User Performance in Complex Bi-manual Haptic Manipulation with 3 DOFs vs. 6 DOFs

Rene Weller*
Clausthal University

Gabriel Zachmann†
Clausthal University

ABSTRACT

We present the results of a comprehensive user study that evaluates the influence of the degrees of freedom on the users' performance in complex bi-manual haptic interaction tasks. To do that, we have developed a novel multi-player game that allows the qualitative as well as the quantitative evaluation of different force-feedback devices simultaneously. The game closely resembles typical tasks arising in tele-operation scenarios or virtual assembly simulations; thus, the results of our user study apply directly to real-world industrial applications. The game is based on our new haptic workspace that supports high fidelity, two-handed multi-user interactions in scenarios containing a large number of dynamically simulated rigid objects; moreover, it works independent of the objects' polygon count. The results of our user study show that 6 DOF force-feedback devices outperform 3 DOF devices significantly, both in user perception and in user performance.

1 INTRODUCTION

Haptics is an emerging technology; it adds the sense of touch to applications in fields like tele-operations, medical simulations, or virtual assembly tasks that are known from the automotive and aircraft industry. In these areas, force-feedback already helps to improve human-computer, as well as human-human interactions in multi-user scenarios for almost two decades.

For a long time, haptic devices were bulky, expensive, and could be installed and handled only by experts. This has changed but in the last few years, when the first low-cost haptic devices entered the market, which were designed especially for desktop use. Besides typical consumer electronic applications like games or online shops, where the sense of touch could be a decision criterion for selecting products, these low-cost devices could also be used to improve the quality of training skills or enhance the desktop of each constructing or design engineer.

However, if a whole engineering office should be equipped with haptic devices cost could be still a limiting factor, even if they are low-cost machines. The cost of haptic devices mainly depends on the number of actuators. Consequently, the low-cost devices for the mass market usually support only 3 DOFs. Obviously, real - world object manipulations comprises not only forces with 3 DOFs but also torques with 3 DOFs. Therefore, rendering these kinds of interactions faithfully requires much more expensive 6 DOF haptic devices.

This raises the question whether or not the enhanced experience is worth the additional cost for the 6 DOF devices, which is precisely the question that this paper endeavors to answer.

Intuitively, it seems obvious that users operating with full 6 DOFs should perform much better than users that are provided only 3 DOFs. In fact, the influence of the DOFs in human-computer interaction is still an active field of research, with partly contradictory results, even if they do not include haptics and are restricted to single-hand interactions. However, this paper not only presents a qualitative analysis, but also quantitative methodologies to assess the influence of full 6 DOF force and torque rendering objectively.

In order to conduct our user studies, we have implemented a haptic workspace that provides high-fidelity 6 DOF force-feedback in object manipulation scenarios containing a large number of dynamically simulated rigid objects. In addition, it supports different kinds of haptic (and non-haptic) devices for bi-manual multi-user interactions. We have implemented a new collision detection technique to meet the special requirements of haptic devices for a very high simulation frequency.

It is a challenge to define a task that does not favor one of the input methods in advance. In our case, this means we need a task that can be solved with 3 DOF devices as well as with 6 DOF devices with the same level of success. Moreover, we need a task that requires coordinated bi-manual interactions from the users. Therefore, we have developed a simple haptic multi-player game that requires complex, two-handed manipulations of two players within the same environment at the same time.

In order to evaluate the users' performance, we recorded all paths of all objects, including those of the users' hands, for later quantitative and qualitative analysis. Moreover, we utilized a questionnaire to evaluate some of the "softer" factors of such a haptic workspace.

The results support our initial hypothesis, that 6 DOF haptic devices outperform 3 DOF haptic devices with respect to user perception and also user performance. This might encourage device manufacturers to spend more efforts in the development of cheaper 6 DOF haptic devices for desktop use.

2 RELATED WORK

Haptic user interfaces have been actively applied to the domain of human-computer interaction in virtual environments for almost two decades. Many user studies have shown that providing haptic feedback during virtual interaction tasks has positive effects on the perceived realism.

For instance, [2] developed a multimodal shared virtual environment. The experiments showed that force-feedback during collaboration with a remote partner contributes to the feeling of "sense of togetherness", which is a kind of presence. Moreover, force-feedback also helps to improve the user performance. Other authors obtained very similar results with respect to multi-user haptic interactions. Experiments cover a wide spectrum of tasks reaching from training of motor skills in surgery [5], rehabilitation tasks [7], tele-operation [13] to computer games [14]. Moreover, haptic systems can also help to enhance the emotional immersion in real-time messaging. [15] developed a virtual hug system that supports 3D virtual worlds like Second Life.

Furthermore, some bi-manual haptic workspaces have been developed already: [12] used two SPIDAR-G devices that provide 6 DOF motion and 6 DOF force-feedback. A simple 3D pointing

---

*e-mail: rwe@tu-clausthal.de
†e-mail:zach@tu-clausthal.de

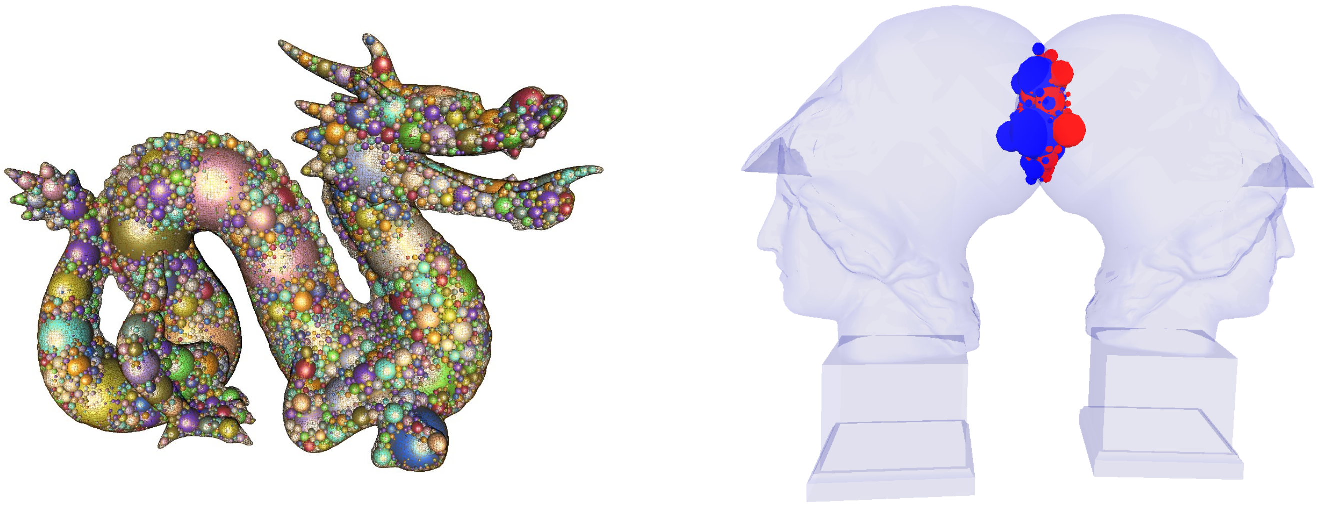

Figure 1: *Left: ISTs are based on sphere packings of the object. Right: Such sphere packings enable us to compute an approximation of the penetration volume if the objects overlap.*

task was used to evaluate the system. The results indicate that bi-manual haptic interactions are more intuitive and efficient with respect to task completion time than single-handed manipulations. Two-handed haptic interaction has also shown to be a promising way for shape modelling applications: [1] was able to ensure an enhanced precision during interaction; [3] applied a two-handed tracking system and Phantom devices to help users control their gestures during sketching 3D shapes directly in 3D space.

In addition, there exists a large body of work on two-handed interaction in general, without a special focus on haptics. For instance, [8] has shown that two-handed interaction combines two types of advantages: first, twice as many degrees of freedom simultaneously available to the user can result in increased motion efficiency; second, single-handed interaction often requires a higher level of abstraction because of an unnatural, mental composition task. Consequently, bi-manual interaction can reduce the cognitive load. [16] was partly able to validate these assumptions. They conducted a user study to test two-handed freeform deformations using datagloves. The results show an improvement of the user's perception, but only if the degree of symmetry was high.

However, the effect of the degrees of freedom on the user's perception is still an active field of research. [6] proposed a theoretical principle to capture the control structure of an input device: a device that is able to move directly across all dimensions is called an *integral* device, while a device that constrains the user's movement along a single dimension is called a *separable* device. This is an extension to a theoretical framework proposed in [4] called the *perceptual structure* of objects and tasks by structuring its attributes into integral and separable attributes. They supported this theory by showing that user performance increases if the perceptual structure of the object being manipulated matches the control structure of the device. However, the matter does not seem to be settled yet, since [10] obtained completely opposite results when conducting a simple manipulation experiment using a dataglove for an integral device versus a touchscreen for a separable device: the results suggest that the simultaneous manipulation of all DOFs does not necessarily lead to better performance. [11] validated these results when investigating 3D manipulation using a 2D multitouch screen.

However, all of the experiments mentioned in the above two paragraphs were conducted without any force-feedback. Consequently, it is impossible to extend the findings directly to haptic environments. [10], for example, explains his results by real-world constraints that reduce the interaction dimensionality in the real world, such as gravity. But, with haptic devices it is easy to model these physical constraints as well.

To our knowledge, there is very little work on the comparison of haptic devices with different degrees of freedom. [18] presented a study about the effect of torque-feedback on purely haptic perception of the location of objects in virtual environments. Usually, research concentrated mostly on analyzing devices with an asymmetric number of sensors and actuators. For instance, [17] found that for tasks like drawing or tracing, devices with 3 DOFs of force and an additional 3 DOFs of positioning can approximate the performance of full force and torque feedback.

An extended abstract of preliminary results of our user study has been published in [22].

3 THE HAPTIC WORKSPACE

The main challenge when doing haptic rendering is the extremely high frequency that is required: While the temporal resolution of the human eye is limited to approximately 30 Hz, the bandwidth of the human tactile system is about 1000 Hz. In most haptic scenarios, the computational bottleneck remains the collision detection, whereas the force computation can be done relatively fast.

In order to achieve such a high simulation rate, the heart of our haptic workspace is our new geometric data structure, called *Inner Sphere Trees (ISTs)*, that not only allows us to detect collisions between pairs of massive objects at haptic rates, but also enables us to define a novel type of contact information that guarantees *stable* forces and torques [20].

3.1 Inner Sphere Trees

The main idea of the ISTs is that we do not build an (outer) hierarchy based on the polygons on the boundary of an object, like most other bounding volume hierarchies do, but we fill the interior of the

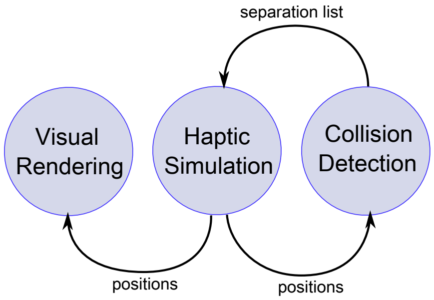

Figure 2: *The simulation thread in our haptic workspace computes the collision forces based on the* separation list*, which captures the current collision information. This list is generated in the collision detection thread. Conversely, the haptic thread passes the new positions of the objects to the collision and the (visual) rendering thread.*

model with a set of *non-overlapping* spheres that cover the object's volume densely. On top of these inner BVs, we create a hierarchy in order to accelerate the collision detection queries. This enables us to define a novel extent of intersection, the *penetration volume* (See Figure 1). The penetration volume corresponds to the amount of water being displaced by the overlapping parts of the objects and, thus, leads to a physically motivated and *continuous* penalty force.

Our ISTs and, consequently, the collision detection algorithm are independent of the geometry complexity. Moreover, they support all kinds of object representations, including polygon meshes and NURBS surfaces, whereas their memory consumption is very modest.

Our collision detection scheme together with a novel penalty force approach that is based on the penetration volume, enable us to treat physically based simulation and haptic rendering in a common way. The only difference between dynamic objects and user-controlled objects is, that the forces for the latter are rendered to the haptic device instead of using them for the simulation. For further information we refer the interested reader to [21].

For visual output we use an open source scenegraph[1] that supports shading and multi-monitor output.

Even if the ISTs are very fast, it is not possible to guarantee constant time intervals for the collision detection. Therefore, we extended the algorithm with a time critical approach and included multithreading support.

3.2 Time-critical multi-threaded computation of penetration volume

In cases of interpenetrating objects, the computation of the penetration volume can run slower than the required 1000Hz, because it might have to visit many nodes during traversal, especially in cases with heavy overlaps. Consequently, an answer of this query type can not be guaranteed within a predefined time budget as it is needed for haptic applications. Moreover, the force computation requires time, too.

On the other hand, almost all currently available CPUs include multiple cores or, at least, support functions to accelerate multithreading.

One appropriate strategy to realize time-critical traversals is a decoupling of the force computation and the collision detection by running them asynchronously in different threads. In the following, we present more details about our approach, which is based on separation lists (see [9, 19]).

We divide the work into several independent threads: 1) a *haptic simulation thread*, which is responsible to handle the user input and computes the forces, and 2) a *collision detection thread*, in which separation lists are generated. Depending on the application it is, of course, possible to add other threads, e.g., a rendering thread.

During runtime, the collision detection thread only maintains a separation list and passes it to the haptic thread. In return, the haptic thread passes the current positions of the simulated objects to the collision detection thread for the next query. The haptic thread then uses the current separation list to compute the force, until the next collision detection query is finished.

Usually, especially in haptic simulations running at 1 kHz, the spatial coherence is high and thus, the separation lists between two synchronizations do not differ very much.

4 THE DESIGN OF THE STUDY: A HAPTIC GAME

Usually, when designing haptic user studies, some kind of object docking or path following task is used. Unfortunately, these kinds of tasks are not very well suited when one wants to compare the influence of the degrees of freedom because depending on the dock or the path, one of the devices is favored in advance. For example, if a docking task requires a rotation of the object, it is impossible to solve it with a 3 DOF device that does not support changes of the orientation. On the other hand, if the task does not require changes of the object's orientation, there would be no need for a 6 DOF device. Moreover, these tasks usually can be solved with a single-handed device. Consequently, there is no need for coordinations between both hands, which is essential in bi-manual interaction tasks.

Consequently, we had to design a new kind of experiment that supports a fair comparison of devices with different degrees of freedom and additionally requires complex bi-manual interactions not only as an option, but as a necessity. Therefore, we use a kind of indirect and imprecise docking task. This means the objects to place are not directly glued to the haptic tool but must be controlled indirectly following a physically-based simulation. Moreover, the objects do not have to be placed precisely into a predefined docking station, but into a wider goal.

This indirect interaction metaphor that we propose resembles closely typical tasks arising in bi-manual tele-operation scenarios or virtual assembly simulations. Thus, the analysis of the users' performance in this experiment allows for conclusions of practical relevance.

In detail, we have implemented a simple two-player haptic game that is based on our haptic workspace. The players sit face-to-face at a table with two monitors in between (See Figure 3). Each player operates the two identical force-feedback devices on his side, one for each hand. In order to evaluate the differences between 3- and 6 DOF interactions, one of the players uses two 3 DOF devices[2], whereas his opponent operates two 6 DOF devices[3].

We used these kinds of force-feedback devices, because they have comparable specifications (see Figure 4), they are both designed for desktop use, and there is no other pair of devices that differs in DOFs yet has similar specs.

The playing field is a room with a set of complex objects with different shapes lying on the ground. Each player has a "well" in

---

[1] OpenSG, www.opensg.org

[2] Novint Falcon, www.novint.com

[3] Haption Virtuose 6D Desktop, www.haption.com

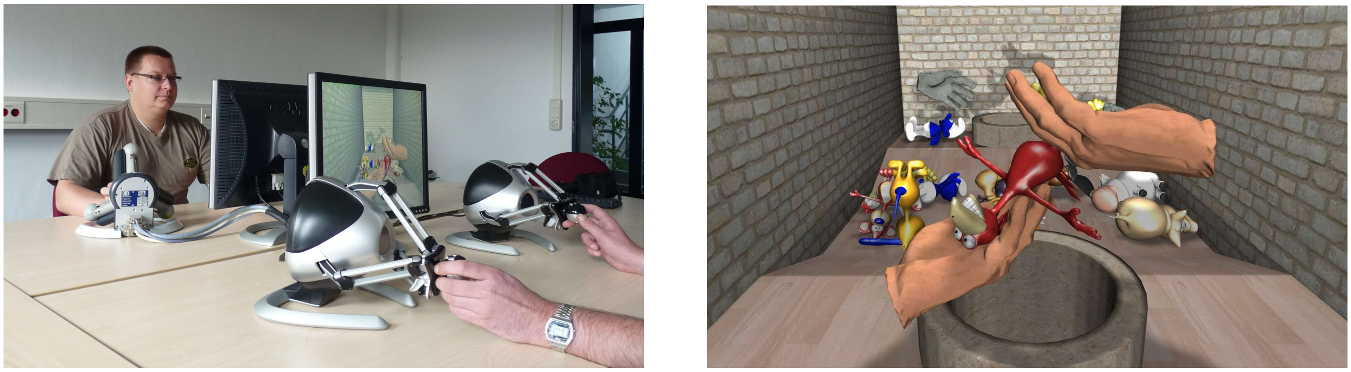

Figure 3: *The two-player set up with four haptic devices for our user study (left). The playing field of our haptic game (right).*

front of him and controls two rigid virtual hands with his two force-feedback devices. The goal of the game is to pick up the objects and place them in the player's own well in front of him. Figure 3 gives an overview of the setup with the four haptic devices; it also shows a typical view of the playing field.

Even if the task is the same for both players, different strategies can lead to the goal depending on the degrees of freedom of the devices. In tests prior to the final study's design, the 6 DOF operators usually picked up a single object and directly placed it in the well. On the other hand, the 3 DOF users shoveled some of the objects to the front of the well and tried to push them up the well's walls (we dubbed this the "shovel technique"). Consequently, the success of both techniques can be tweaked by the hight of the well and the number of objects in the scene. In order to guarantee a fair comparison we adjusted the parameters such that with both techniques the chance to win and the chance to pocket an objects is almost the same for both input devices. Additionally, we chose the objects such that their size and form factor forces the users to really use coordinated bi-manual interactions.

For two reasons it is essential that we do not take the winning rate or the number of pocketed objects as distinctive measure: the same probability to win with both kinds of devices proves the fairness and comparability of our results and moreover, the winning rate could also influence the answers of the questionnaire subconsciously.

In order to maintain fairness we also implemented the facility to turn the virtual hands with the 3 DOF devices by mapping rotations to the buttons on the haptic handle (see picture in Figure 4), because it could be complicated for the 3 DOF users to pick up or shovel the objects with the hands remaining in their initial orientation due to the rigidity of the controlled virtual hands. The device has four buttons; we used three of them to change the pitch, yaw, and roll of the virtual hand, while the fourth button changes the direction of the rotation. In addition to the general learning period when operating unknown devices, this relatively complex control paradigm for the three rotational degrees of freedom required some training. Thus each round of the game started with a training phase that ends when both players managed to pocket an object. However, the results of our user study show that almost all participants used the possibility to change the hand's orientation only in the training phase in order to bring the hands into a comfortable orientation. During the game they only made very few attempts to adjust the orientation.

For the evaluation, we recorded the forces and torques acting on the user-controlled hands and, additionally, we tracked the covered distances and rotations. This data allows to derive conclusions about the efficiency of the haptic interaction. In addition, we recorded the time for the training phase. Moreover, we conducted a user interview after the game using a questionnaire, where we asked the users about the quality of the feedback and their preferences with respect to 3 DOFs vs. 6 DOFs.

The setting of a game was chosen to ensure that, due to the competitiveness, the users are highly concentrated on the challenge and not on the potentially unknown and fascinating devices. After finishing a round, the players swap seats. Thus, each player plays with both kinds of devices. Due to this, we were able to test a large amount of subjects in a relatively small time interval, and moreover, we could keep the learning phase relatively short.

5 THE USER STUDY

In the following, we will give an overview of the user study that we conducted using our haptic game described above.

5.1 Participants and Protocol

We tested a total number of 47 participants, aged 17 to 34 years. Half of them were high school students visiting our department of computer science, the others were scientific employees with the department. Of the participants, 33 were male and 14 female, 3 were left-handed and 44 right handed. 27 of them play computer games regularly, and almost all have some experience in gaming, except 4 who stated they never played a computer game before. Only 5 participants use VR devices regularly. 8 subjects did not play our haptic game for the first time, because they already helped in the pre-test phase to improve the game design, but only two of them played it more often than twice. Only these 8 persons had made experiences with haptic devices before, 6 of them during the pre-test-phase.

The participants entered the room with the experimental setup in groups of 4 persons. They were given a short verbal introduction about the game, the experiment, and the special properties and features of the devices, such as the dead-man protection of the 6 DOF device or the mapping of rotations to the buttons of the 3 DOF device.

After this short introduction and a few seconds for the subjects to assume the right and comfortable grasping of the haptic handles, the training phase started immediately. The time for the training phase was restricted to maximally 3 minutes but could end earlier if both players managed to pocket an object. Like the training phase, the game also lasted 3 minutes. During the game, the players received feedback about the score and the time limit by a heads-up display on the screen. After completing the game, the subjects were asked to answer a questionnaire and rate the intuitiveness of control, the

	3 DOF	6 DOF
Manufacturer	Novint	Haption
Model	Falcon	Virtuose 6D Desktop
Translational Workspace	102mm x 102mm x 102mm	Sphere with 120mm in diameter
Rotational Workspace	-	35° in the 3 directions
Maximum force in translation	10N	15N
Maximum torque in rotation	-	0.5 Nm
Price	200$	30,000$

Figure 4: *The haptic devices that we used in our evaluation: The 3 DOF Novint Falcon (left) and the 6 DOF Haption Virtuose 6D Desktop (right). The specifications of both force feedback devices show a comparable workspace and a comparable amount of maximum translational force. The 6 DOF device can additionally render torques.*

quality of the force feedback, and so on, on a five-point Likert scale. The Likert scale has suitable symmetry and equidistance for the use of parametric analysis.

5.2 Results

The groupwise introduction and the attendance of other persons in the room during the test could distract the players. However, the results of our survey show that the concentration during the game was rated very high (3 DOFs: M=4.32, SD=.837, 6 DOFs: M=4.23, SD=1.026, with the Likert scale ranging from “Heavy distractions”=1 to “No distractions”=5). Also the training time (3 DOFs: M=2.51, SD=.655, 6 DOFs: M=2.81, SD=.680, with the Likert scale ranging from “Too short”=1 over “Perfect”=3 to “Too long”=5) and the playing time (3 DOFs: M=2.64, SD=.705, 6 DOFs: M=2.57, SD=.683, with the same Likert scale) was rated as sufficient overall.

As mentioned in the introduction, we hypothesized that 6 DOF haptic devices are better suited for complex bi-manual haptic interactions than 3 DOF devices with respect to intuitiveness and the naturalness of the control paradigms, the quality of the force-feedback, and other parameters. A paired-samples t-test was conducted to compare the measured values and the results of the survey in 3 DOF and 6 DOF conditions.

Overall, the results support our hypothesis that object manipulation using force-feedback with 6 DOFs is more natural and more intuitive: from our survey, we get a highly significant difference in the scores for naturalness of control in the 3 DOF (M=2.83, SD=.816) and 6 DOF (M=3.55, SD=.717) case; t(46)=-6.425, p<0.001 with the Likert scale reaching from “Not natural”=1 to “Perfect natural”=5. We get a similar highly significant result for the intuitiveness of control (3 DOF (M=3.28, SD=.877) and 6 DOF (M=4.04, SD=.779); t(46)=-4.741 p<0.001 (Likert scale from “Not intuitive”=1 to “Perfectly intuitive”=5)). Also, the quality of the force-feedback shows highly significant differences between 3 DOF (M=2.98, SD=1.011) and 6 DOF (M=3.66, SD=.867) conditions; t(46)=-4.761 p<0.001 (Likert scale from “Unsatisfiable”=1 to “Perfect”=5). However, the mediocre absolute values show that there is still room for improvements regarding the naturalness and the quality of the forces and torques.

Even though most subjects rated the time given for the training phase as sufficient for both kinds of devices, the paired-samples t-test shows a significant difference between 3 DOF (M=2.51, SD=.655) and 6 DOF (M=2.81, SD=.680) conditions; t(46)=-2.625, p=0.012. This further supports the results about the intuitiveness of control and the higher naturalness.

In the training phase, the time measured until a player manages to pocket the first object also supports the user’s experience we observed through the questionnaire: they needed significantly more time to learn the handling of the 3 DOF devices (M=94.66, SD=69.370) than the 6 DOF devices (M=60.74, SD=51.809); t(46)=2.954, p=0.005.

In order to guarantee a fair comparison, we adjusted the task so that the 3 DOF operators and the 6 DOF operators can win with the same chance. The measured results support the validity of our calibration: overall, there were 20 rounds of all games won using a 3 DOF device, and 18 rounds won using a 6 DOF device (9 rounds were a tie).

The number of objects that were pocketed by users using the 6 DOF devices was slightly larger (M=5.94, SD=4.532) than the number of objects pocketed by users using the 3 DOF devices (M=5.64, SD=4.321). However, there is no statistically significant difference between the number of pocketed objects with respect to the DOFs.

Additionally, a one-way between-subjects ANOVA was conducted to compare the effect of experience on the number of pocketed objects: there was a significant difference between the group that has haptic experience, which is exactly the group that played the game more than once, and the participants that played the game only for the first time (Unexperienced 3 DOF: N=39, M=4.95, SD=3.692, Experienced 3 DOF N=8, M=9.00, SD=5.757, F(1,46)=6.538, p=0.014, Unexperienced 6 DOF: N=39, M=5.08, SD=3.608, Experienced 6 DOF N=8, M=10.13, SD=6.334, F(1.46)=9.814, p=0.003). In both cases, 3 DOF and 6 DOF, the experienced users was able to pocket significantly more objects than the unexperienced users. However, they were still not able to pocket significantly more objects with 6 DOF than with 3 DOF or vice versa. Also these results show that the calibration of our experiment works correctly: the task can be solved with both kinds of devices with the same succession rate. This implies the fairness of the game.

Even if the chance to win the game is independent of the degrees of freedom, we expected differences in the users’ performance due to the different techniques: as already mentioned in the section before, the 3 DOF users usually shoveled the objects on the ground into the direction of the well, whereas the 6 DOF users precisely picked up the objects. These different strategies directly affects the efficiency of the haptic interactions. The "shovel"-technique can be successful, but it is inefficient with respect to the covered distances, because the users need a higher frequency of forward and backward moving of their hands.

This hypothesis is supported by our measured data: the distances covered by the 6 DOF device that was used with the dominant hand (M=295.8, SD=134.0) is significantly (t(46)=-12.034, p<0.001) shorter compared to the paths of the 3 DOF device used with the dominant hand (M=724.1, SD=235.0). For the non-dominant hand, we obtain almost the same picture (3 DOF (M=374.0, SD=291.5) and 6 DOF (M=605.0, SD=251.4); t(46)=-5.991, p<0.001).

Figure 6 shows the z-position of the virtual hand in the scene, which is controlled by the user. One can clearly see the typical, high-frequency “shovelling” of the 3 DOF user and the relatively smooth motion of the 6 DOF user. Moreover, the plots reveal another typical strategy of the 3 DOF users: they tried to distract the 6 DOF users when they had managed to grab an object. You can see this, for instance, at the 5000-th sample position: here, the 3 DOF user tried to knock the object out of the 6 DOF user's hand.

The above mentioned distance measures for the dominant and the non-dominant hand have some other impacts, too: the distance covered by the dominant hand of the 3 DOF users is significantly longer than that of their non-dominant hand (dominant hand: M=724.1, SD=235.0; non-dominant hand: M=605.0, SD=251.4; t(46)=3.368, p=0.002). Surprisingly, we get the opposite result when looking at the 6 DOF paths (dominant hand: M=295.8, SD=134.0; non-dominant hand: M=374.0, SD=291.5), even if the result is not statistically significant.

Further experiments will have to show if this is an impact of the strain due to the reduced degrees of freedom, or if it is a result of the special "shovel" strategy facilitated by this game.

With the 6 DOF device, the rotation of the user's real hands is mapped directly to the device, whereas with the 3 DOF device, the rotation virtual hand is mapped to the buttons as described above. In other words, with the 6 DOF device, an integral set of object parameters (position and orientation) is mapped to an integral task (moving the end-effector of the device), while with the 3 DOF device the set of object parameters is treated as a separable set [4, 6].

This has, of course, consequences on the strategies that users employ. Usually, the 3 DOF users first brought their virtual hands in a suitable orientation and changed it only very seldomly during the game, whereas the 6 DOF users rotated their real and virtual hands continuously. Figure 6 shows a typical situation. Additionally, we computed the Euler angles and accumulated all rotational changes. This shows significant differences, using the paired-samples t-test, for both the dominant and non-dominant hands (6 DOF dominant: M=90.0, SD=64.0; and 3 DOF dominant: M=15.1, SD=16.0; t(46)=7.495, p<0.001; 6 DOF non-dominant: M=85.9, SD=27.6; and 3 DOF non-dominant: M=13.6, SD=11.5; t(46)=14.883, p<0.001). This suggests that mapping of rotations to buttons cognitively overwhelmed users in time-critical tasks requiring precision motor control.

We used 6 different objects in our game, all of them are cartoon animals (see Figure 5). We chose these objects, because their extremities, like the wide-spreaded arms, oversized feet and ears, or the tails, should simplify the grasping of the objects by clamping them between the fingers of the virtual hands (this facilitated object manipulation considerably). Surprisingly, the only object without strongly protruding extremities, the rhino model, was pocketed most often. We tested the significance with a chi^2-test and obtained a significance level of p<0.01 with the 3 DOF devices, and even p<0.001 with the 6 DOF devices. We believe that this can be a hint that the abstraction between the simple handle of the force-feedback device and the detailed virtual hand cognitively overloads the users, but this has to be investigated in more depth in future studies.

All other factors we investigated, like the age, the sex, and the handedness do not have any significant effects on the user's performance. Even the experience in gaming or with other virtual reality devices does not have any effect. We checked this by using one-way between-subjects ANOVA tests. Eight participants that started with the 6 DOF devices in the first round, and then switched to the 3 DOF devices in the second round, stated after the swap of seats that it was really hard and unnatural to cope with the reduced feasibilities of the 3 DOF devices. Conversely, there was not a single user starting with the 3 DOF device who complained about the extended degrees of freedom after swap of seats. However, the analysis of the users' questionnaires does not show any significant differences between users starting with 3 DOFs and ending with 6 DOFs, or vice-versa, with respect to the rating of the different devices.

6 CONCLUSIONS AND FUTURE WORK

We presented a new multi-user haptic workspace with support for a large number of haptic devices and a likewise number of dynamic objects with a high polygon count. Its multithreaded architecture guarantees a constant simulation rate of 1KHz that is required for stable haptic interactions. Based on our workspace we have implemented a haptic multi-player game with complex bi-manual haptic interactions that we use for a quantitative and qualitative analysis of haptic devices with respect to their number of sensors and actuators.

We conducted a user evaluation with 47 participants. The results show that 6 DOF devices outperform 3 DOF devices significantly, both in user perception and in objective data analysis. For example, the learning phase is much shorter and the users judged the 6 DOF device to be much better with regard to the quality of forces and the intuitiveness of control. However there is still place left for improvements of the haptic devices: The overall rating of force quality and also naturalness of control is rated only mediocre.

However, there are still some challenges left for the future: Further studies are necessary to find the best trade-off between cost and performance regarding bi-manual complex haptic interactions. This could include asymmetric set-ups of the haptic devices, e.g. 6 DOF for the dominant hand and cheaper 3 DOF for the other hand.

Finally, there are also some challenging extensions for our haptic workspace, e.g. it would be nice to extend our approach also to deformable objects or to include networking functionality in order to support tele-operations.

ACKNOWLEDGMENT

This work was partially supported by DFG grant ZA292/1-1 and BMBF grant Avilus / 01 IM 08 001 U. Special thanks go to Jérôme Perret of Haption for lending the 6 DOF devices for our user study.

REFERENCES

[1] F. T. Attar, R. V. Patel, and M. Moallem. Hived: A passive system for haptic interaction and visualization of elastic deformations. *World Haptics Conference*, 0:529–530, 2005.

[2] C. Basdogan, C.-H. Ho, M. A. Srinivasan, and M. Slater. An experimental study on the role of touch in shared virtual environments. *ACM Trans. Comput.-Hum. Interact.*, 7:443–460, December 2000.

[3] D. F. Keefe, R. C. Zeleznik, and D. H. Laidlaw. Drawing on air: Input techniques for controlled 3D line illustration. *IEEE Transactions on Visualization and Computer Graphics*, 13(5):1067–1081, 2007.

[4] W. R. Garner. *The Processing of Information and Structure*. Potomac, 1974.

[5] M. Hutchins, D. Stevenson, M. Adcock, and P. Youngblood. Using collaborative haptics in remote surgical training. In *in Proc.First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems(WHCŠ05*, pages 481–482. IEEE Computer Society, 2005.

[6] R. J. K. Jacob, L. E. Sibert, D. C. McFarlane, and M. P. Mullen, Jr. Integrality and separability of input devices. *ACM Trans. Comput.-Hum. Interact.*, 1:3–26, March 1994.

[7] Y. Jung, S.-C. Yeh, and J. Stewart. Tailoring virtual reality technology for stroke rehabilitation: a human factors design. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI '06, pages 929–934, New York, NY, USA, 2006. ACM.

[8] A. Leganchuk, S. Zhai, and W. Buxton. Manual and cognitive benefits of two-handed input: an experimental study. *ACM Trans. Comput.-Hum. Interact.*, 5:326–359, December 1998.

Figure 5: *Screenshots of the objects we used in the game. Surprisingly, the rhino was pocketed significantly more often than the other objects.*

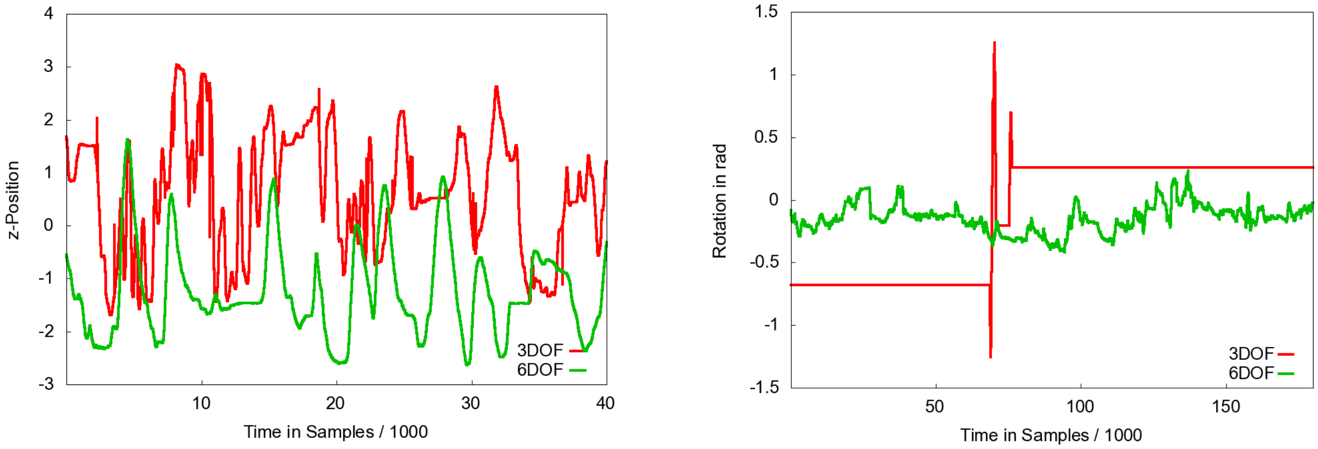

Figure 6: *Typical data recorded from the users' interaction during the game. The plot on the left shows the position of the haptic handle in z direction, which is controlled by the users' dominant hand with the 3 DOF (red) and 6 DOF device (green). Clearly, one can see the typical high frequencies caused by the "shovelling technique" that is often applied by 3 DOF users, whereas the 6 DOF users interact more precisely. Moreover, one can see how the 3 DOF user tries to distract the 6 DOF user at sample time 10k. The right plot shows the roll-angle of the 3 DOF (red) and the 6 DOF (green) users. The 6 DOF users typically rotate their virtual hands continuously, while the 3 DOF users let their hands in almost the same orientation at all times.*

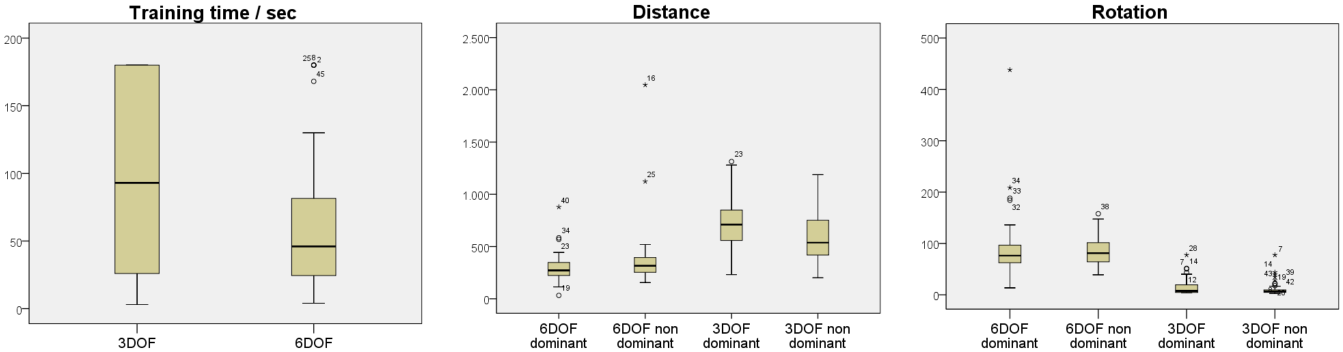

Figure 7: *Left: the time that the users needed to pocket the first object during training with respect to 3 DOF and 6 DOF devices. Clearly, the 3 DOF users needed significantly more time. Middle: the distances covered by the users' dominant and non-dominant virtual hands. Clearly, the paths of the 3 DOF users are significantly longer than the paths of the 6 DOF users. Moreover, they prefer to use their dominant hand. Surprisingly, the 6 DOF users cover a slightly longer path with their non-dominant hand. Right: the total amount of rotations applied by the users during the game, which was obtained by accumulating the changes of the Euler angles. Obviously, the 3 DOF users avoid to rotate their virtual hands, probably because the orientation of the virtual hands is mapped to the buttons of the end-effector of the force-feedback device. (Usually, they brought it in a comfortable position during the training phase and did not change it during the game.*

[9] T.-Y. Li and J.-S. Chen. Incremental 3d collision detection with hierarchical data structures. In *VRST '98: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 139–144, New York, NY, USA, 1998. ACM.

[10] A. C. M. Veit and D. Bechmann. Influence of degrees of freedom's manipulation on performances during orientation tasks in virtual reality environments. In *VRST 2009: The 16th ACM Symposium on Virtual Reality and Software Technology*, Kyoto (Japan), November 2009.

[11] A. Martinet, G. Casiez, and L. Grisoni. The effect of dof separation in 3d manipulation tasks with multi-touch displays. In *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*, VRST '10, pages 111–118, New York, NY, USA, 2010. ACM.

[12] J. Murayama, L. Bougrila, Y. K. Akahane, S. Hasegawa, B. Hirsbrunner, and M. Sato. Spidar g&g: A two-handed haptic interface for bimanual vr interaction. In *Proceedings of EuroHaptics 2004*, pages 138–146, 2004.

[13] N. Stylopoulos and D. Rattner. Robotics and ergonomics. *Surg Clin North Am*, 83(6):1321–1337, December 2003.

[14] D. Swapp, V. Pawar, and C. Loscos. Interaction with co-located haptic feedback in virtual reality. *Virtual Reality*, 10:24–30, 2006. 10.1007/s10055-006-0027-5.

[15] D. Tsetserukou. Haptihug: A novel haptic display for communication of hug over a distance. In *EuroHaptics (1)*, pages 340–347, 2010.

[16] M. Veit, A. Capobianco, and D. Bechmann. Consequence of two-handed manipulation on speed, precision and perception on spatial input task in 3d modelling applications. *Journal of Universal Computer Science (j-jucs), Special Issue on Human-Computer Interaction*, 14(19):3174–3187, 2008.

[17] L. N. Verner and A. M. Okamura. Force & torque feedback vs force only feedback. In *WHC '09: Proceedings of the World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 406–410, Washington, DC, USA, 2009. IEEE Computer Society.

[18] S. Wang and M. A. Srinivasan. The role of torque in haptic perception of object location in virtual environments. In *HAPTICS '03: Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'03)*, page 302, Washington, DC, USA, 2003. IEEE Computer Society.

[19] R. Weller and G. Zachmann. Kinetic separation lists for continuous collision detection of deformable objects. In *Third Workshop in Virtual Reality Interactions and Physical Simulation (Vriphys)*, Madrid, Spain, 6–7 November 2006.

[20] R. Weller and G. Zachmann. Inner sphere trees for proximity and penetration queries. In *2009 Robotics: Science and Systems Conference (RSS)*, Seattle, WA, USA, June/July 2009.

[21] R. Weller and G. Zachmann. A unified approach for physically-based simulations and haptic rendering. In *Sandbox 2009: ACM SIGGRAPH Video Game Proceedings*, New Orleans, LA, USA, Aug. 2009. ACM Press.

[22] R. Weller and G. Zachmann. 3-dof vs. 6-dof - playful evaluation of complex haptic interactions. In *IEEE International Conference on Consumer Electronics (ICCE), 2011 Digest of Technical Papers*. IEEE Computer Society, Jan. 2011.