 45*, pages 47–50, Sydney, Australia, 2005. Australian Computer Society, Inc.

[29] P. Fuchs, G. Moreau, J.-P. Papin, and A. Berthoz. *Le traité de la réalité virtuelle*. Les Presses de l'Ecole des Mines, 2001.

[30] N. Griffith and M. Fernström. LiteFoot - a floor space for recording dance and controlling media. pages 475–481, Ann Arbor, Michigan, 1998.

[31] J. Hollerbach. Locomotion interfaces. In K. M. Stanney, editor, *Handbook of Virtual Environments*, page 1232. 2002.

[32] J. M. Hollerbach, D. Checcacci, H. Noma, Y. Yanagida, and N. Tetsutani. Simulating side slopes on locomotion interfaces using torso forces. In *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'03)*, page 91. IEEE Computer Society, 2003.

[33] J. M. Hollerbach, R. Mills, D. Tristano, R. R. Christensen, W. B. Thompson, and Y. Xu. Torso force feedback realistically simulates slope on Treadmill-Style locomotion interfaces. *The International Journal of Robotics Research*, 20(12):939–952, Dec. 2001.

[34] J. M. Hollerbach, Y. Xu, R. R. Christensen, and S. C. Jacobsen. Design specifications for the second generation sarcos treadport locomotion interface. *Haptic Symposium, Proceedings of ASME Dynamic Systems and Control Division*, 69:1293—1298, 2000.

[35] S. Howison, J. Ockendon, and J. Oliver. Deep-and shallow-water slamming at small and zero deadrise angles. *Journal of engineering mathematics*, 42(3):373–388, 2002.

[36] L. Hutchings. System and method for measuring movement of objects. united states patent 5724265., Mar. 1998.

[37] V. Interrante, B. Ries, and L. Anderson. Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, pages 167—170, 2007.

[38] H. Iwata and T. Fujii. VIRTUAL PERAMBULATOR: a novel interface device for locomotion in virtual environment. In *Proceedings of the 1996 Virtual Reality Annual International Symposium (VRAIS 96)*, page 60. IEEE Computer Society, 1996.

[39] H. Iwata, H. Yano, H. Fukushima, and H. Noma. CirculaFloor: a locomotion interface using circulation of movable tiles. In *Proceedings of the 2005 IEEE Conference 2005 on Virtual Reality*, pages 223–230. IEEE Computer Society, 2005.

[40] H. Iwata, H. Yano, H. Fukushima, and H. Noma. Circulafloor: A locomotion interface using circulation of movable tiles. In *Proceedings of IEEE Virtual Reality Conference*, pages 223–230, 2005.

[41] H. Iwata, H. Yano, and F. Nakaizumi. Gait master: A versatile locomotion interface for uneven virtual terrain. In *Proceedings of IEEE Virtual Reality Conference*, pages 131–137, 2001.

[42] H. Iwata, H. Yano, and H. Tomioka. Powered shoes. In *ACM SIGGRAPH 2006 Emerging technologies*, page 28, Boston, Massachusetts, 2006. ACM.

[43] H. Iwata, H. Yano, and M. Tomiyoshi. String walker. In *ACM SIGGRAPH 2007 emerging technologies*, page 20, San Diego, California, 2007. ACM.

[44] H. Iwata and Y. Yoshida. Path reproduction tests using a torus treadmill. *Presence: Teleoper. Virtual Environ.*, 8(6):587–597, 1999.

[45] H. Iwata and Y. Yoshida. Path reproduction tests using a torus treadmill. *Presence: Teleoperators and Virtual Environments*, 8(6):587–597, 1999.

[46] H. Iwata and Y. Yoshida. Path reproduction tests using a torus treadmill. *Presence: Teleoperators and Virt. Env.*, 8(6):587–597, 1999.

[47] E. Johnstone. A MIDI foot controller - the PodoBoard. In *Proceedings of International Computer Music Conference*, Montréal, Canada, 1991.

[48] L. Kohli, E. Burns, D. Miller, and H. Fuchs. Combining passive haptics with redirected walking. In *Proceedings of the 2005 international conference on Augmented tele-existence*, pages 253–254, Christchurch, New Zealand, 2005. ACM.

[49] Y. Kume, A. Shirai, M. Kusahara, and M. Sato. Foot interface: fantastic phantom slipper. In *ACM SIGGRAPH 98 Conference abstracts and applications*, page 114, Orlando, Florida, United States, 1998. ACM.

[50] N. N. Latypov. The virtusphere, 2006.

[51] J. J. LaViola, D. A. Feliz, D. F. Keefe, and R. C. Zeleznik. Hands-free multi-scale navigation in virtual environments. In *Proceedings of the ACM symposium on Interactive 3D graphics*, pages 9–15. ACM, 2001.

[52] A. Lécuyer, J.-M. Burkhardt, J.-M. Henaff, and S. Donikian. Camera motions improve the sensation of walking in virtual environments. In *Proceedings of IEEE Virtual Reality Conference*, pages 11–18, 2006.

[53] A. Lécuyer, J.-M. Burkhardt, J.-M. Henaff, and S. Donikian. Camera motions improve the sensation of walking in virtual environments. In *Proceedings of IEEE Virtual Reality*, pages 11–18, 2006.

[54] M. Lesser. Thirty years of liquid impact research: a tutorial review. *Wear*, 186:28–34, 1995.

[55] M. Longuet-Higgins. An analytic model of sound production by raindrops. *Journal of Fluid Mechanics*, 214(1):395–410, 1990.

[56] M. Meehan, B. Insko, M. Whitton, and F. Brooks Jr. Physiological measures of presence in stressful virtual environments. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 645–652. ACM, 2002.

[57] M. Minnaert. Xvi. on musical air-bubbles and the sounds of running water. *Philosophical Magazine Series 7*, 16(104):235–248, 1933.

[58] E. R. Miranda and M. Wanderley. *New Digital Musical Instruments: Control And Interaction Beyond the Keyboard*. A-R Editions, Inc., 1st edition, July 2006.

[59] J. Monaghan. Smoothed particle hydrodynamics. *Annual review of astronomy and astrophysics*, 30:543–574, 1992.

[60] W. Moss, H. Yeh, J. Hong, M. Lin, and D. Manocha. Sounding liquids: Automatic sound synthesis from fluid simulation. *ACM Transactions on Graphics (TOG)*, 29(3):21, 2010.

[61] M. M
"uller, D. Charypar, and M. Gross. Particle-based fluid simulation for interactive applications. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 154–159. Eurographics Association, 2003.

[62] M. M
"uller, B. Solenthaler, R. Keiser, and M. Gross. Particle-based fluid-fluid interaction. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 237–244. ACM, 2005.

[63] N. Nitzsche, U. Hanebeck, and G. Schmidt. Motion compression for telepresent walking in large target environments. *Presence: Teleoperators and Virtual Environments*, 13(1):44–60, 2004.

[64] N. Nitzsche, U. D. Hanebeck, and G. Schmidt. Motion compression for telepresent walking in large target environments. *Presence: Teleoper. Virtual Environ.*, 13(1):44–60, 2004.

[65] H. Noma. Design for locomotion interface in a large-scale virtual environment. ATLAS: ATR locomotion interface for active Self-Motion. In *Proc. of the ASME Dynamic Systems and control division*, pages 111–118, 1998.

[66] H. Noma, T. Sugihara, and T. Miyasato. Development of ground surface simulator for Tel-E-Merge system. In *Proceedings of the IEEE Virtual Reality 2000 Conference*, page 217. IEEE Computer Society, 2000.

[67] R. Nordahl. Evaluating environmental sounds from a presence perspective for virtual reality applications. *EURASIP Journal on Audio, Speech, and Music Processing*, 2010:4, 2010.

[68] R. Nordahl, A. Berrezag, S. Dimitrov, L. Turchet, V. Hayward, and S. Serafin. Preliminary experiment combining virtual reality haptic shoes and audio synthesis. In *Proc. Eurohaptics*, 2010.

[69] R. Nordahl, A. Berrezag, S. Dimitrov, L. Turchet, V. Hayward, and S. Serafin. Preliminary experiment combining virtual reality haptic shoes and audio synthesis. *Haptics: Generating and Perceiving Tangible Sensations*, pages 123–129, 2010.

[70] R. Nordahl, S. Serafin, and L. Turchet. Sound synthesis and evaluation of interactive footsteps for virtual reality applications. In *Proc. IEEE VR 2010*, 2010.

[71] R. J. Orr and G. D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *CHI '00 extended abstracts on Human factors in computing systems*, pages 275–276, The Hague, The Netherlands, 2000. ACM.

[72] S. Papetti, M. Civolani, F. Fontana, A. Berrezag, and V. Hayward. Audio-tactile display of ground properties using interactive shoes. In *Proc. 5th Int. Haptic and Auditory Interaction Design Workshop*, pages 117–128, Sept. 2010.

[73] J. Paradiso, C. Abler, K. yuh Hsiao, and M. Reynolds. The magic carpet: physical sensing for immersive environments. In *CHI '97 extended abstracts on Human factors in computing systems: looking to the future*, pages 277–278, Atlanta, Georgia, 1997. ACM.

[74] J. Paradiso, K. Hsiao, and E. Hu. Interactive Music for Instrumented Dancing Shoes. *Proc. of the 1999 International Computer Music Conference*, pages 453–456, 1999.

[75] J. Paradiso, E. Hu, and K. yuh Hsiao. The CyberShoe: a wireless multisensor interface for a dancer's feet. *In Proc. of International Dance and Technology 99, Tempe AZ*, pages 26—28, 2006.

[76] R. Pausch, T. Burnette, D. Brockway, and M. E. Weiblen. Navigation and locomotion in virtual worlds via flight into hand-held miniatures. In *Proceedings of ACM SIGGRAPH*, pages 399–400. ACM, 1995.

[77] R. Pausch, T. Burnette, D. Brockway, and M. E. Weiblen. Navigation and locomotion in virtual worlds via flight into hand-held miniatures. In *Proceedings of SIGGRAPH'95*, pages 399–400, 1995.

[78] T. Peck, M. Whitton, and H. Fuchs. Evaluation of reorientation techniques for walking in large virtual environments. In *Virtual Reality Conference, 2008. VR '08. IEEE*, pages 121–127, 2008.

[79] A. Pelah and J. Koenderink. Editorial: Walking in real and virtual environments. *ACM Transactions on Applied Perception (TAP)*, 4(1):1, 2007.

[80] R. F. Pinkston. A touch sensitive dance floor/MIDI controller. *The Journal of the Acoustical Society of America*, 96(5):3302, Nov. 1994.

[81] A. Polaine. The flow principle in interactivity. In *Proceedings of the 2nd Au. conf. on Interactive Entertainment*, page 158, 2005.

[82] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proceedings of Eurographics*, 2001.

[83] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proceedings of Eurographics*, 2001.

[84] S. Razzaque, D. Swapp, M. Slater, M. C. Whitton, and A. Steed. Redirected walking in place. In *Proceedings of the workshop on Virtual environments 2002*, pages 123–130, Barcelona, Spain, 2002. Eurographics Association.

[85] S. Razzaque, D. Swapp, M. Slater, M. C. Whitton, and A. Steed. Redirected walking in place. In *Proceedings of the EGVE '02 workshop on Virtual environments*, pages 123–130, 2002.

[86] B. Richardson, K. Leydon, M. Fernstrom, and J. A. Paradiso. Z-Tiles: building blocks for modular, pressure-sensing floorspaces. In *CHI '04 extended abstracts on Human factors in computing systems*, pages 1529–1532, Vienna, Austria, 2004. ACM.

[87] E. Richardson. The sounds of impact of a solid on a liquid surface. *Proceedings of the Physical Society. Section B*, 68:541, 1955.

[88] A. Rovers and H. van Essen. FootIO - design and evaluation of a device to enable foot interaction over a computer network. In *Eurohaptics Conference, 2005 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2005. World Haptics 2005. First Joint*, pages 521–522, 2005.

[89] J. Schacher and M. Neukom. Ambisonics Spatialization Tools for Max/MSP. In *Proceedings of the International Computer Music Conference*, 2006.

[90] H. Schmidt, S. Hesse, R. Bernhardt, and J. Krüger. HapticWalker—a novel haptic foot device. *ACM Trans. Appl. Percept.*, 2(2):166–180, 2005.

[91] M. Schwaiger, T. Thummel, and H. Ulbricht. A 2d-motion platform: The cybercarpet. In *Proceedings of Second Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 415–420, 2007.

[92] M. Schwaiger, T. Thümmel, and H. Ulbrich. Cyberwalk: Implementation of a ball bearing platform for humans. In *Human-Computer Interaction. Interaction Platforms and Techniques*, pages 926–935. 2007.

[93] S. Serafin, L. Turchet, R. Nordahl, S. Dimitrov, A. Berrezag, and V. Hayward. Identification of virtual grounds using virtual reality haptic shoes and sound synthesis. In *Proc. of Eurohaptics symposium on Haptics and Audio-visual environments*, 2010.

[94] M. Slater, M. Usoh, and A. Steed. Depth of presence in immersive virtual environments.

[95] M. Slater, M. Usoh, and A. Steed. Taking steps: the influence of a walking technique on presence in virtual reality. *ACM Trans. Comput.-Hum. Interact.*, 2(3):201–219, 1995.

[96] M. Slater, M. Usoh, and A. Steed. Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Transactions on Computer-Human Interaction*, 2(3):201–219, 1995.

[97] M. Slater, M. Usoh, and A. Steed. Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Trans. on Computer-Human Interaction*, 2(3):201–219, 1995.

[98] P. Srinivasan, D. Birchfield, G. Qian, and A. Kidane. Design of a pressure sensitive floor for multimodal sensing. In *Proceedings of the Ninth International Conference on Information Visualisation*, pages 41–46. IEEE Computer Society, 2005.

[99] K. Stanney and al. *Handbook of Virtual Environments: Design, Implementation, and Applications*. Lawrence Erlbaum Associates, 2002.

[100] K. M. Stanney et al. *Handbook of Virtual Environments: Design, Implementation, and Applications*. Lawrence Erlbaum Associates, 2002.

[101] F. Steinicke, G. Bruder, T. Ropinski, and K. H. Hinrichs. Moving towards generally applicable redirected walking. In *Proceedings of the Virtual Reality International Conference (VRIC)*, pages 15—24. IEEE Press, 2008.

[102] J. N. Templeman, P. S. Denbrook, and L. E. Sibert. Virtual locomotion: Walking in place through virtual environments. *Presence: Teleoper. Virtual Environ.*, 8(6):598–617, 1999.

[103] L. Terziman, A. Lécuyer, S. Hillaire, and J. M. Wiener. Can camera motions improve the perception of traveled distance in virtual environments? In *Proceedings of IEEE Virtual Reality Conference*, pages 131–134, 2009.

[104] H. Tomoyuki, K. Masato, K. Shingo, T. Shuntaro, N. Fumikazu, and I. Minoru. The ball array treadmill. *Proceedings of the Virtual Reality Society of Japan Annual Conference*, 6th:211–212, 2001.

[105] J. Torchinsky. Giant atari 2600 joystick. `http://www.vgg.com/jason/jdt/projects.html`, 2006. [Online; accessed 6-January-2011].

[106] S. Toshiaki, A. M, M. Tsutomu, and A. M. Terrain surface simulator ALF. *Proceedings of the Virtual Reality Society of Japan Annual Conference*, 4th:335–338, 1999.

[107] L. Turchet, R. Nordahl, A. Berrezag, S. Dimitrov, V. Hayward, and S. Serafin. Audio-haptic physically based simulation of walking sounds. In *Proc. of IEEE International Workshop on Multimedia Signal Processing*, 2010.

[108] L. Turchet, R. Nordahl, and S. Serafin. Examining the role of context in the recognition of walking sound. In *Proc. of Sound and Music Computing Conference*, 2010.

[109] L. Turchet, R. Nordahl, S. Serafin, A. Berrezag, S. Dimitrov, and V. Hayward. Audio-haptic physically-based simulation of walking on different grounds. In *Multimedia Signal Processing (MMSP), 2010 IEEE International Workshop on*, pages 269–273. IEEE, 2010.

[110] L. Turchet, S. Serafin, S. Dimitrov, and R. Nordahl. Physically based sound synthesis and control of footsteps sounds. In *Proceedings of Digital Audio Effects Conference*, 2010.

[111] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks, Jr. Walking > walking-in-place > flying, in virtual environments. In *Proceedings of SIGGRAPH '99*, pages 359–364, 1999.

[112] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and J. F. P. Brooks. Walking \> walking-in-place \> flying, in virtual environments. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 359–364. ACM Press/Addison-Wesley Publishing Co., 1999.

[113] M. Usoh, E. Catena, S. Arman, and M. Slater. Using presence questionnaires in reality. *Presence: Teleoperators And Virtual Environments*, 9(5):497–503, 2000.

[114] Y. Visell, J. Cooperstock, and K. Franinovic. The ecotile: An architectural platform for audio-haptic simulation in walking. In *Proc. of the 4th Intl. Conf. on Enactive Interfaces*, 2007.

[115] Y. Visell, J. Cooperstock, B. Giordano, K. Franinovic, A. Law, S. McAdams, K. Jathal, and F. Fontana. A vibrotactile device for display of virtual ground materials in walking. *Lecture Notes in Computer Science*, 5024:420–426, 2008.

[116] Y. Visell, J. Cooperstock, B. Giordano, K. Franinovic, A. Law, S. McAdams, K. Jathal, and F. Fontana. A vibrotactile device for display of virtual ground materials in walking. In *Haptics: Perception, Devices and Scenarios*, pages 420–426. 2008.

[117] Y. Visell, A. Law, J. Ip, S. Smith, and J. Cooperstock. Interaction capture in immersive virtual environments via an intelligent floor surface. In *Virtual Reality Conference (VR), 2010 IEEE*, pages 313–314. IEEE, 2010.

[118] C. Ware and S. Osborne. Exploration and virtual camera control in virtual three dimensional environments. In *Proceedings of ACM SIGGRAPH*, pages 175–183, 1990.

[119] J. Wendt, M. C. Whitton, and F. Brooks. Gud wip: Gait-understanding-driven walking-in-place. In *Proceedings of IEEE Virtual Reality*, pages 51 –58, 2010.

[120] B. Williams, G. Narasimham, B. Rump, T. McNamara, T. Carr, J. Rieser, and B. Bodenheimer. Exploring large virtual environments with an hmd when physical space is limited. In *Proceedings of the ACM Symposium on Applied Perception in Graphics and Visualization*, pages 41–48, 2007.

[121] B. Williams, G. Narasimham, B. Rump, T. P. McNamara, T. H. Carr, J. Rieser, and B. Bodenheimer. Exploring large virtual environments with an HMD when physical space is limited. In *Proceedings of the ACM symposium on Applied perception in graphics and visualization*, pages 41–48, Tubingen, Germany, 2007. ACM.

[122] S. Zhai. *Human performance in six degree of freedom input control.* PhD thesis, University of Toronto, 1995.

[123] Y. Zhang, J. Pettré, Q. Peng, and S. Donikian. Data based steering of virtual human using a velocity-space approach. In *Proceedings of Motion in Games*, Lecture Notes in Computer Science 5884, pages 170–181. Springer, 2009.