

3D Interaction Techniques for Bimanual Haptics in Virtual Environments

Anthony Talvas, Maud Marchal, Gabriel Cirio, Anatole Lécuyer

Abstract Bimanual haptics is a specific kind of multi-finger interaction that focuses on the use of both hands simultaneously. Several haptic devices permit this form of interaction, but they are subject to a certain number of limitations for interacting with virtual environments (VEs), notably workspace size issues and manipulation difficulties with single-point interfaces. Interaction techniques exist to overcome these limitations and allow users to perform specific two-handed tasks, such as the bimanual exploration of large VEs and grasping of virtual objects. These interaction techniques will be reviewed in this chapter.

1 Introduction

Bimanual haptics implies the use of both hands for haptic interaction with virtual or remote environments. It is more specific than the general *multi-touch* or *multi-finger* classes of interaction in that it does not include interaction with fingers of the same hand only, but requires at least one finger of each hand. What makes two-handed interaction more specifically interesting to study is that most tasks executed in our daily lives are intrinsically bimanual [14], from simple tasks such as striking a match to more complex ones such as playing an instrument. As such, two-handed interaction feels more natural to use than the one-handed kind.

This chapter focuses specifically on bimanual haptic interaction with virtual environments (VEs), which is a domain that is slowly emerging while already bearing promising applications. In the medical field, it can be used in surgery training, notably for minimally invasive surgery as illustrated by the bimanual surgical simulator interface [29] and the *da Vinci* Surgical System Simulator [18]. Another medical application is rehabilitation, as shown by a bimanual haptic desktop platform for upper-limb post-stroke rehabilitation [17]. Medicine, however, is not the only target

IRISA/INRIA Rennes, France, e-mail: {anthony.talvas,maud.marchal,anatole.lecuyer}@inria.fr

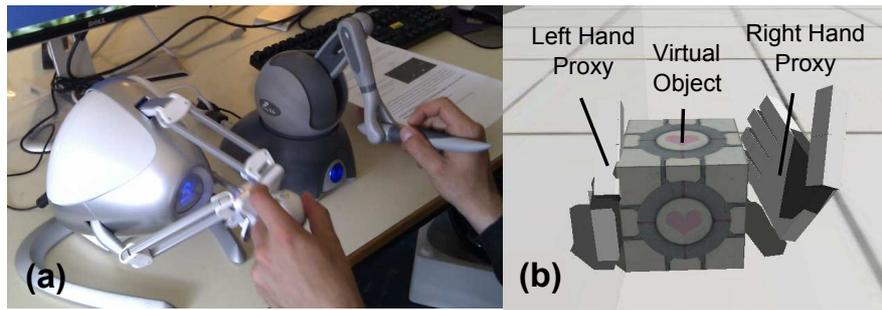


Fig. 1 Example of bimanual haptic interaction with a virtual environment. (a) Bimanual haptic setup made of two single-point devices. (b) Proxies and virtual object.

field: non-medical applications include industrial prototyping [15] and 3D modeling software [9].

Several haptic devices are suitable for bimanual interaction, and can be subdivided into two categories: single-point and multi-finger interfaces. While the latter category allows to handle virtual objects in a similar fashion to actual hands, the former makes the task more challenging as single-point devices are represented in the virtual environment by a couple of god-objects or proxies, which is the strict minimum to allow grasping of an object (Figure 1). Notably, a grasped object tends to slip from virtual hands if the contacts between proxies and objects are not strongly maintained, which shows the need for techniques to maintain more easily these contacts during grasping.

Another issue is that, regardless of which category they belong to, most of the bimanual haptic interfaces have small workspaces, which is a strong limitation when a user wants to carry out tasks as simple as picking an object and moving it around in a large VE. Several workspace extension techniques were proposed to address this issue, notably the *Bubble* technique [6] which proved to be suitable for simultaneous grasping of objects and exploration in a VE with a bimanual whole-hand haptic interface [21]. More recently, novel metaphors adapted this technique to any couple of 3DoF single-point haptic interfaces [26], also allowing grasping and navigation at the same time.

2 Related work

The hardware and software solutions that allow bimanual haptic interaction with VEs are currently subject to some limitations, which will be overviewed in this section. Interaction techniques and metaphors, which combine hardware, software and knowledge of human cognition to allow a user to perform a given task, can notably help to overcome these limitations. Very few genuinely two-handed haptic interactive techniques were developed to this date, but some techniques from unimanual

haptic interaction and two-handed interaction in 3D environments, not necessarily haptic, can be adapted for a bimanual haptic case. This section will provide a global overview of the different interactive techniques suitable for bimanual haptic interaction with VEs.

2.1 Hardware and Software Limitations for Bimanual Haptics

Several haptic devices allow bimanual interaction with VEs. However, they share the same drawback of having workspaces fairly limited in size, thus preventing the interaction with large VEs without using techniques to increase these workspaces. This problem can be handled at the hardware layer, by improving the existing hardware or by using additional devices.

A straightforward and general approach consists in increasing the workspace provided by the devices, either by using large serial linkages, string-based devices with large frames, or redundant degrees of freedom (DoF) [7]. For unimanual interaction, human-scale devices could be developed using that approach. However, as far as bimanual devices are concerned, such scales could not be obtained and the highest reach that was achieved so far was that of human arms, in the cases of the DLR interface [15] and the Haptic Workstation [21].

Another solution for providing larger workspaces is the use of mobile haptic interfaces, *i.e.* haptic devices fixed on a mobile robot. Bimanual examples of these are the Mobile Haptic Grasper [4] and VISHARD7-based mobile interface [22]. However, even though they potentially provide an infinite planar workspace [11], these devices are still limited in vertical reach.

Other hardware approaches address the workspace issue by providing additional DoF to the user to handle navigation in the VE. For instance, a 3DoF foot pedal was used for controlling the motion of a two-armed robot in a remote environment [23]. The latter approach may be the only one to provide an infinite workspace in all directions, but is fully dependant on the availability of such pieces of hardware.

As far as manipulation capabilities of haptic devices are concerned, whether unimanual or bimanual, those devices can be divided into two main classes. The first category is that of single-point interfaces, such as the SPIDAR G&G [19], the DLR bimanual haptic interface [15] or the more widespread PHANToM series. The second category is that of multi-fingered interfaces, such as the MasterFinger-2 [12], SPIDAR-8 [28] and Bimanual HIRO [8], which enable interaction through 4, 8 and 10 fingertips respectively. Within this category of multi-finger devices, the Haptic Workstation is a special case, not only providing interaction through the fingertips but also through the palms of both hands [21].

Multi-finger interfaces are well suited for grasping, as more contact points with objects means better handling of the said objects: it was shown that 3 contact points with friction, or 7 contact points without friction, were the minimum requirements for form closure of an object, *i.e.* the ability to fully restrain an object [1]. However, manipulating virtual objects with two single-point haptic interfaces is a different

matter. It was shown that 2 contact points could allow full manipulation of an object, but only if the contacts are soft-finger, meaning being able to support moments around the contact normal. Even so, handling virtual objects with single-point devices is very much like attempting to handle an object with the fingertips of index fingers, and as such a challenging task.

Concerning the software layers that link the aforementioned devices to the VEs, there are various haptic APIs that allow to do so, which can be divided into two categories. Device-specific APIs usually grant low-level access to one model or series of interfaces, generally allowing the use of the same device in a dual way. However, it is not possible to conjointly use devices of two different manufacturers using only one of these APIs, but it is possible to combine the respective APIs of the two devices. For instance, HDAL SDK [20] and OpenHaptics [24] can be used to operate conjointly a Novint Falcon and a PHANToM device.

This process can be bothersome, though, and thus generic APIs were developed to support a wider range of devices, not limited to a single manufacturer. Some of them have clearly shown their bimanual capabilities, such as H3DAPI [25], which was used for a bimanual haptic 3D modeling software [9], and Haptik Library [5], which was part of the software architecture of the Mobile Haptic Grasper [4]. Both cases used twice the same device, but these APIs can also be used with two different devices ; notably, examples shown later in this chapter used Haptik Library to handle a Novint Falcon and a PHANToM Omni.

However, while these APIs support the use of two devices at the same time, they do not truly support bimanual capabilities *per se*. There are for instance no solutions to reduce the impact on the user of different workspace sizes, output forces, or degrees of freedom in each hand. Similarly, while methods exist to prevent collision between the hands of the user and the two devices [8], they are not yet implemented in these APIs either.

2.2 Interaction Techniques for Bimanual Haptics

The hardware approaches for increasing the workspaces of bimanual haptic hardware do effectively allow the exploration of large VEs to some extent. However, on the one hand, most of these techniques do not increase those workspaces infinitely, and on the other hand, such devices are not necessarily widespread. This leads to the use of interaction techniques for handling the extension of these workspaces. They have the advantage of being generic and applicable to any haptic device available to the user with no further requirements, although the majority of currently existing techniques are not specific to two-handed haptics.

A first technique consists in applying a scaling factor to match the real workspace provided by the haptic devices with a virtual volume defined in the VE [10], although reducing the accuracy of motions in the virtual space. Another approach is the *clutching* technique, which consists in holding down a button to temporarily interrupt the coupling between the device and the proxy while the user recenters the

device. The *Dual Shell* method is an extension of this technique, that automatically handles the clutching when predefined boundaries are reached, without requiring the potentially counterintuitive manipulation of a button [16]. The use of rate control was also proposed to control the velocity of the virtual proxy through the position of the haptic device [30]. This technique infinitely increases the workspace in all directions, although seeming to be far from intuitive.

A last interaction technique is the *Bubble* metaphor, which is originally a unimanual haptic workspace extension technique [6]. It uses direct control of the position of the proxy in the VE while the device remains inside predefined spherical boundaries (defined as the *bubble*). When the device moves outside this *bubble*, the latter is translated throughout the virtual environment, with a small elastic force that gives the feel of the surface of the *bubble* to the user. For a unimanual painting task, a user experiment showed that this technique was faster, more accurate and more appreciated than the scaling factor and clutching techniques.

The bimanual case was not yet thoroughly investigated with the *Bubble* technique, though it was used for bimanual interaction with complex VEs through the Haptic Workstation [21]. This implementation allowed users to translate and rotate the camera by moving both hands outside the bubble in the same direction. It showed to be efficient for simultaneous navigation and manipulation with this specific device, as it allows interaction with the palms and fingers of both hands. However, it remains difficult to use with single-point interfaces, as picked objects tend to be frequently dropped during the translations of the virtual workspace through rate control. This is particularly true when using two different interfaces with physical workspaces of different size and shape, a scenario that has been scarcely studied.

Aside the question of how to explore large VEs with bimanual interfaces, next comes the matter of manipulating virtual objects with these devices. Bimanual haptic manipulation techniques vary greatly depending on the task that is to be executed. For instance, the multi-modal mesh manipulation system [9], a bimanual 3D modeling software, uses a tool-object metaphor. It can assign a grabbing task to the non-dominant hand and different manipulation tools to the dominant hand, or a tool in each hand for bimanual manipulation of the mesh (stretching, folding, tearing, *etc.*). Similarly, the Bimanual Haptic Simulator for medical training [27] assigns a palpation task to the non-dominant hand and a needle insertion task to the dominant hand.

However, while these techniques work for a specific task, the question of how to switch between different tasks and techniques remains open. In the context of 3D interaction techniques, different methods were proposed to apply transitions between different tools, or subtasks, proposed to the user in the case of the Responsive Workbench [3]. They defined explicit transitions, like picking up a tool in a toolbox, with a default behaviour when no tool has been picked up yet, that should be specific of the application. Another example of explicit transition, more practical in the sense that it does not require several movements between the workspace and the toolbox, is the use of gestures with the hands. Implicit transitions were also defined, in which a switch from an subtask to another occurs in a seamless way, almost imperceptible to the user. An example of this is the switch from a unimanual grabbing technique

to a bimanual grab-and-twirl technique, that occurs naturally as the user reaches in with the other hand to help the manipulation.

There is a more general two-handed manipulation task that several studies focused on, which is the grasping of arbitrary virtual objects. A technique was proposed to detect when a user attempts to grasp an object, based on a study on the segmentation of a grasping task with multi-finger interfaces [13]. The study distinguished three major steps with specific forces applied on the grasped object: approach (no force applied on the object), gripping (a horizontal force being applied), and lifting (a vertical component being added). The information on the forces applied by each finger on a virtual object can thus allow the controller to detect grasping and simulate it accordingly.

Beyond detecting grasping, a straightforward interaction technique for handling bimanual multi-finger manipulation of virtual objects is the use of virtual hands, such as spring-damper hands, which proved to be suitable to interact with the objects of a virtual house [21]. However, as mentioned earlier, manipulating objects with two single-point interfaces is intrinsically more challenging: to the best of our knowledge, no manipulation technique tackles this problem.

3 Interaction Techniques for Improving Bimanual Interaction with Single-Point Haptic Interfaces

Bimanual interaction with VEs using single-point haptic interfaces has been scarcely studied to this date. Thus a set of interaction techniques was proposed for improving the exploration of large VEs with these interfaces, as well as the manipulation of objects with two 3DoF haptic devices represented by simple proxies [26].

This section will present these techniques, which include two haptic exploration techniques, a grasping detection method, and two haptic manipulation techniques. The exploration techniques are the *double bubble*, which allows free motion with both hands in a VE, and a viewport adaptation method that maintains both virtual proxies on screen. The grasping detection method allows to determine when a user attempts to effectively pick an object between both hands. Finally, the manipulation techniques are the *magnetic pinch*, which uses a simulated spring to keep the virtual proxies from dropping a picked object, and the *joint control*, which solves issues related to different control modes between the two hands.

3.1 Double Bubble

In the *double bubble* technique, the workspace of each haptic device is defined by two areas, each associated to a control mode (Figure 2). An inner area controls the proxy directly in position, and an outer area, starting at the boundaries of the

inner area and extending up to the physical limits of the device, controls the virtual workspace in speed within the VE (Figure 3).



Fig. 2 Approximate *bubble* and physical workspace sizes of a Novint Falcon and a PHANToM Omni. Inner rectangles represent the *bubbles*, while outer polygons represent the physical workspaces.

Besides using two interfaces instead of one, two major differences separate the *double bubble* from the previously mentioned *bubble* technique. The first difference is the use of a rectangular parallelepiped for the boundaries of the *bubbles* instead of a sphere, to better fit the physical workspaces of the devices. We can notably think of PHANToM devices which have a higher width than their height or depth. The second difference is the presence of a visual feedback added to the haptic feedback when the devices leave the boundaries, in the form of a trail behind the rate-controlled proxies (Figure 3).

An issue that was observed with several users attempting to use the technique was that they got confused whenever the right proxy was on the left of the screen and vice versa. A way to prevent this is to simulate an invisible “separation plane” that prevents the centers of the bubbles from crossing, by negating the horizontal component of the *bubble* velocities when such situation is about to happen. Given the left *bubble* of center $\mathbf{l} = (l_x, l_y, l_z)$ and its displacement at the next simulation step $\mathbf{d} = (d_x, d_y, d_z)$, and the right *bubble* of center $\mathbf{r} = (r_x, r_y, r_z)$, the constraint is applied following Equation (1). The same constraint is applied for the right bubble.

$$\mathbf{d} \leftarrow (r_x - l_x, d_y, d_z) \quad \text{if} \quad l_x + d_x > r_x. \quad (1)$$

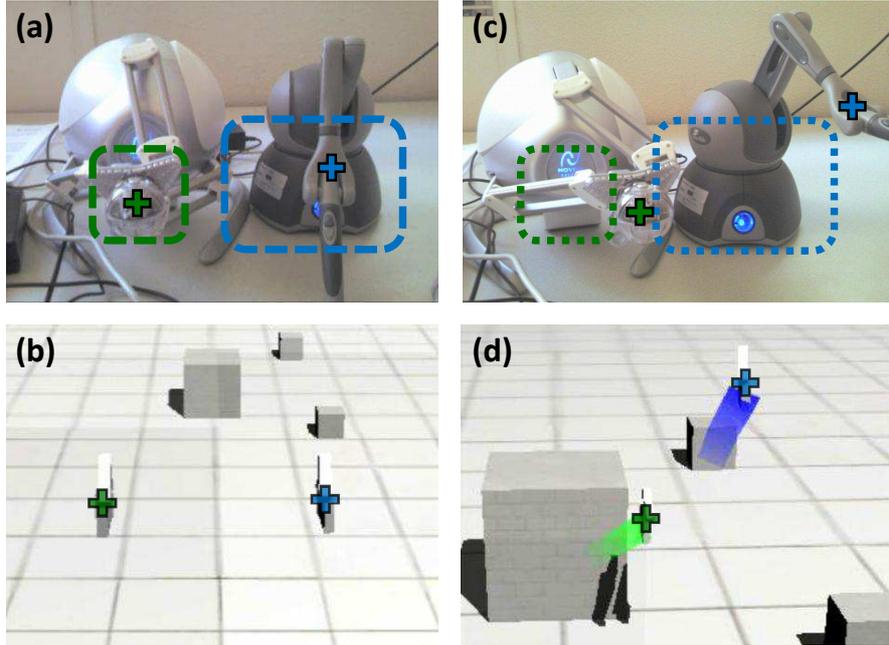


Fig. 3 Control modes of the *double bubble*. (a-b) Devices inside the *bubbles* : position control. (c-d) Devices outside the *bubbles* : rate control.

3.2 Viewport Adaptation

Since each device is attached to a *bubble* independent from the other, a method is required to keep both proxies on the screen, as they can move infinitely in completely opposite directions. Thus, a method was developed to ensure that both *bubbles* remain in view no matter the distance between them (Figure 5). This is accomplished by setting the distance of the camera to the center of the scene to a value proportional to the distance between the leftmost border of the left *bubble* and the rightmost border of the right *bubble*, plus an arbitrary margin (Figure 4). Given the left *bubble* of center $\mathbf{l} = (l_x, l_y, l_z)$ and width w_l , and the right *bubble* of center $\mathbf{r} = (r_x, r_y, r_z)$ and width w_r , the position of the camera is computed following Equations (2-4).

The center of scene \mathbf{s} is first computed from both workspace centers following:

$$\mathbf{s} = \left(\frac{l_x + r_x}{2}, \frac{l_y + r_y}{2}, \max(l_z, r_z) \right). \quad (2)$$

The width of the displayed scene w_s is then computed from the widths of both workspaces w_l and w_r , as well as an arbitrary margin m that ensures that the virtual workspace boundaries do not leave the borders of the screen:

$$w_s = \sqrt{(r_x - l_x)^2 + (r_y - l_y)^2} + w_l/2 + w_r/2 + 2m. \quad (3)$$

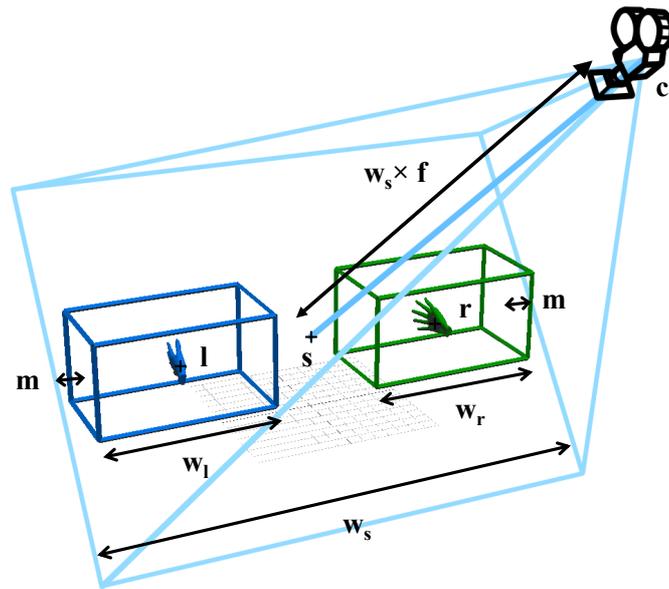


Fig. 4 Computation of the camera position for viewport adaptation.

Finally, the position of the camera c is computed following:

$$c = s + w_s \times d \times a. \quad (4)$$

where d is a scalar that depends on the camera field of view, and a is an arbitrary vector that determines the angle from which the scene is displayed.

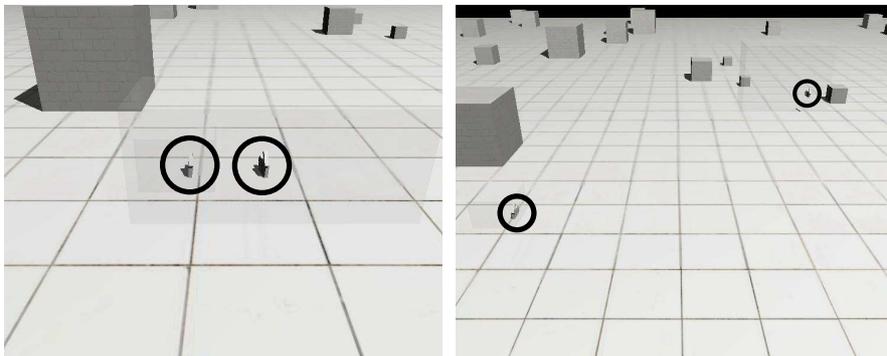


Fig. 5 Result of the viewport adaptation with different relative positioning of the proxies (circled).

A limitation of this method is that it does not allow rotations of the viewport. A way to handle these rotations is to use the “separation plane” mentioned in the pre-

vious section, by translating the x component of the displacement of the *bubble* into an angular velocity that will be applied to the vector \mathbf{a} (Figure 6). The displacement of both proxies can be used, making the viewport rotation faster if both hands push in opposite directions simultaneously.

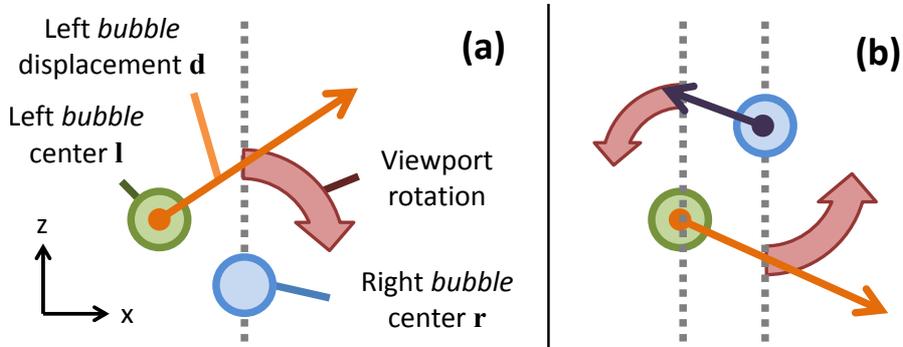


Fig. 6 Rotation of the viewport by “pushing” on the separation plane with the proxies. (a) One-handed case. (b) Two-handed case.

3.3 Grasping Detection

A grasping detection method is required in order to detect when a user is actually attempting to pick an object and not simply touching it. Three conditions are considered to determine whether both hands are grasping an object or not, according to the contact normals, the contact forces, and the relative position of both hands (Figure 7):

1. The angle between the contact normals must be under a certain threshold.
2. Both contact forces must exceed a threshold in order to discriminate simple contacts with an object from a true intent of grasping the object.
3. Two cylinders projected from both proxies following the contact normal must intersect. Their radii match the approximate sizes of the proxies by default, but can be tuned to make the detection more or less sensitive.

3.4 Magnetic Pinch

Once a grasping situation is detected, the *magnetic pinch* takes effect, which “magnetizes” both hands to the picked object to prevent unintentional drops from happening. A visual feedback is also added to the haptic feedback in the form of red

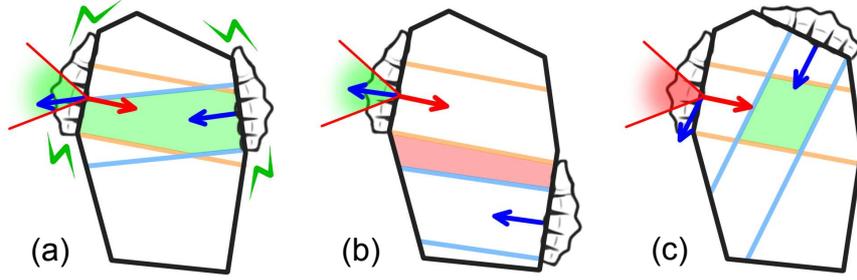
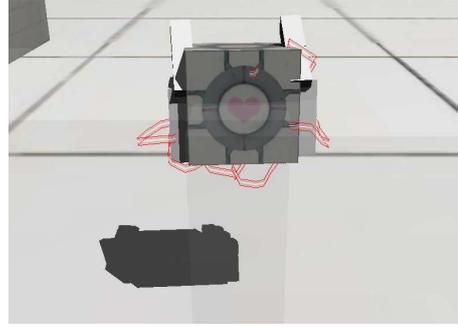


Fig. 7 Different cases of dual contact with a virtual object, one being appropriate for grasping and two not being as such: (a) Normals nearly colinear and hands face-to-face, (b) Hands not in front of each other, (c) Normals far from colinearity.

bolts, emphasizing the activation of the technique to the user (Figure 8). Two implementations were considered for the magnetic pinch: one based on springs and the other based on constraints.

Fig. 8 Visual feedback of the *magnetic pinch*, symbolized by red bolts between virtual proxies and picked object.



3.4.1 Spring-based

The first version simulates a spring pulling both hands towards the picked object. For each haptic device, a force \mathbf{F}_h is generated following:

$$\mathbf{F}_h = -k_h \times \left(1 - \frac{g_s}{\|\mathbf{o} - \mathbf{p}\|}\right) \times (\mathbf{o} - \mathbf{p}). \quad (5)$$

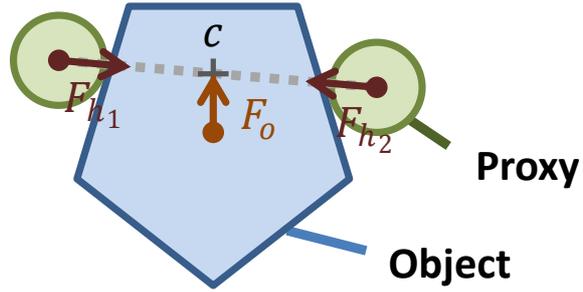
\mathbf{p} is the position of the first interface, \mathbf{o} is the position of the second interface, g_s is the size of the grasped object (the distance between the two contact points when the grasping is initiated), and k_h is the stiffness of the spring. The spring is

removed as soon as the user gives enough force to end the contact of the hands with the object, hence dropping it.

Additionally, the position of the grasped object \mathbf{g}_p can be constrained to the central point between the positions of the two virtual proxies \mathbf{l} and \mathbf{r} , further reducing the risk of unwanted drops. For this, another spring is used, of stiffness k_o with a force \mathbf{F}_o following:

$$\mathbf{F}_o = -k_o \times \left(\frac{\mathbf{l} + \mathbf{r}}{2} - \mathbf{g}_p \right). \quad (6)$$

Fig. 9 Forces applied by the spring-based approach of the *magnetic pinch*. F_{h_1} and F_{h_2} are the forces applied on the centers of mass of both proxies to pull them towards each other. F_o is the force applied on the center of mass of the picked object to pull it towards c , the middle point between both virtual proxies.



This approach successfully stabilizes the grasping of light virtual objects. However, lifting heavier objects can still be problematic with this approach, as well as correctly handling the rotations of the picked object.

3.4.2 Constraint-based

The second approach for the *magnetic pinch* constrains the relative position and orientation of both proxies relative to the picked object. Both constraints are removed as soon as at least one proxy provides an outwards force whose value exceeds an arbitrary threshold. This threshold determines the weight of the objects that can be picked with the magnet effect: the higher the value, the heavier the objects that can be lifted. However, higher values also imply that strong forces must be applied to release even lighter objects, which can give an unnatural feel of being abnormally “glued” to the object in these cases. This effect could be reduced by dynamically modulating the threshold with the weight of the picked object.

Not only does this approach allow to carry heavier virtual objects, but it also allows the rotation of these objects on two degrees of freedom, even with 3DoF devices, by using the relative position between both proxies. In addition, if at least one of the two devices has 6DoF sensing (possibly underactuated), then it is possible to handle the third rotational degree of freedom of the object.

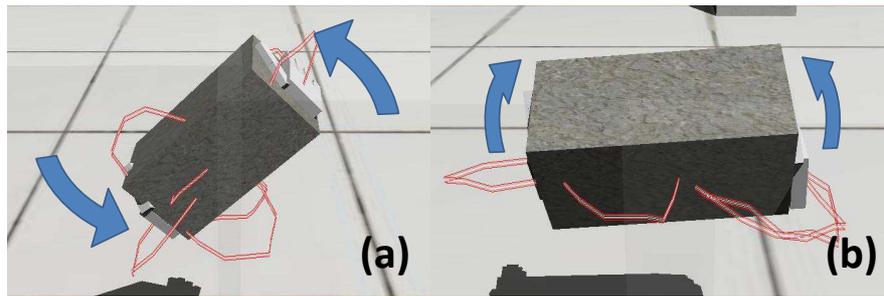


Fig. 10 Rotation of a virtual object using the constraint-based approach of the *magnetic pinch*: (a) Rotation using the relative position between both hands. (b) Rotation using the torque of the right device.

3.5 Joint Control

The *double bubble* metaphor may introduce a difference in control modes and/or scaling factors when activated (Figure 11). In order to reduce the impact of these differences when pick-and-placing a virtual object, the notion of *joint control* was introduced. During a grasping situation, both devices use a common control/display ratio (average of both) and common bubble size (minimal dimensions), and enter rate control simultaneously when at least one device leaves its *bubble*. This technique allows easier exploration of a VE when holding an object between virtual hands controlled by two different haptic interfaces.

4 Evaluation

An experiment assessed the efficiency of the previously mentioned techniques, involving a simple pick-and-place task, where users had to pick a cube and place it at a given position. The *double bubble* technique was compared to the *clutching* technique for workspace extension, and the benefits of the *magnetic pinch* and *joint control* were also measured for grasping facilitation.

4.1 Method

4.1.1 Population

Thirteen participants (2 females and 11 males) aged from 20 to 26 (mean = 22.8, sd = 1.7) performed the experiment. None of the participants had any known perception disorder. All participants were naïve with respect to the proposed techniques, as well as to the experimental setup and the purpose of the experiment.

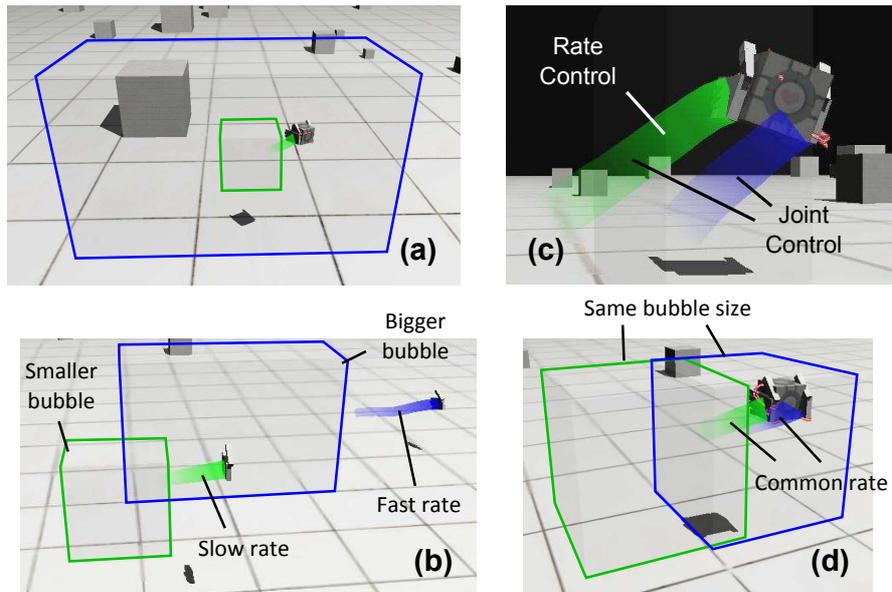


Fig. 11 Illustration of *joint control*. (a) Carrying an object without *joint control* ; case where the left *bubble* is in rate control and not the right *bubble*. (b) Difference in *bubble* size and workspace translation speed without *joint control*. (c-d) Carrying an object with *joint control*.

4.1.2 Experimental Apparatus

The participants were seated at 1m in front of a 24 inch widescreen monitor. The experiment was conducted using two different haptic interfaces. The participants manipulated a Falcon (Novint Technologies Inc., Albuquerque, New Mexico, USA) in their left hand, and a PHANTOM Omni (Sensable Technologies, Wilmington, Massachusetts, USA) in their right hand, both placed in front of the screen as shown in Figure 12. Visual feedback was rendered at a refresh rate of 50 Hz, while the haptic rendering rate was 1,000 Hz. Physical simulation was performed using Nvidia PhysX at a rate of 1,000 Hz to match the update frequency of the haptic loop. A virtual coupling mechanism [2] was used between the haptic interfaces and the virtual proxies by simulating a spring-damper system between each haptic device and its corresponding proxy.

4.1.3 Virtual Environment

The VE was composed of a 100m-wide ground plane with four potential target planes, of 1m of width, placed at the corners of a 6m-wide square around the center of the VE. The target plane of each trial was colored in red, and the other planes were colored in white. The cube to be manipulated had a width of 30cm and a mass



Fig. 12 Apparatus used in the experiment.

of 3g, and was placed at the center of the VE. The proxies controlled by each haptic device were physically represented by cubes of 20cm of width, and were positioned 2m away from each other and 5m away from the central cube at the start of each trial. The cube was thus lying beyond the limits of the workspaces. The proxy controlled by the left device was visually represented by a blue left hand, and the right proxy was represented by a green right hand. Figure 13 shows the scene as displayed at the beginning of a trial.

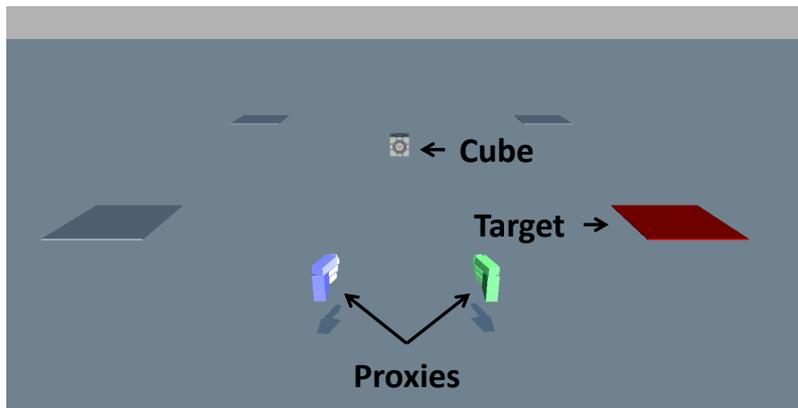


Fig. 13 Virtual environment used in the experiment.

4.1.4 Procedure

At the start of each trial, both haptic devices and proxies were set to their starting positions. The subject had to pick the cube from both sides, carry it towards the red target and make the cube contact with the target, thus ending the trial. A black screen warned the subject about the beginning of the next trial.

4.1.5 Experimental Conditions

A within-subject design was used to evaluate the four different conditions. In the control condition *Ctrl*, the participants were able to use the *clutching* technique when they reached the limits of the workspaces. The three other conditions corresponded to: (1) *DB* (*double bubble*), (2) *MP* (*clutching with magnetic pinch/joint control*) and (3) *DB+MP* (a combination of *double bubble* and *magnetic pinch/joint control*). All the conditions were tested 44 times (11 times per target). The order between the different conditions was counterbalanced across participants, and for each condition, the order between the targets was randomized. The experiment lasted around 1 hour.

4.1.6 Collected Data

For each trial and each participant, the completion time and number of drops were recorded. The completion time is the time elapsed between the moment the proxies leave their starting positions and the moment the cube touches its target plane. The number of drops is the number of hits recorded between the cube and any part of the ground plane that is not the target plane. At the end of the experiment, participants had to complete a subjective questionnaire in which they had to grade the different techniques according to different criteria. The participants could rate the criteria from 1 (very bad) to 7 (very good). The different criteria were: (1) Global appreciation, (2) Efficiency, (3) Learning, (4) Usability, (5) Fatigue, and (6) Realism.

4.2 Experiment Results

4.2.1 Completion Time and Number of Drops

A statistical analysis was conducted from the completion time data collected during the experiment. For each participant, statistics (mean M , standard deviation SD) were computed on the 44 trials in each condition. A Friedman test on the completion time (in seconds) revealed a significant effect of the technique ($\chi^2 = 27.66$, $p < 0.001$). Follow-up post-hoc analysis revealed that completion time in both the *MP* ($M = 14.16$, $SD = 7.14$) and *DB/MP* ($M = 8.43$, $SD = 2.91$) conditions were sig-

nificantly shorter than in the control ($M = 21.41$, $SD = 13.19$) and DB ($M = 20.06$, $SD = 14.63$) conditions ($p < 0.001$ in all cases), and that the $DB+MP$ condition led to significantly shorter times than the MP condition as well ($p < 0.001$).

Similarly, a statistical analysis was conducted on the number of drops for all trials of each participant. A Friedman test showed a significant effect of the technique ($\chi^2 = 25.52$, $p < 0.001$). Post-hoc analysis showed that the MP ($M = 4.22$, $SD = 9.45$) and DB/MP ($M = 2.36$, $SD = 2.33$) conditions led to significantly less drops than the control ($M = 7.88$, $SD = 6.37$) and DB ($M = 8.79$, $SD = 6.77$) conditions ($p < 0.001$ in all cases).

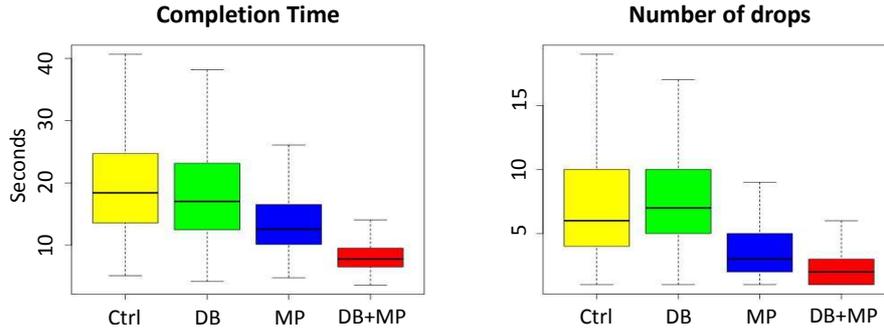


Fig. 14 Box plots of the completion times and number of drops for all conditions. They are delimited by the quartile (25% quartile and 75% quartile) of the distribution of the condition over the individuals. The median is represented for each trial.

4.2.2 Subjective Questionnaire

A Friedman test was performed to analyse the answers of the participants to the subjective questionnaire. The reported p-values were adjusted for multiple comparisons (alpha-level $p=0.05$). A significant effect was found for 5 criteria: Global appreciation ($\chi^2 = 4.62$, $p < 0.001$), Efficiency ($\chi^2 = 4.92$, $p < 0.001$), Learning easiness ($\chi^2 = 4.50$, $p < 0.001$), Use easiness ($\chi^2 = 4.80$, $p < 0.001$) and Fatigue ($\chi^2 = 4.46$, $p < 0.001$).

Post-hoc analysis showed that the $DB+MP$ condition was preferred to both the control and DB for all criteria: Global appreciation ($p < 0.001$ and $p < 0.001$ respectively), Efficiency ($p < 0.001$ and $p < 0.001$), Learning ($p < 0.001$ and $p < 0.001$), Usability ($p < 0.001$ and $p < 0.001$) and Fatigue ($p < 0.001$ and $p < 0.001$). The MP condition was also preferred over the control and DB for 3 criteria: Global appreciation ($p = 0.029$ and $p = 0.028$), Learning ($p = 0.032$ and $p = 0.009$) and Usability ($p = 0.027$ and $p = 0.008$), plus a fourth criteria for the DB : Efficiency ($p = 0.020$).

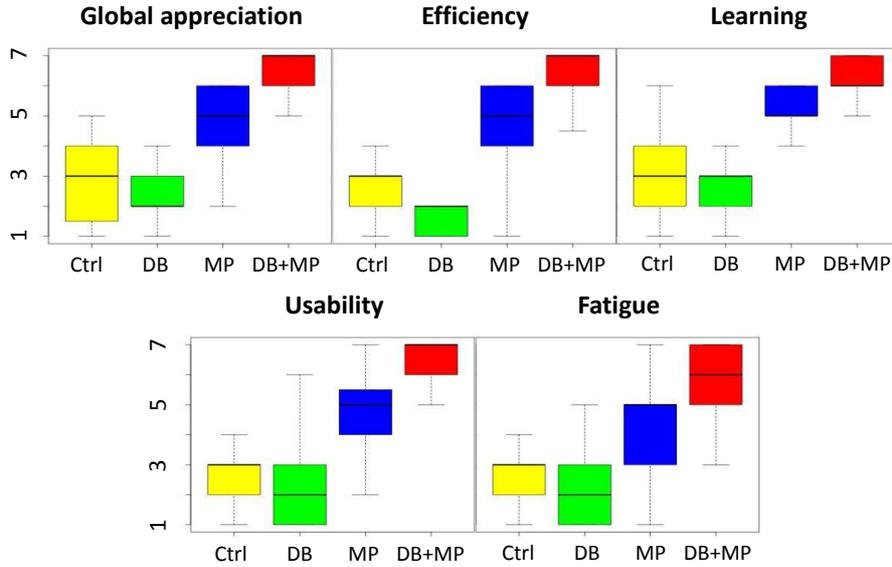


Fig. 15 Box plots of the subjective ratings for the significant criteria, for all conditions. They are delimited by the quartile (25% quantile and 75% quantile) of the distribution of the condition over the individuals. The median is represented for each trial.

4.3 Discussion

The conducted experiment showed that the *magnetic pinch* and *joint control* improved performances and subjective appreciation for a pick-and-place task over the *double bubble* and *clutching* navigation techniques, while the combination of these techniques with the *double bubble* led to the best results.

The *double bubble*, used alone, performed as good as the *clutching* technique without outperforming it, in terms of completion time, drop rate, and subjective appreciation. The technique allows to translate the workspace in a VE in a smoother way than the *clutching* technique, by removing the need to move the devices back and forth several times.

The experiment showed that the *magnetic pinch* and *joint control* significantly reduced completion times and dropping rates compared to the conditions that did not use them. In addition, the subjective appreciation also favored the conditions which used these techniques over those that did not, globally as well as for learning and usability. These results strongly indicate that the *magnetic pinch* and *joint control* techniques, by stabilizing the grasping of a virtual object with virtual proxies, are efficient for facilitating pick-and-place tasks. Additionally, while the *magnetic pinch* inherently adds an unrealistic behaviour through the magnetic attraction, it does not seem to hinder the global realism of the scene, as no significant difference in the participants perception of realism was reported for the different conditions.

The *double bubble* showed its full potential when used jointly with the *magnetic pinch* and *joint control*, outperforming the combination of the latter techniques with *clutching*. The *double bubble* allows users to perform the task in a simpler and faster way than the *clutching* technique, which imposes frequent stops of both proxies to recenter the two haptic devices.

5 Perspectives

Haptic interaction with virtual environments using two hands is a topic that is being increasingly investigated but is still far from reaching its full potential. Interaction techniques dedicated to that field are not numerous, as in a majority of cases unimanual techniques are used in a dual way rather than developing actual bimanual techniques and metaphors. Notably, this means that the same tasks and possibilities are given to both hands, not taking into account the intrinsic differences between them and their ability to perform asymmetric tasks. Exceptions are studies that focus on specific tasks such as medical training or 3D modeling, which assign a different tool to each hand prior to the beginning of the task, with possibly two-handed interactions between these tools.

All haptic devices suited for bimanual interaction have fairly limited physical workspaces, raising the need for interaction techniques to extend them. Most of the existing haptic exploration techniques that allow to do this are unimanual, notably the *clutching* and *Bubble* techniques. The *Double Bubble* metaphor extended the latter to a bimanual use, allowing independent translation of the virtual workspaces of both devices within a large virtual environment through hybrid position/rate control modes. While the *clutching* technique leads to “jerky” motions with the need to make frequent stops when the boundaries of the workspaces are reached, the *double bubble* allows much smoother motions in virtual environments. These techniques only allow translational movements originally, but two approaches allow to rotate the viewport as well: the bimanual implementation of the original *bubble* technique (at the expense of lateral translations), as well as the separation plane of the *double bubble*.

While multi-finger interfaces enable good manipulations of virtual objects through their multiple contacts, it is not the case for single-point devices. The grasping of objects with these devices can be tedious, thus raising the need for interaction techniques such as the *magnetic pinch* to simplify the lifting and carrying of objects with only two contact points. Applying a small magnet effect does help for maintaining the grip of the object, although some issues remain. Notably, holding an object with two haptic devices allows to translate the said object in all directions, as well as rotate it using the relative position between the devices, but there is a degree of freedom that cannot be controlled at all with 3DoF devices. Also, while this technique was developed for single-point interfaces, its relevance for multi-finger interaction and the question of how to adapt it to that case remain to be studied. Finally, the

base principles work perfectly for rigid bodies, but whether they could work and/or be adapted for interaction of deformable bodies is yet to be explored.

It was shown that more issues arise when exploration and grasping happen simultaneously, mostly with single-point devices. With *clutching*, drops tend to happen during the temporary decoupling between devices and proxies, while with the *double bubble*, the differences of velocities in rate control as well as times of activation of rate control can lead to unwanted drops as well. Thus, techniques like the *joint control* are required to handle these two classes of haptic interaction simultaneously. However, there are still several issues to be addressed. Notably, it is still impossible to rotate the viewport with an object in hand using the *double bubble*.

6 Conclusion

This chapter focused on haptic interaction with VEs using both hands, by overviewing several interaction techniques suited for two-handed haptics. Among them, workspace extension techniques such as the *clutching* and *double bubble* allow bi-manual haptic exploration of large VEs, the latter leading to smoother motions than the former. While multi-finger interfaces allowed good handling of virtual objects through their multiple contact points, single-point devices strongly benefit from interaction techniques such as the *magnetic pinch* that maintain the contacts between virtual proxies and picked objects. It is also apparent that there is a need for interaction techniques such as the *joint control* at the interface of both navigation and manipulation, as handling them separately leads to recurrent issues. Future work will develop these techniques, notably by applying them to devices or proxies with more contact points, to other types of simulations such as deformable objects, or to interfaces with more degrees of freedom. Finally, interaction techniques presented in this chapter were used in simple virtual environments, with a very simple task, and future work will assess their applicability and usefulness for real applications such as medical training or industrial prototyping.

References

1. F. Barbagli, K. Salisbury, and R. Devengenzo. Enabling multi-finger, multi-hand virtualized grasping. In *Proc. of IEEE International Conference on Robotics and Automation*, volume 1, pages 809 – 815 vol.1, sept. 2003.
2. J.E. Colgate, M.C. Stanley, and J.M. Brown. Issues in the haptic display of tool use. In *Proc. of International Conference on Intelligent Robots and Systems*, volume 3, pages 140 –145 vol.3, aug 1995.
3. L. D. Cutler, B. Fröhlich, and P. Hanrahan. Two-handed direct manipulation on the responsive workbench. In *Proc. of Symposium on Interactive 3D graphics, I3D '97*, pages 107–114. ACM, 1997.
4. M. de Pascale, A. Formaglio, and D. Prattichizzo. A mobile platform for haptic grasping in large environments. *Virtual Reality*, 10:11–23, 2006. 10.1007/s10055-006-0026-6.

5. M. de Pascale and D. Prattichizzo. The haptik library: A component based architecture for uniform access to haptic devices. *IEEE Robotics & Automation Magazine*, 14(4):64–74, December 2007.
6. L. Dominjon, A. Lécuyer, J.-M. Burkhardt, G. Andrade-Barroso, and S. Richir. The "bubble" technique: Interacting with large virtual environments using haptic devices with limited workspace. In *Proc. of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, WHC '05, pages 639–640, Washington, DC, USA, 2005. IEEE Computer Society.
7. L. Dominjon, J. Perret, and A. Lécuyer. Novel devices and interaction techniques for human-scale haptics. *Vis. Comput.*, 23(4):257–266, March 2007.
8. T. Endo, T. Yoshikawa, and H. Kawasaki. Collision avoidance control for a multi-fingered bimanual haptic interface. In *Proc. of International Conference on Haptics - Generating and Perceiving Tangible Sensations: Part II*, EuroHaptics'10, pages 251–256. Springer-Verlag, 2010.
9. A. Faeth, M. Oren, J. Sheller, S. Godinez, and C. Harding. Cutting, deforming and painting of 3d meshes in a two handed viso-haptic vr system. In *Proc. of IEEE Virtual Reality*, pages 213–216, march 2008.
10. A. Fischer and J. M. Vance. Phantom haptic device implemented in a projection screen virtual environment. In *Proc. of Workshop on Virtual Environments*, EGVE '03, pages 225–229, New York, NY, USA, 2003. ACM.
11. A. Formaglio, D. Prattichizzo, F. Barbagli, and A. Giannitrapani. Dynamic performance of mobile haptic interfaces. *IEEE Transactions on Robotics*, 24(3):559–575, june 2008.
12. P. García-Robledo, J. Ortego, J. Barrio, I. Galiana, M. Ferre, and R. Aracil. Multifinger haptic interface for bimanual manipulation of virtual objects. In *Proc. of IEEE International Workshop on Haptic Audio-visual Environments and Games*, pages 30–35, nov. 2009.
13. P. García-Robledo, J. Ortego, M. Ferre, J. Barrio, and M.A. Sánchez-Urán. Segmentation of bimanual virtual object manipulation tasks using multifinger haptic interfaces. *Instrumentation and Measurement, IEEE Transactions on*, 60(1):69–80, jan. 2011.
14. Y. Guiard. Asymmetric division of labor in human skilled bimanual action: the kinematic chain as a model. *Journal of Motor Behavior*, 19(4):486–517, dec. 1987.
15. T. Hulin, M. Sagardia, J. Artigas, S. Schaetzle, P. Kremer, and C. Preusche. Human-scale bimanual haptic interface. In *Proc. of 5th International Conference on Enactive Interfaces*, pages 28–33, nov. 2008.
16. M. Isshiki, T. Sezaki, K. Akahane, N. Hashimoto, and Sato M. A proposal of a clutch mechanism for 6dof haptic devices. In *Proc. of 18th International Conference on Artificial Reality and Telexistence*, pages 57–63, 2008.
17. Siqiao Li, A. Frisoli, C.A. Avizzano, E. Ruffaldi, L.I. Lugo-Villeda, and M. Bergamasco. Bimanual haptic-desktop platform for upper-limb post-stroke rehabilitation: Practical trials. In *Proc. of IEEE International Conference on Robotics and Biomimetics*, pages 480–485, dec. 2009.
18. S. Loi-Wah, F. Van Meer, Y. Bailly, and C.K. Yeung. Design and development of a da vinci surgical system simulator. In *Proc. of International Conference on Mechatronics and Automation*, pages 1050–1055, aug. 2007.
19. J. Murayama, L. Bougrila, Y. Luo, K. Akahane, S. Hasegawa, B. Hirsbrunner, and M. Sato. Spidar g&g: A two-handed haptic interface for bimanual vr interaction. In *Proc. of EuroHaptics*, pages 138–146, 2004.
20. Novint. www.novint.com/index.php.
21. R. Ott, V. De Perrot, D. Thalmann, and F. Vexo. Mhaptic: a haptic manipulation library for generic virtual environments. In *Proc. of International Conference on Cyberworlds*, CW '07, pages 338–345, Washington, DC, USA, 2007. IEEE Computer Society.
22. A. Peer and M. Buss. A new admittance-type haptic interface for bimanual manipulations. *IEEE/ASME Transactions on Mechatronics*, 13(4):416–428, aug. 2008.
23. A. Peer, U. Unterhinninghofen, and M. Buss. Tele-assembly in wide remote environments. In *Proc. of 2nd International Workshop on Human-Centered Robotic Systems*, 2006.

24. Sensable Technologies. <http://www.sensable.com/>.
25. SenseGraphics. <http://www.h3dapi.org/>.
26. A. Talvas, M. Marchal, C. Nicolas, G. Cirio, M. Emily, and A. Lécuyer. Novel interactive techniques for bimanual manipulation of 3d objects with two 3dof haptic interfaces. In *Proc. of EuroHaptics (1)*, pages 552–563, 2012.
27. S. Ullrich, D. Rausch, and T. Kuhlen. Bimanual haptic simulator for medical training: System architecture and performance measurement. In *Joint Virtual Reality Conference of EuroVR - EGVE (2011)*, 2011.
28. S. Walairacht, Y. Koike, and M. Sato. String-based haptic interface device for multi-fingers. In *Proc. of IEEE Virtual Reality*, 2000.
29. K. Waldron and K. Tollon. Mechanical characterization of the immersion corp. haptic, bimanual, surgical simulator interface. In Bruno Siciliano and Paolo Dario, editors, *Experimental Robotics VIII*, volume 5 of *Springer Tracts in Advanced Robotics*, pages 106–112. Springer Berlin / Heidelberg, 2003.
30. S. Zhai. User performance in relation to 3d input device design. *SIGGRAPH Comput. Graph.*, 32(4):50–54, November 1998.