Lionel Dominjon
Jérôme Perret
Anatole Lécuyer

Novel devices and interaction techniques for human-scale haptics

L. Dominjon (✉)
P&I Laboratory, ENSAM Angers
dominjon@ingenierium.com

J. Perret
Haption SA.
jerome.perret@haption.com

A. Lécuyer
Bunraku Project, INRIA/IRISA
anatole.lecuyer@irisa.fr

Abstract Today, most haptic devices are designed for desktop applications. Therefore, when dealing with human-scale virtual environments, haptic interaction usually remains limited to only a small portion of the virtual world. This paper describes novel technologies aimed at addressing this issue. We investigate hardware solutions, with the design of devices offering a large workspace, as well as software solutions, with the description of a new interaction paradigm called haptic hybrid control. First evaluations demonstrate encouraging results and suggest that such solutions are relevant, depending on the targeted application.

Keywords Virtual reality · Human-scale haptics · Force-feedback · Haptic device · Interaction technique · Haptic hybrid control

1 Introduction

Industrial applications of virtual reality technologies have been rapidly developing over the past decade. In many, such as ergonomics studies or assembly simulations on a virtual production line, the need for haptic interaction with large-scale virtual environments (VE) is very strong. Indeed, the interaction volume required to perform such tasks is generally quite large. Therefore, haptic devices used in such simulations should provide a workspace sufficient to match the whole task workspace, without limiting haptic interaction with the virtual environment. Haptic interaction involved in such simulations is generally referred to as *human-scale haptics* [3].

Unfortunately, most current force-feedback devices are designed for desktop applications and are not compatible with human-scale haptics. To address this issue, specific software interaction techniques have been developed such as adding a scaling factor [8] that amplifies the motion of the user in the virtual environment. These interaction techniques aim at allowing haptic interaction within a task workspace that is larger than the actual workspace of the device. But they are sometimes awkward and difficult to use.

This paper describes and discusses novel human-scale haptic devices, and also proposes novel interaction techniques that may be associated with these devices to artificially extend their workspace.

Therefore, the paper starts with an overview of related work focusing on the design of human-scale haptics. It is followed by a description of two novel haptic devices specially designed for human-scale haptic interaction: a large-scale cable-driven haptic device, and a redundant platform supporting a force-feedback arm. Then it describes a new paradigm for haptic interaction with large virtual environments called haptic hybrid control (HHC). Two recent interaction techniques based on the concept of HHC are then discussed: the bubble technique (for large translations), and haptic hybrid rotations (for large rotations). The paper

ends with a general discussion, a conclusion and a description of potential perspectives.

2 Related work

For many virtual reality applications, especially industrial ones, human-scale haptics has become highly desirable. This issue may be addressed in several ways.

First, some research has been on designing grounded devices featuring a large workspace. Mechanical architectures usually rely on large serial linkages, such as the device developed in the GROPE project [4]. This project is commonly known as the first project pioneering haptic feedback with a large-scale virtual environment. The Large Haptic Interface for Aeronautic Maintainability (LHIfAM) [2] is a more recent attempt at human-scale haptics with a grounded device. A major problem with such devices concerns visual occlusion issues when used in an immersive virtual environment such as CAVE-like displays. By fitting the user body, wearable exoskeleton architectures can minimize visual occlusions issues with the virtual objects. The L-Exos device from Frisoli et al. [9] is a four-degrees of freedom (dof) ground-based wearable exoskeleton bound to the user's arm and offering a workspace close to the one of the human arm (approx. 70%). According to the authors, the device is very effective for simulating the touch by hand of large objects or the manipulation within the whole workspace of the arm.

But the use of serial architectures featuring large linkages may sometimes impair the apparent stiffness of the system. In addition, backlash in gears may also result in an accuracy loss at the end-effector. In the case of L-Exos for example, the accuracy at the end-effector is approximately equal to 10 mm. Another problem with large linkage structures concerns the apparent inertia of the device. Thus, expensive lightweight materials featuring a high stiffness, such as carbon fibres, are often used in combination with the implementation in the control loop of inertia compensation techniques.

Therefore, other architectures have been investigated. Another possible architecture consists in replacing the classical serial structure with parallel free-moving cables for force transmission. Such interfaces are usually called tendon-driven (or cable-driven) devices. Bouguila and coworkers' Big-SPIDAR [3] is an example of an efficient tendon-driven device dedicated to human-scale haptics. The use of very small diameter cables avoids visual occlusion when used for example in CAVE-like displays. The stringed haptic workbench from Tarrin et al. [20] is another version adapted to haptic interaction on a two-screen Holobench. The system provides a haptic workspace that is large enough to cover almost the whole visual space of the Holobench without disturbing stereoscopic viewing. In theory, the workspace of such a system is only limited by the length of the cables. In practice, however, technical issues get in the way: the amount of cable that can be stored on a motor pulley without degrading the performance of force transmission and length measurement is indeed very limited. A second, more fundamental limitation appears when trying to build a 6-dof force-feedback device based on the SPIDAR, as the workspaces in rotation and translation are interdependent.

A second way of dealing with human-scale haptics consists in increasing the number of degrees of freedom of the devices. Such haptic devices may be divided into two categories, depending on the way they are controlled: mobile haptic interfaces [15] and redundant devices [16].

The former are usually set as desktop devices mounted on a mobile platform which allows for overcoming the stationary location dependence of grounded systems. The basic idea is that the mobile platform actively follows the movements of the operator, thus allowing the haptic device to always remain in the centre of its workspace. This allows the haptic device to render forces on the operator in a configuration close to its maximal structural stiffness. The main advantage of this approach is that the device may not be grounded. In this case, the resulting haptic workspace is not limited any longer by the workspace of the device, but by the volume the platform can move within. Thus, mobile haptic interfaces can provide very dextrous haptic feedback in a small but movable workspace. The main issue with such systems concerns the design of a global control scheme, which can drive the moving platform according to the user movements, while placing the haptic device in its optimal configuration for the best dynamic response. Barbagli et al. [1] have proposed two different mobile haptic interfaces: a general purpose interface based on the holonomic Nomad mobile robot and a PHANToM Premium 1.5, and a more limited interface based on a non-holonomic Pioneer mobile robot and a PHANToM Premium 1.5 as well.

The second category corresponds to redundant devices. Redundant haptic devices are interfaces featuring more actuators than degrees of freedom. The redundancy enables the designers to integrate shorter segments, thus keeping a fairly high force and stiffness. Of course, putting several actuators in series would eventually reduce the overall stiffness of the system, but this reduction would increase linearly with the number of actuators. The hyper redundant system VISHaRD10 (Virtual Scenario Haptic Rendering Device with 10 actuated dof) from Ueberle et al. [21] is an example of a highly redundant device and allows high force tasks in large volumes. VISHaRD10 provides a large cylindrical workspace of $\varnothing 1.7\,\text{m} \times 0.6\,\text{m}$, and a maximum payload of 7 kg which is sufficient to attach additional haptic displays.

Another approach to human-scale haptics relies on the use of software interaction techniques designed to artificially extend the workspace of the haptic device.

A first interaction technique is based on the concept of clutching [13]. It is inspired by the use of a classical 2D mouse. When reaching the limits of the mouse's workspace, the user may lift (de-clutch) the mouse, in order to put it down on a new location (clutch). This technique was implemented in haptic APIs (Application Programming Interface) such as in the VIRTUOSE API [11]. When the user reaches an uncomfortable posture with the force-feedback interface, he/she may de-clutch and freeze the virtual cursor in the VE by pressing a button. Then he/she can move the haptic device, reach a more comfortable position, and then clutch again by releasing the button to unfreeze the virtual cursor. From some point of view, angular clutching behaves like a natural action for the user. Indeed, clutching preserves directional compliance between the user's motion and the manipulated object [18]. This means that when the user rotates the device, the virtual object rotates in the same direction (around the same axis). However clutching is also frustrating and can notably impair user performance. Indeed, Zhai pointed out that de-clutching and re-clutching processes take time to complete, which may significantly decrease user performance [23]. In addition, performing clutching operations may make the interaction seem unnatural, reducing the feeling of engagement [10].

To reduce or cancel the need for clutching, non-isomorphic mappings may be used. This means that the user moves objects in ways that have no analog in the physical world [14]. Motion scaling is an example of non-isomorphic mapping. Fischer and Vance [8] used such mappings for their integration of a 3-dof PHANToM Premium 1.5 haptic interface in the C6 (a CAVE-like system). They used an amplification of the user's motion, i.e. a motion scaling between the haptic workspace and the VE [19]. The link between the two spaces is defined by a scaling factor equal to the ratio: *largest dimension of the workspace of the haptic device* to *largest dimension of the virtual environment*. Regarding 6-dof interaction, rotation scaling was described by Poupyrev et al. [18]. It is based on quaternion compositions, and may be organized in two categories: absolute and relative mappings. Absolute mapping consists in amplifying the global orientation of the device, whereas relative mapping consists in amplifying the rotation that occurred since the previous simulation step. The main issue with absolute mapping is probably the fact that it does not preserve directional compliance. Thus, unpredicted rotations of the object may occur. Nevertheless, relative mappings are not a robust solution for extending the angular range of a device because they do not preserve nulling compliance. This results in a drift of the initial orientation, which makes the scaling factor impossible to calibrate. On the contrary, absolute mappings do preserve nulling compliance, which makes it possible to calibrate the device to map a given angular range onto the device angular range.

3 Design of human-scale haptic devices

Human-scale haptics has been a major research topic for a few years, both as a technical challenge and as a market potential. The criteria for success include both theoretical efficiency and user-oriented (or customer-oriented) factors, which condition the "marketability" of the product. A commercial company would notably like to know under which conditions the customer is able to use the product eight hours a day. As an answer, we can suggest the following criteria for the acceptance of (human-scale) haptic devices:

- *Safety*. The device must be safe to use, and the user must feel safe when using it. The first point can be managed by the hardware, e.g. by integrating a "deadman" feature, such as a switch that cuts off the motor power supply when the device handle is released. The second point is more difficult to address, and is linked with the absence of vibrations, of disquieting noises, and of unwanted motion in general.
- *Robustness*. The device must be strong enough to sustain intensive mechanical stress due to repetitive use by different operators, some of them possibly untrained and thus not aware of its fragility.
- *Ease of use*. The device must be simple to operate, i.e. easy and fast to switch on and work with. Complex procedures like geometric calibration, system check, and power management should be avoided.

As a consequence, we decided to apply those criteria to two different hardware approaches. Thus, we describe hereafter our achievements and results in the development of a cable-driven SPIDAR-like interface, and in the integration of a redundant linear positioning axis for an existing haptic device.

3.1 Cable-driven haptic device

From a user's point of view, a preliminary analysis of the characteristics of classical cable-driven haptic interfaces such as the SPIDAR [3] could stress the following potential limitations:

- *Calibration*. Traditionally, the system calibration is carried out by placing the end-effector in a position perfectly known and repeatable. The movement needed to reach that position must be done with active force-feedback, so that the cables are pulled. However, since the current position is not known before the end of the calibration phase, no high-level safety measures can be implemented.

- *Winding quality*. The quality of the cable winding on the motor pulley affects greatly both the precision of the position measurement and the force-feedback. When two motor loops overlap each other, the apparent winding radius is changed, and thus the length measurement. Moreover, when power is switched on, such overlapping loops tend to slip back, causing sharp blows in the cables.
- *Unreeling*. When motor power is switched off, the cables unreel partially under their own weight and stiffness or due to the user movements, and often entangle themselves. The integration of brakes on the motors only solves part of the problem, and reduces the system flexibility.

Therefore, we have developed a new generation of actuators for cable-driven systems, called Inca (see Fig. 1). This hardware solution meets each of the three aforementioned items. Calibration at a fixed position is replaced by absolute encoders. Winding quality is improved by using special pulleys with an advancing spiral thread, so that the cable loops are stored regularly. Finally, unreeling could be prevented by a spring-based passive tension device.

These functions are all integrated into a modular architecture, so as not to impair the flexibility of the SPIDAR concept.

Each Inca module integrates (1) one motor with encoder, driving two spiral pulleys, and (2) one absolute position sensor, in the form of a multi-turn potentiometer. The first spiral pulley is for the transmission cable, while the second receives a tension cable, to be attached to a linear spring outside of the module. Thanks to this spring, it is possible to implement a dead-man switch at the end-effector (motor power is on only when the user holds the handle) without facing unreeling problems. The technical specifications of the Inca module are given in Table 1.

Fig. 1. Inca module

Table 1. Technical specifications of an Inca module

|	
Peak force on transmission cable	37.5 N
Max continuous force on transmission cable	12.5 N
Winding length of transmission cable	2.5 m
Force of tension cable	5 N
Winding length of tension cable	1.0 m
Length measurement resolution	7.10×10^{-6} m
Dimensions	$0.1 \times 0.2 \times 0.3$ m
Weight	2.5 kg

Fig. 2. Inca6D adapted to a visual workbench

Inca modules can be combined just as the motors of a classical SPIDAR system, either by four (Inca 3D) or by eight (Inca 6D). Figure 2 shows an example of a 6-dof haptic device built from eight Inca modules, and adapted to a holobench visual display.

Using the Inca technology, it is possible to build haptic interfaces with a large workspace in translation, up to 2 m to a side approximately. However, the workspace in rotation is rather small, and depends a lot on the geometry of the end-effector. Using a near-spherical manipulation tool with cables attached to the vertices of a tetrahedron, it is reduced to about 10° in each direction. With different geometries of the attachment points, it can be enlarged to some 30°, but at the cost of the translation workspace, which is then reduced.

3.2 Redundant moving platform

Our second development was inspired both by the redundant actuation approach and the moving platform solution. We decided to place a Virtuose 6D35-45 haptic interface on a 1-dof positioning system (see Fig. 4).

The workspace in translation of the standard Virtuose 6D35-45 haptic interface is a complex shape, with a maximum width of 1016 mm and a maximum height of 900 mm (see Fig. 3). But, as can be seen in Fig. 3, it is difficult to work in the extreme positions, and the *operational workspace* of the Virtuose is usually defined as a cube 450 mm in size.

Therefore, we developed a back driveable linear axis, based on a rotational actuator and a cable transmission (see Fig. 4). The technical specifications of the linear positioning axis are given in Table 2. Placed on the positioning system along the central axis, with a linear displacement of 2000 mm, the total workspace takes

Fig. 3. Perspective and top views of the workspace of the Virtuose 6D35-45


Fig. 4. Virtuose 6D35-45 on the linear positioning axis (courtesy of CEA/List)

the shape of a cylinder based on an elliptic section of 900×1016 mm, and most of its volume becomes operational.

From the point of view of the system dynamics, the linear axis is closer to a moving platform, as its inertia and level of friction are one order of magnitude higher than those of the haptic interface. But thanks to its back

Table 2. Technical specifications of linear positioning axis

Peak force	50.0 N
Max continuous force	25.0 N
Linear displacement	2.0 m
Position resolution	3×10^{-6} m

driveability, it can also be considered for the control laws as a redundant joint of the complete haptic system.

In our implementation, for simplicity reasons, we chose the redundant scheme. However, the moving platform scheme could also be selected as it presents several advantages, especially a high stiffness and a reduced inertia. Moreover, in future work, we could potentially reduce the speed of the linear axis, thus enhancing safety. The back driveability is a strong safety factor, as we can monitor the position errors and switch off and brake down the motors in case of collision, but the inertia alone is a real danger in itself, which can be reduced only by lowering the motion speed.

3.3 Control schemes

In many ways, the frontier between the redundant actuation approach and the moving platform approach is rather thin. From a system point of view, they are indeed identical. Truly, differences can be found in the hardware and/or software architecture, but they cannot be regarded as definitely discriminating. From our point of view, the difference lies in the control algorithms alone.

In the case of redundancy, all joints are part of the same control loop, and share the same control dynamics. The output force (and torque) is a result of the combined action of all actuators, and the redundancy is a purely internal issue, which can be seen as an optimisation problem, to minimize a given energy function. For example, the configuration of the redundant joints can be chosen so as to move away from joint limits. Here, all joints must share the same properties, such as back driveability if there is no force sensor in the loop.

In the other case, the haptic device and the moving platform are controlled by two different loops. A first, high bandwidth loop controls the haptic device as usual, but accepts a dynamic offset. A second control loop manages the motion of the platform at a slower pace, with the objective of maximizing the workspace of the haptic device, while providing the dynamic offset to the first loop. Here, the two subsystems do not have to share the same properties, and back driveability is not necessary for the moving platform.

The first control scheme is rather simple to implement, and gives quick results. Depending on the mechanical architecture, new issues can arise, such as internal oscillations around a singular configuration, or inertial forces due to internal movements. Nevertheless, the main drawback comes from the reduction of the global stiffness of the system, due to the combined contribution of all joints.

The second scheme calls for a clear separation of the control bandwidth of both subsystems, in order to avoid any dynamic coupling between them. Under that condition, the global system stiffness at high frequency is not hampered by the presence of the slow moving platform. However, the choice of control parameters for the moving platform often turns out to be difficult: if too slow, then it may not prevent the user from reaching the joint limits, and if too fast, then it may affect the behaviour of the haptic device.

One benefit of our redundant moving platform is that both control schemes (mobile platform vs. redundant control) may be implemented, which will make possible to rigorously compare them in the same situation.

3.4 Preliminary conclusion

We have presented the design of two novel force-feedback devices dedicated to human-scale haptics. The first informal tests that were carried out are encouraging as both devices seem to feature a workspace compatible with human-scale haptics, large admissible forces and a high structural stiffness. Future work is now necessary to formally evaluate these haptic devices. Our first informal tests also find that for some applications, the workspace of these devices, as large as it is, might not always be large enough.

4 Design of interaction techniques for human-scale haptics

The novel human-scale haptic devices presented above open a wide variety of new applications featuring haptic feedback. But sometimes, for a given application, the user is likely to need an even larger workspace than the one that can be provided by the device. To add more versatility to the use of our devices, we have also developed advanced

interaction techniques allowing the user to overcome the physical limitations of the devices.

4.1 Haptic hybrid control

The techniques we have developed are based on the concept of *haptic hybrid control* (HHC). This novel interaction paradigm involves a conventional hybrid control [12], which we propose to improve by simulating the behaviour of an elastic device during the rate control phases. In the following sections, we will briefly describe the fundamentals of standard hybrid control, its inherent limitations, and the concept of haptic hybrid control as well as implementation details.

4.1.1 Hybrid position/rate control

Hybrid control, as proposed by Hollis et al. [12], consists in dynamically adjusting the order of the transfer function according to the position of the device in space. A usual approach consists in dividing the space around the device into two zones: an inner zone, close to the neutral position of the device, and an outer zone, around the inner zone. When the device operates in the inner zone, it is supposed to be quite far from its mechanical stops, which means that interaction is not likely to be disturbed by any contact with the stops. Therefore, the virtual object may be safely manipulated with natural 1 : 1 motion mapping. On the contrary, when the device operates in the outer zone, it is supposed to be close to its physical stops and the manipulation transfer function must allow the user to keep moving the virtual object while limiting its hand motion. The order of this transfer function must therefore be greater than 1, i.e. either rate or acceleration control. Human performance analysis suggest that the use of a rate-control is more adapted [17].

4.1.2 Limitations of hybrid control

In the implementation proposed by Hollis et al. [12], position control and rate control of the manipulated object are both achieved with an isotonic interface. Nevertheless, according to Zhai and Milgram [22], rate control of an object using an isotonic device may be complex, due to the absence of the self-centering effect. Thus, the user has to always accurately remember the location of the neutral position of the device when he wants to stop the motion of the object he manipulates. Using an isotonic device for rate control requires a high level of concentration and might be perceived by the user as very frustrating.

4.1.3 Concept of haptic hybrid control

In order to address the aforementioned issue with hybrid control, we propose to mimic the duality present in the transfer function of the system (position or rate control)

on the resistance mode of the interface. Therefore, according to Zhai and Milgram [22], a device adapted to such a hybrid control should feature both an isotonic behaviour when operating in the position control zone (inner zone), and an elastic behaviour when in the rate control zone (outer zone). To our knowledge, such a device does not exist. Therefore, we propose to use the force-feedback of the haptic interface to simulate the behaviour of an elastic device when its tip is located in the outer zone, i.e. close to the mechanical stops. We call this extension haptic hybrid control.

4.1.4 Implementation of haptic hybrid control

The concept of haptic hybrid control may be implemented in several ways according to, for example, the kind of application it is supposed to be integrated into, or the expertise level of the user. We have identified four main parameters characterizing HHC:

- *Shape and size of the isotonic workspace.* This zone is a privileged natural interaction zone since inside this region the manipulated objects are position-controlled. Thus, it is often desirable to maximize this volume by using a shape accurately fitting the physical limitations of the workspace of the haptic device, especially in the case of human-scale haptics.
- *Visual representation of the boundary between the two control zones.* In order to help the user in identifying the current control mode (position or rate), the boundary between the rate control (elastic) zone and the position control (isotonic) zone may be visually displayed. This can help the user to anticipate the change in control mode when the device gets close to the boundary between the two zones. The user's viewpoint generally being outside the isotonic zone, it is recommended to display the boundary in semi-transparency: otherwise it would be impossible to see inside this zone. In addition, the semi-transparency provides the user with occlusion cues that are useful for the user to perceive the location of the isotonic workspace with respect to its surroundings. To avoid overloading of the visual scene, it is possible to make the transparency level variable according to the position of the device. For example, as the device gets closer to the boundary, the opacity increases, starting from fully transparent when the device operates close to its neutral position.
- *Profile of the force used to simulate the behaviour of an elastic device.* This force is supposed to pull the device back inside the isotonic zone when crossing the boundary of the rate-control zone. According to the definition of an elastic interface [22], possible force profiles are not restricted to Hooke's law. Nevertheless, simulated forces have to respect a few constraints. First, to ensure a correct force continuity and thus avoid force jumps when switching from one control zone to the other,

the magnitude of the generated force has to be null at the level of the boundary between the two control zones. Then, the force should increase as the device gets deeper into the rate control zone. Implicitly, this force also constrains the device inside a "safe" volume, preventing the user from ever reaching the mechanical stops of the device.

- *Transfer function of the velocity vector used when in rate control mode.* The velocity vector of the manipulated object may be constant, like implementations of rate control in most real time strategy video games, or dynamic, depending on different criteria or events in the simulation. Another possibility is to allow the user to directly control the velocity of the manipulated object. This can be achieved by basing the velocity of the object on the position of the device. For instance, the velocity vector could increase with the penetration depth of the device in the rate control zone. In order to avoid discontinuities in the object motion when switching from one control mode to the other, it is advised to set a null velocity when the device is located on the boundary between the two zones. Saturating the velocity at a maximum value might also be desirable to avoid involuntary overshootings.

4.1.5 Preliminary conclusion

We have presented the concept of haptic hybrid control, which is based on a hybrid position and rate control, coupled with the simulation of the behaviour of an elastic device during the rate control phases. In the two following parts, we describe two applications of the HHC: one interaction technique called the bubble technique, aiming at enlarging the control of the *position* of an object, and another interaction technique called haptic hybrid rotations to improve *orientation* control.

4.2 Two applications of haptic hybrid control

4.2.1 The bubble technique

The bubble technique [22] was designed to propose a solution to the problem of haptic interaction with a virtual environment that is larger than the workspace of the device. The name of the technique is inspired by the visual representation of the boundary of the isotonic workspace of the device, which actually looks like a soap bubble.

Characterizing the isotonic workspace of a device is pretty simple. Indeed, any volume in 3D space may bind this workspace. In order to provide an isotropic interaction, we propose to use a spherical volume centred on the device neutral position. The radius of the sphere is chosen as large as possible to maximize the isotonic workspace of the device. Meanwhile, it must also ensure that the device would never reach its mechanical stops.

Fig. 5. The Bubble Technique: visual representation of the bubble

In accordance with Sect. 4.1.4, the boundary between the position control zone and the rate control zone is visually displayed as a semi-transparent sphere, looking like a bubble (see Fig. 5). To avoid visual overload of the virtual environment when not necessary, we propose to decrease the opacity of the bubble as the device gets closer to its neutral position.

The elastic behaviour of the device when in rate control mode is achieved by setting a force that is proportional to the penetration depth of the tip of the device in the rate control zone, oriented towards the centre of the sphere. Thanks to the force-feedback of the device, the user can have the feeling of sliding on the inner surface of the sphere.

Regarding the last parameter stated in Sect. 4.1.4, the transfer function of the rate control, we propose to use a non-linear control law that is supposed to allow precise and low amplitude motion while keeping the ability to move the manipulated object very quickly. Informal tests resulted in the choice of a third degree polynomial control law.

An evaluation of the bubble technique [23] was carried out to compare it with two other common interaction techniques: the scaling technique, consisting in amplifying the motion of the user, and the indexing technique, consisting in providing the user with a clutching mechanism. Participants were asked to paint a virtual model as fast and as precisely as possible inside a CAVE-like display, using a desktop haptic device. The results showed that the bubble technique enabled both the quickest and the most precise paintings. It was also the most appreciated technique.

4.2.2 Haptic hybrid rotations

Haptic hybrid rotations [6] were designed to overcome the hardware angular limitations of force-feedback devices. The problem becomes particularly important in VR applications where the user needs a full control over the object's orientation, such as assembly simulation. Unlike the translation space, finding a metaphor to represent the angular isotonic workspace of a device in an intuitive manner is a complex problem. Indeed, rotation space is not a vector space, but a closed curved space that may be represented by a 4D sphere. Therefore, representing a rotation space in 3D is not trivial. Thus, we propose to decouple the three rotational degrees of freedom into a roll component (1-dof) and a pitch and yaw component (2-dof) and to handle them separately. Regarding the roll component, we propose to bind this dof with two virtual springs (see Fig. 5). The roll isotonic workspace is thus limited by the two springs. Regarding the other two components, we propose to use a virtual cone to bind the pitch and yaw isotonic workspace (see Fig. 5).

The bounding cone related to the pitch and yaw component is visually displayed as a semi-transparent cone. Thus it provides the user with a good visibility of the isotonic workspace in every condition. As for the bubble technique, we propose to make the opacity of the cone variable according to the orientation of the device. Regarding the roll springs, they are displayed as two curved coil springs (see Fig. 5).

The simulation of an angular elastic device is achieved by applying a torque as soon as the device crosses the surface of the cone or stresses the roll springs. The torque we propose to generate is proportional to the angular penetration depth of the device inside the rate-control zone.

Regarding the transfer function of the rate control, the challenge is to try to automatically continue the motion initiated by the user, in order to provide the impression of a smooth and continuous rotation of the manipulated object. For that purpose, regarding the roll component, we propose to consider a rotation axis that is co-linear with the main axis of the device. Regarding the pitch and yaw components, we suggest to consider a rotation axis that is orthogonal to both the main axis of the device and the longitudinal axis of the cone. Concerning the magnitude of the velocity vector, we also propose a simple control law proportional to the penetration depth of the device in the rate control zone.

Finally, we carried out an experimental evaluation [6] consisting in an assembly simulation task. The evaluation showed that, in our experimental conditions, this technique was up to twice as fast as the other two compared techniques (scaling and indexing techniques, adapted to rotation), without any significant loss of precision. Haptic hybrid rotations were also well appreciated by the participants, mainly regarding their global appreciation of the different techniques.

4.3 Preliminary conclusion

We have proposed two applications of haptic hybrid control: the bubble technique, to control the position of an object over a large VE, and haptic hybrid rotations

Fig. 6. Haptic Hybrid Rotations: visual representation of the handle of the device with the virtual roll springs and the virtual cone

to control its orientation. Both interaction techniques aim at allowing an accurate and intuitive interaction with a human-scale VE, while cancelling the problem

of the physical limitations of the haptic device being used. Even if the first usability evaluations we ran concerning these new interaction techniques suggest very encouraging results (see [6] and [7] for more details), further evaluations are required to complete their evaluation

5 Conclusion

In this paper, we have described the design of two human-scale haptic devices based on different architectures: a parallel cable-driven device, Inca, and a redundant moving platform associated with a Virtuose 6D35-45 force-feedback device. We have also proposed a novel interaction paradigm, the haptic hybrid control, which aims at enhancing the versatility of human-scale devices by allowing the user to overcome their physical limitations.

This way, the combination of novel human-scale haptic devices and appropriate interaction techniques may open a wide variety of new applications featuring haptic feedback.

References

1. Barbagli, F., Formaglio, A., Franzini, M., Giannitrapani, A., Prattichizzo, D.: An experimental study of the limitations of mobile haptic interfaces. In: Experimental Robotics IX – STAR (Springer Tracts on Advanced Robotics). Springer, Berlin Heidelberg New York (2004)
2. Borro, D., Savall, J., Amundarain, A., Gil, J.J., Garcia-Alonso, A., Matey, L.: A large haptic device for aircraft engine maintenance. IEEE Comput. Graph. Appl. **24**, 70–74 (2004)
3. Bouguila, L., Ishii, M., Sato, M.: Scalable SPIDAR: A haptic interface for human-scale virtual environments. In: First International Workshop on Haptic Human-Computer Interaction (2001)
4. Brooks, F., Ouh-Young, M., Batter, J., Kilpatrick, J.: Project GROPE – haptic displays for scientific visualisation. In: Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'90) (1990)
5. Dominjon, L., Lécuyer, A., Burkhardt, J.-M., Richir, S.: Haptic hybrid rotations: Overcoming hardware angular limitations of force-feedback devices. In: Proceedings of the IEEE International Conference on Virtual Reality (IEEE VR), Washington, US (2006)
6. Dominjon, L., Lécuyer, A., Burkhardt, J.-M., Andrade-Barroso, G., Richir, S.: The "bubble" technique: interacting with large virtual environments using haptic devices with limited workspace. In: Proceedings of World Haptics Conference (joint Eurohaptics Conference and Haptics Symposium), Pisa, Italy (2005)
7. Dominjon, L., Lécuyer, A., Burkhardt, J.-M., Richir, S.: A comparison of three techniques to interact in large virtual environments using haptic devices with limited workspace. Lect. Notes Comput. Sci. **4035**, 288–299 (2006)
8. Fischer, A., Vance, J.: PHANToM haptic device implemented in a projection screen virtual environment. In: Proceedings of the 7th International Immersive Projection Technologies Workshop and 9th Eurographics Workshop on Virtual Environments, Zurich, Switzerland (2003)
9. Frisoli, A., Rocchi, F., Marcheschi, S., Dettori, A., Salsedo, F., Bergamasco, M.: A new force-feedback arm exoskeleton for haptic interaction in virtual environments. In: Proceedings of the World Haptics Conference (joint Eurohaptics Conference and Haptics Symposium), Pisa, Italy (2005)
10. Hand, C.: Survey of 3-D interaction techniques. Comput. Graph. Forum **16**, 269–281 (1997)
11. HAPTION: VIRTUOSE API V2.0 Programming Manual (2004)
12. Hollis, R.L., Salcudean, S.E.: Lorentz levitation technology: a new approach to fine motion robotics, teleoperation, haptic interfaces, and vibration isolation. In: Proceedings of the International Symposium for Robotics Research (1993)
13. Johnsen, E., Corliss, W.: Human Factors Applications in Teleoperator Design and Operation. Wiley, New York (1971)
14. Mine, M.: Virtual Environment Interaction Techniques. University of North Carolina TR95-018 (1995)
15. Nitzsche, N., Hanebeck, U., Schmidt, G.: Design issues of mobile haptic devices. J. Robot. Syst. **20**, 549–556 (2003)
16. Patel, R., Shadpey, F.: Control of Redundant Robot Manipulators: Theory and Experiments, vol. 316. Springer, Berlin Heidelberg New York (2005)
17. Poulton, E.C.: Tracking Skill and Manual Control. Academic, New York (1974)
18. Poupyrev, I., Weghorst, S., Fels, S.: Non-isomorphic 3D Rotational Techniques. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,

The Hague, The Netherlands, pp. 540–547. ACM, New York (2000)

19. Preusche, C., Hirzinger, G.: Scaling issues for teleoperation. In: Proceedings of the 5th PHANToM Users Group Workshop, Aspen, Colorado (2000)

20. Tarrin, N., Coquillart, S., Hasegawa, S., Bouguila, L., Sato, M.: The Stringed HapticWorkbench: a New HapticWorkbench Solution. In: Proceedings of EUROGRAPHICS (2003)

21. Ueberle, M., Mock, N., Peer, A., Michas, C., Buss, M.: Design and control concepts of a hyper redundant haptic interface for interaction with virtual environments. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Workshop on Touch and Haptics, Sendai, Japan (2004)

22. Zhai, S., Milgram, P.: Human performance evaluation of manipulation schemes in virtual environments. In: Proceedings of VRAIS'93: 1st IEEE Virtual Reality Annual International Symposium, Seattle, WA, (1993)

23. Zhai, S., Milgram, P., Buxton, W.: The influence of muscle groups on performance of multiple degree-of-freedom input. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Common Ground Vancouver, British Columbia, Canada, pp. 308–315. ACM, New York (1996)

LIONEL DOMINJON received the Engineering Diploma from IFMA, the French Institute for Advanced Mechanical Engineering, France, in 2002. He received his Ph.D. degree in Computer Science from University of Angers, France, in 2006. Then he became a research and development engineer at CLARTE, a center for applied research in virtual reality in 2006.

JÉRÔME PERRET received his Ph.D. in Robotics & Artificial Intelligence in 1995 at the LAAS/CNRS, Toulouse, France. He worked for five years in the company Cybernetix, France, first as research engineer and then as group leader. He created the company Haption in 2001, together with a former research engineer at the French Nuclear Agency (CEA). As General Manager of Haption, he is responsible for the management of research activities and for the international business development.

ANATOLE LÉCUYER has been a Research Scientist at INRIA since 2002. He graduated from Ecole Centrale Lille, France, in 1996, and received his M.S. in 1996 from University of Lille 1, France. He received his Ph.D. in Computer Science in 2001 from University of Paris XI, France. He spent one year in the French Commission for Atomic Energy, Fontenay-aux-Roses, France, as a research engineer. At IRISA, Rennes, France, he is a member of the BUNRAKU research team. His main research interests include: virtual reality, 3D interaction, haptic feedback, pseudo-haptic feedback and brain-computer interfaces.