

3D Worlds for Collaborative Work

Leiden Institute of Advanced Computer Science

Bart van der Drift

Table of contents

1 – Introduction	–	–	–	–	–	–	–	–	–	Pg. 2
2 – What is presence?	–	–	–	–	–	–	–	–	–	Pg. 3
3 – What influences the amount of presence felt?	–	–	–	–	–	–	–	–	–	Pg. 4
4 – A comparison between presence and pretense	–	–	–	–	–	–	–	–	–	Pg. 6
5 – State of the art virtual reality	–	–	–	–	–	–	–	–	–	Pg. 7
6 – State of the art Collaborative Virtual Environments	–	–	–	–	–	–	–	–	–	Pg. 9
7 – CVEs in the future	–	–	–	–	–	–	–	–	–	Pg. 11
8 – Research question	–	–	–	–	–	–	–	–	–	Pg. 13
9 – Tools and Engines	–	–	–	–	–	–	–	–	–	Pg. 15
10 – Prototype and experiment	–	–	–	–	–	–	–	–	–	Pg. 17
11 – Conclusion and future research	–	–	–	–	–	–	–	–	–	Pg. 20
References	–	–	–	–	–	–	–	–	–	Pg. 21

1 – Introduction

Knowledge workers are less bound to one office than they used to be in the past. Take, for example, a technology consultant. Such a consultant will be abroad, on the road, at clients, or at home a large percentage of his/her time. Other examples of work that is highly distributed across the globe are interdisciplinary research and large research projects. Typically, persons involved in such research projects are working on the same topic but living thousands of kilometers apart from each other. One option is to travel a lot, but this has obvious disadvantages (cost, time, environmental pollution, jet lag, etc.). However, physical presence (or at least the impression of physical presence) is a very important part of collaboration, and therefore one of the major benefits of traveling as a solution to spatially distributed projects. In the recent years there has been a surge of technical solutions to facilitate remote working, such as video conferencing, *Second Life* meetings, wiki pages, blogs, etc. The purpose of this project is to investigate what the impact of these technologies are for the future workspace. How can these technologies facilitate work effectiveness, social bonding, pleasure in work, etc.


In this project we will look at one specific aspect of *Computer Supported Collaborative Work* (CSCW), i.e., the potential role of 3D worlds for collaborative remote work. An important question we will try to answer is to what extent such worlds can give the impression of presence, how this impression arises and which factors have an influence on the amount of presence felt. We will also try to make a comparison between presence and pretense, since we believe there are some interesting similarities between them. Furthermore, we take a look at Virtual Reality and Collaborative Virtual Environments that are around today, and what they might bring in the future. We then propose some experiments, discuss their positives and negatives, and explain what choices we made for our experiment. We describe the tools that are available to develop the virtual world needed for this experiment, from general game engines to collaboration-based tools, and elaborate on which one we chose for our research. Finally, we present the prototype we made and propose a way to execute the experiments with our product.

This is the report of the Bachelor Project of Bart van der Drift (Leiden Institute of Advanced Computer Science/LIACS), supervised by Joost Broekens (Telematica Institute) and Walter Kosters (LIACS). I'd also like to thank Gerwin de Haan from TU Delft, who helped a great deal with finding a good research question and coming up with possible experiments, as well as some other things, and the Multiverse community, who helped me out when I was stuck, and gave some great advice when I needed it.

2 – What is presence?

A human has to think ahead in a lot of situations. When you are about cross an unstable looking bridge, you might imagine what would happen if the bridge would collapse. This means you are visualising a possible future reality in which the bridge will break down. A human is constantly creating mental models about his/her environment for existing options through cognitive processes [1]. When you are concentrating on such a mental model, you are disregarding certain sensory input information from the current reality, such as the visual observation that the bridge is still intact [2]. When sensory input from the current reality is disregarded, the feeling of actually being in the non-existing mental-model world can arise.

These mental models are not necessarily products of reality. They can also originate from a non-existing world. This can be, for example, an imaginary world (fantasy or dream), a described world (book) or a virtual world (Virtual Reality or videogame). In the context of a virtual world, disregarding reality means ignoring the hardware, audio from the real world and other ‘real’ things that do not correspond to the virtual world. When all these things happen, we get a feeling of *presence* [3]. The process that leads to presence can be schematically written down as follows [3]:



3 – What influences the amount of presence felt?

Since both a real and imaginary world can give rise to presence, but the intensity can be different, there must be certain factors that influence the amount of presence felt. Following the diagram above, these factors must have an influence to some degree on one or more of the following: the world, the user's cognitive processes, the user's interpretation of the generated mental models, the amount of focus of the user and the extend to which the user disregards reality. Two concepts were introduced by Witmer and Singer to divide these factors.

The first concept is *immersion*. They define (psychological) immersion as “a psychological state characterized by perceiving oneself to be enveloped by, included in, and interacting with an environment that provides a continuous stream of stimuli and experiences” [4]. Immersion seems to be a result of the world, cognitive processes and interpretation of mental models.

The second concept is *involvement*: “Involvement is a psychological state experienced as a consequence of focusing one's energy and attention on a coherent set of stimuli or meaningfully related activities and events” [4]. This can be seen as a combination of focus and disregarding reality.

A user can be immersed in a virtual environment in two ways. First, through the system displaying the sensory data depicting his or her surroundings. Important is the way the surroundings and the user's body are displayed with respect to the users position and orientation in the virtual world. Secondly, the way the user interacts with the world is important [5]. Body tracking devices, such as electromagnetic sensors allow the user to act in a natural way [6]. When the user is using these tools, the mental models created concerning the virtual movement are consistent with the mental models concerning the user's body movement [5].

Research indicates that using video to communicate in a collaborative world results in more presence than using audio, and using audio results in more presence than using text based communication [7].

The use of Augmented Reality results in a stronger sense of presence over desktop systems. However, a stronger sense of presence does not necessarily lead to more engagement. Here we assume engagement and involvement are more or less the same. Research has even indicated that a stronger sense of presence can obstruct the engagement, because a more realistic experience leads to higher expectations, since users relate the experience to a real-life situation, instead of a simulated situation. These expectations can not always be met, due to bugs and haptic inconsistencies (“I should have felt them when I reached out my hand”), which disappoints the users. When using desktop systems, users expect less. It should be noted however, that the test users of these experiments had no experience with Augmented Reality when the experiments took place, and didn’t know what to expect. It is suggested that “mediation may be necessary for some players to fully engage with certain interactive media experiences” [8].

A factor enhancing involvement is multimodality: combining visual, auditory and haptic sensory input. Trimodal is better than bimodal, and bimodal is better than unimodal. The reason for this is that “multimodal VE users start their cognitive process faster, thus, in a similar exposure time they can pay attention to more informative cues and subtle details in the environment and integrate them creatively” [9].

Another factor determining the amount of involvement is the nature of the relationship between users, in case of a multi-user system. Research concerning videogames indicated that “when compared to playing against a computer, playing against another human elicited higher spatial presence, engagement, anticipated threat, post-game challenge appraisals, and physiological arousal, as well as more positively valenced emotional responses. In addition, playing against a friend elicited greater spatial presence, engagement, and self-reported and physiological arousal, as well as more positively valenced facial EMG responses, compared to playing against a stranger” [10].

4 – A comparison between presence and pretense

Children learn to play roles and pretend at a very early age. Three important developments enable preschoolers to create joint make-believe worlds with others [11]. These are the abilities to:

- (1) manage multiple roles as playwrights and actors,
- (2) invent novel plots, and
- (3) deliberately blur the boundary between reality and pretense.

The first ability is not specifically relevant to presence. However, the second and third abilities show interesting similarities with the processes that cause a sense of presence. Inventing plots is very similar to the process of imagining possibilities from certain situations, which is done by giving interpretations to mental models. Deliberately blurring the boundary between reality and pretense is very similar to disregarding properties of the reality that do not correspond to the world being imagined.

Pretense is different from ‘Error acting as if’. The latter is, for example, when you jump up from your chair because you think you see a spider, while there actually is no spider at all [12]. An interesting question is whether presence is pretence, or error acting as if. In other words: Do we feel presence by *deliberately* pretending to be in a virtual world, or do we falsely believe that we are inside a world that is not real? How conscious is this process? Perhaps the process starts consciously, and slowly turns into error acting as if, as you more and more forget about the world around you. We will not try to give an answer to this question in this paper, we just leave this interesting question open for possible future research.

5 – State of the art virtual reality

These days, the technology behind virtual reality has advanced to a point where we are able to develop productive and meaningful virtual reality applications. Many new virtual reality systems are developed every day, for more and more applications. For this reason, it is very difficult to say what state of the art virtual reality exactly is at this moment. However, we will try to give an impression of what's hot in the world of virtual reality at this moment, and what new things are developed, while keeping the focus on presence-related research. Virtual reality can be used for many different purposes, such as education or training, therapy and entertainment.

Educational applications of virtual reality have shown that training in these environments can significantly improve the performance of surgeons in the operating room. These surgeons work faster, make less errors and show improvement in their economy of movement [13].

Virtual reality training of pilots has been very common for years. Virtual reality can even be used to improve the roadside crossing judgments of children [14].

In therapy, virtual reality can be a very effective tool to treat a great variety of diseases and disorders. Researchers are finding that some of the best applications of virtual reality focus on therapy. Presence can be of great value in medical sessions. Virtual realities have even shown to be a remedy against pain, both physical and psychological. For example, virtual reality has been used to relieve discomfort of burned patients during treatment. Other well-known medical applications are treatments against phobias and Post-Traumatic Stress Disorder [15].

At this time, virtual reality is not available for the masses yet. However, there are a lot of applications of virtual reality in the field of entertainment in theme parks and cinemas. For example, Disney has been researching virtual reality since 1992. They learned many lessons including: “Physical interfaces help to immerse guests in an experience”, “Facile controls are critical for guest acceptance; focus on the guest's natural skills” and “Immersive theatres are important tools for virtual design” [16].

Recent years, the addition of smell to virtual reality has frequently been a subject of research. Smell gives users a uniquely compelling experience of presence in the virtual world [17]. Most attempts to realize these ‘olfactory displays’ have involved capturing and synthesizing an odour. These processes still pose many challenging problems. The

difficulty is mainly due to the way human olfaction works. A set of so-called ‘primary odours’, of which all other odours can be made, has not been found. Also, many existing interactive olfactory displays simply diffuse the scent into the air, which does not provide the ability of spatio-temporal control of olfaction. However, the use of tubes or air cannons in combination with nose tracking can provide this ability [18].

Another important addition to virtual reality is touch. The majority of today's haptic interfaces are designed for hand-based interaction with virtual environments. However, there are several real-life tasks that require a person to interact with the environment using one's feet. Researchers have developed systems to simulate walking in a virtual environment, as well as an approach to foot based interactions, intended for users in sitting position [19].

Recently, Apple has filed a patent for a three-dimensional display system. This “three-dimensional display system provides a projection screen having a predetermined angularly-responsive reflective surface function. Three-dimensional images are respectively modulated in coordination with the predetermined angularly-responsive reflective surface function to define a programmable mirror with a programmable deflection angle.” [20]

It could be very interesting if Apple would get involved into virtual reality development. They might be able to reach a large variety of people, and maybe even make virtual reality technology available to the masses over time.

6 – State of the art Collaborative Virtual Environments

As the technology of virtual reality has developed, so have the systems supporting collaborative work based on these technologies. These systems are called Collaborative Virtual Environments (CVEs). CVEs that use a visual representation of the participants can roughly be divided into two categories: those that use video feeds of the participants, and those that use avatars to depict the participants. Both approaches have their advantages and disadvantages.

By capturing video of a participant, a one on one representation of the person is generated. However, this technique does not allow views from different perspectives (without changing the camera position), and hence cannot be considered fully 3D. Another problem that arises when video feeds are used is the difficulty of establishing mutual eye contact, when the camera and display are not positioned in one line. This issue can be solved by using multiple cameras from different angles [21].

Using avatars allows participants to be represented in full 3D, and mutual eye contact can easily be established. Even though rendering real-time avatars that are indistinguishable from real life persons is yet to be achieved, technologies behind this are developing at a high rate, largely influenced by the videogame industry. Many believe that user embodiment, the provision of users with appropriate body images so as to represent them to others and also to themselves, is a key issue for collaborative virtual environments [22].

One of the first CVEs was DIVE. “The Distributed Interactive Virtual Environment (DIVE) is an internet-based multi-user VR system where participants navigate in 3D space and see, meet and interact with other users and applications” [23]. This system is accessible from a desktop PC but can also be used immersively. It comprises many disjoint virtual worlds, linked by portals, and world and object descriptions can be downloaded on demand from the Web.

Some years after DIVE, MASSIVE was released [24]. This system mainly focuses on videoconferencing over WANs. MASSIVE-2 and MASSIVE-3 have been used to create a variety of systems, including public participation in online art and performance. This has

also led to the idea of inhabited television [25], which is a fusion of CVEs and traditional television.

NPSNET is a networked 3D virtual environment developed by the U.S. Navy. NPSNET-IV was the first 3D virtual environment that incorporated both the IEEE 1278 Distributed Interactive Simulation (DIS) application protocol and the IP Multicast network protocol for multi-player simulation over the Internet [28].

SPLINE is an environment for online communities. It can be used to create end-user CVE applications. The construction mechanism of SPLINE is very interesting, for example allowing buildings that are larger from the inside than from the outside [25]. However, there doesn't seem to be a lot of information about this system on the web anymore.

The Blaxxun Community Platform is another tool for managing and visiting communities on the web. You can create your own 'second life' like community. Even a mobile version for smart phones and PDAs is available [29].

Croquet is an open source software development environment and software infrastructure for creating and deploying deeply collaborative multi-user online applications and metaverses on and across multiple operating systems and devices [30].

Project Wonderland is a toolkit for creating collaborative 3D virtual worlds. Within those worlds, users can communicate with high-fidelity, immersive audio, and can share live applications such as web browsers, OpenOffice documents, and games [31].

The relationship of members of a team collaborating through a CVE is very important. Team cohesion and trust have a positive effect on the performance of a team. Virtual teams seem to be more task-focused and less social-focused than traditional teams, although this task-focus seems to lessen over time. In global virtual teams, cultural and language differences are very common. This can lead to difficulties that have a negative impact on the collaboration activities. Even differences among team members from different regions of the same country can cause problems. These problems can be solved by actively understanding and accepting the differences. Another factor that determines the performance of a virtual team is technical expertise of the members. A lack of technical expertise and the inability to cope with technical problems have a negative effect on individual satisfaction with the team experience and performance. Consistent training among team members can improve the performance of the team [26].

7 – CVEs in the future

We see that virtual reality is becoming more regular. In research, but also in real life applications. Although virtual reality is not accessible to the masses yet, it is expected to be in the future. A lot of thinking is done about what the role of virtual collaborative environments can be, and how these technologies might further develop. It is possible that complex systems able to filter, manipulate and generate nonverbal or even verbal interaction will exist some day [27].

For example, when engaging in a virtual environment using rendered avatars, automatically mimicking others, regulating eye contact, body language and even facial expressions, or collecting data about interactions are all theoretically possible. Rendering a consistent image to every participant may not even be necessary. In other words, one participant can be shown something different than the other. For example, when a leader is negotiating with member A and member B, and A has his legs crossed and B has his arms crossed, the leader could be crossing his legs as well in the world presented to A, and crossing his arms in the world presented to B at the same time, in order to mirror both members. It would also be possible to have eye contact for 70% of the time with both A and B, or different nonverbal communication could be presented to different persons at the same time [27]. In the field of 2D CSCW environments it is suggested that a ‘what you see is what I see’ approach can actually hinder collaboration [25], so this may very well be applicable to CVEs as well.

Another possibility would be changing the representation of your face in the virtual environment to look more like your conversation partner, a person with whom he/she maintains a deep relation of trust, or someone famous, such as a politician or a religious person [27]. The face could even be constructed from scratch, trying to create an avatar that incorporates some user defined qualities. Besides changing the looks and behaviour of avatars, it would also be possible to present the environment in an arbitrary manner, for example by rearranging the position of participants, or changing the room colours or decoration to give participants a feeling of comfort. Again, different situations can be presented to different participants at the same time, managing the environment for each member separately [27].

Before CVEs will have the chance to become available to the masses, developers will most certainly run into many problems that have to be solved. Benford et al. enumerated four challenges that will have to be faced [25].

The first one is scalability. Supporting many different users to interact in an environment is very complex and depends of multiple system bottlenecks. Online gaming has shown that it is possible to allow a large group of users to interact with each other in real-time, highly detailed environments. However, when supporting collaborative communication, a lot more data needs to be sent, for example information about body language and high quality audio samples, and delay should be minimal.

The second challenge is distributed architectures. Multiple architectures are used to update users with the changes of the world, which can be divided into three main categories. The first one is server/client, in which each participant's application communicates only with a common server program that is responsible for passing messages on to other clients as appropriate. The second one is peer-to-peer unicast, in which each participant's application sends information to other clients as appropriate.

The third one is peer-to-peer multicast, in which the information is sent to all other clients. The third challenge is migrating lessons from 2D interfaces and CSCW. One of these lessons appears to be the hindrance caused by 'what you see is what I see' approaches. Another strongly connected topic is the "space-versus-place debate", in which some people argue CVEs are supposed to be systems in which participants can move independently within a shared coordinate system, where others are displayed as avatars (space), and other people argue that social behaviour is dependent on other important aspects than just these relative positions (place).

The fourth challenge is concerned with new kinds of human factors, which deals with understanding the nature of social interaction in CVEs. This will probably cause new problems in the future since most VR studies have focused on single-user systems until now.

8 – Research question

A lot of research in the field of VR and CVEs is still to be done. We decided to focus on the following research question:

‘What is the effect on the presence felt in a CVE where realistic sound effects are used opposed to a CVE that does not use any sound?’ or ‘What is the effect on the performance of a collaborative task in a CVE where realistic sound effects are used opposed to a CVE that does not use any sound?’. These days, video games all incorporate realistic sound effects. However, this is not always the case for CVEs. It would be interesting to see what effect using realistic sound would have.

We came up with seven different experiments to test our research questions:

1. Rafting: We would like to create a CVE in which two participants are to row across a water. One participant would be on the left, and one on the right, making collaboration a necessity to go straight. Participants are *not* allowed to talk. We assume that participants that can collaborate better would finish in a smaller amount of time. By adding realistic sound effects to the world, especially concerning the paddles touching the water, of half of the pairs of participants, and measuring the time they need to finish, we could investigate whether sound has a positive effect on the performance.
2. Tubes: We would like to create a CVE in which one participant stands in a round chamber, where a certain amount of coloured tubes are located equally divided around him. A second participant is to throw objects down the tubes at the other end in some predefined order, which the first participant is to catch. The participants are to catch a certain amount of objects, after which the clock is stopped. The participants are allowed to talk. By adding realistic sound effects to the world, especially concerning the falling of the objects through the tubes, of half of the pairs of participants, and measuring the time they spend to complete their task, we could investigate whether sound has a positive effect on the performance.
3. Sawing: We would like to create a CVE in which two participants are to saw through a tree using a two-man saw, as fast as possible. To succeed in this, we believe extensive collaboration is required. However, participants are *not* allowed to talk. By adding realistic sound effects to the world, especially concerning the sawing, of half of the pairs of participants, and measuring the time it takes to saw

a tree in half, we could investigate whether sound has a positive effect on the performance.

4. 3D puzzle: We would like to create a CVE in which two participants are to build a cube out of its pieces, which are scattered around an office room, as fast as possible. Participants are allowed to talk. By adding realistic sound effects to the world, especially environmental sounds, such as the sound of a running computer, of half of the pairs of participants, and presenting them a presence questionnaire afterwards, we could measure if realistic sound effects enhance the feeling of presence in a CVE. By measuring the time it takes to complete the puzzle, we could also measure whether realistic sound effects enhance the performance.
5. Search: We would like to create a CVE in which two participants have to find a certain object in an office room as fast as possible. Participants are allowed to talk. By adding realistic sound effects to the world, especially environmental sounds, such as the sound of a running computer, of half of the pairs of participants, and presenting them a presence questionnaire afterwards, we could measure whether realistic sound effects enhance the feeling of presence in a CVE. By measuring the time it takes to find the object, we could also measure whether realistic sound effects enhance the performance.
6. Maze: We would like to create a CVE in which one participant finds himself in a maze, unable to see where to go, and a second participant with an on-top view, who has to guide the first participant through the maze. This can be done by talking, but half of the pairs of participants can also use 3D sound cues for guidance. This is done by clicking with the mouse at a certain location while talking through the microphone. The person in the maze would hear the sound from some direction. By measuring the time it takes to complete the maze, we could also measure whether 3D sound effects enhance the performance.
7. Asynchronous search: We would like to create a CVE similar to experiment 5, however here we would like two participants to search for some special documents in an office room, while there are two types of documents. One type can be viewed only by the first participant, and the other type can be viewed only by the second participant. This is done to enhance the need for collaboration.

For the sake of realism, we considered the use of motion tracking devices such as the Nintendo Wii Remote [37]. Especially for experiments 1, 2 and 3, this would be a valuable addition.

Even though we found all of these experiments interesting, and certainly worth the effort, we decided to proceed with experiment 3. We chose this experiment because we thought its feasibility and simplicity would lead to clear and effectively measurable results, which would make any conclusions more plausible, and because it allowed the use of the Nintendo Wii Remote. We have later changed the sawing of a tree into the cutting down of a tree, due to technical difficulties. However, the principle remained the same.

9 – Tools and Engines

A quick search on the internet tells us that there are a lot of 3D engines out there. For some you need a license, for some you don't. Some are even open source. For our research, we are looking for a development tool to create a prototype of a collaborative system. For these reasons we started searching for a 3D engine that allows extensive modification in a simple way, has a large enough community to get information, has the networking taken care of, and is preferably collaboration-oriented. This greatly reduces the choices we have, but still a lot of options exist. We will enumerate the tools that we considered.

Croquet [30] seems to have everything we wished for. It is open source, which means excellent modifiability, the community is quite large, and it is designed to help develop deeply collaborative programs.

Project Wonderland [31] is almost entirely open source, and it is also intended to be used to make collaborative systems. The major downsides are that it doesn't really exist of any tools of its own: it greatly depends on external tools such as Blender, a 3D content creation tool. Also, there doesn't seem to be a lot of third party support, such as tutorials and guides.

Panda3D [33] is a free 3D engine, largely focussing on game development. There are a lot of tutorials, sample programs etc. on the website. The networking seems to be very low-level though.

Smogworks [34] is an open source 3D game engine. Multiplayer, as well as audio and video, played an important role in the development. The engine doesn't really seem to be in a very advanced state, and the community seems to be small, and no tutorials are provided.

Multiverse [35] is an engine intended for Massively Multiplayer Online (MMO) games. Plenty of tutorials can be found on the website as well as on other sites.

Second Life [36] is a well known 3D world. Its community is enormous and a huge amount of tutorials and guides can be found all over the net. However, it is not a tool to create worlds from scratch. For this reason, it may prove difficult to build exactly what we want.

The table below shows which tools meet the requirements:

Tool/Engine Name	Free of charge	Open Source	Community support	Networking	Collaboration oriented
Croquet	✓	✓	✓	✓	✓
Wonderland	✓	✓	✗	✓	✓
Panda3D	✓	✗	✓	≈	✗
Smogworks	✓	✓	✗	✓	✗
Multiverse	✓	✗	✓	✓	✗
Second Life	✓	✗	✓	✓	✗

On paper, Croquet seems to be the best choice. However, we received some negative opinions about this tool from people who had been working with it for a while, and our first few hours of experience with it couldn't quite convince us either. We decided to use Multiverse. Although not completely open source, the modifiability would be good enough for our project, and though not collaboration oriented, the focus on MMO convinced us the networking and the possibilities for interaction would be sufficient.

10 – Prototype and experiment

We successfully managed to create a prototype with all the functionality we desired. This was achieved by modifying the Multiverse server code, server scripts and client scripts, and using GlovePIE [38] for the Nintendo Wii Remote input scripting.

We ran into some problems during development. We realized that Multiverse is not a finished product yet, when we ran into things that weren't working as they were supposed to, as well as incorrect documentation. We especially experienced problems concerning the modelling and animation, for which the conversion tools weren't working as desired. Some models couldn't be converted at all. Especially larger models couldn't be handled by the conversion tool. The models that could be converted, still gave some problems such as incorrect lighting, wrongly named textures, and so on, forcing us to edit the material files manually. In the end we mostly used models and animations that were developed by the Multiverse team, so they didn't need any conversion or modification. We did change the cutting animation scriptwise to match the motion and the sound, but this also gave some problems as the animation API did not react as it was supposed to according to the documentation. Despite all the setbacks, we now have a system with all the behaviour that we desired. We will describe the behaviour of the prototype by means of the following scenario:

Participant A is sitting behind a desktop system. He enters the virtual world. He finds himself in a world where a tree is standing in front of him (Figure 1). Participant B, who is physically located in another room, and on another desktop system, enters the world as well, at the other side of the tree from where participant A is standing. Both participants can (partially) see each other from a first person perspective on the screen, as well as the tree. They now know they can start cutting down the tree. One of the participants starts by hitting the tree. This is done by simulating a cutting-like motion with the Nintendo Wii Remote. Which participant starts is not regulated. At the first hit, the time is started. The participants are now supposed to take turns on hitting the tree. If they don't follow this rule, in other words, a participant hits the tree two or more times in a row, a penalty point will be added for each mistake made. After a certain amount of hits, the task will be completed. The time and the amount of penalty points are shown on the screen (Figure 2). The participants leave the world.



Figure 1: Participant A has entered the world



Figure 2: The task has been completed

In the scenario above everything is completely automated, so no supervisory intervention is needed. Only the setup of the world and the user accounts is advised to be done by a supervisor. Since the prototype was intended to be used to investigate the influence of sound on the performance of the task, the data shown on the screen needs to be collected by the supervisor after each experiment. When data is collected about two large groups, where half of the groups had the sound enabled and the other half disabled, the data can be interpreted and compared, and conclusions can be drawn about the effect of sound on the performance.

Because two different statistics are collected, namely the elapsed time and the mistakes made, the researcher will have to decide how to combine this data into one unit representing the performance. How this is done greatly depends on how important he/she considers the mistakes to be. One way would be to add the amount of mistakes made to the elapsed time, or multiply the mistakes by some 'mistake significance rate' before adding them to the time, in order to influence how much effect the mistakes have on the performance rating. Many other methods are possible, however. After comparing the statistics for both groups, the effect of sound on the performance can be determined.

A lot of other variants on this experiment can be thought of. It would be possible to vary how large a part of the other participant is visible. This would change the amount of information which is visually presented to the participants, and might influence the usability of the sound for the task. Another possibility is allowing the participants to speak to one another using a microphone. This would enhance the complexity of the collaboration, as well as the diversity of the presented audio. For this reason, the sound of the axe hitting the tree might be processed in a more subconscious way, which could change the influence of the sound on the performance. It can be questioned, however, just how much verbal communication is needed for this simple kind of task. Another way

to push the processing of the audio to a more subconscious level might be by presenting more distracting and/or confusing elements to the participants, for example by adding moving objects to the background, or making the participants hit the tree in a defined place, which changes after each hit. This would probably force the participants to concentrate more on the task and forget about the sound. The complexity of the task can also be increased by making the participants cut down several trees in a row, without predetermining which trees they have to cut down. This would require some extra collaboration, because the participants would have to decide in some way which tree to go to next. Of course the participants would have to be able to move for this experiment, which may be difficult when they are already holding a Wii remote. Of course exclusively the keyboard could be used in this case, but this all requires a basic understanding of the movements in the world, which has to either be taught to the participants in advance, or they should have some experience in first person systems already. This problem does not exist with our original experiment, since the operation is only invoked by using an intuitive motion with the Wii remote.

We might also be interested in the extent to which the participants had the feeling they ‘were there’, in the virtual world. How much presence did they feel? It would be possible to compare the presence felt by the groups that are hearing sounds and those that are not. Some widely accepted questionnaires exist for measuring presence [5]. These questionnaires give a certain numerical rate of presence based on the answers given to the questions in the questionnaire. Comparing these rates between the two groups, we could decide whether sound can have an influence on the presence felt or not.

Perhaps it would even be possible to find a correlation between performance and presence. Especially when the sound does not give any extra information, but only a confirmation of what is already perceived via visual stimuli. This is because if the sound would give information that is not, or hardly, visually noticeable, this would probably increase the performance, but (partially) because of the extra information obtained, and not because of the presence of audio by itself. And if the presence would also be increased by the addition of sound, this could be caused by the audio itself, and not by the information given by this audio. In other words, when the sound is being turned on and off, this is actually not a single factor that changes, but two. The first factor is the presence of audio, and the second factor is the addition of the information that is given by this audio. Then we could conclude that both collaboration and presence get enhanced by the audio, but not that there is a direct relation between presence and collaboration. We believe, however, that the sound in our prototype does not give any extra information, since the fact that the axe is hitting the tree can also be observed in a visual manner. The sound merely confirms what can already be seen. If the performance is better for the group with sound, this could be an indication that the addition of sound enhances the presence, which on its turn increases the performance. Recall that multimodal systems have shown to give rise to more presence than unimodal systems [9].

11 – Conclusion and future research

In this project, we wanted to demonstrate what the effect of sound is on the performance of a task in a CVE. We demonstrated that it is possible to build a prototype in which our research question can be tested by using Multiverse. This makes it very likely that some, if not all, of our other experiments we proposed in Chapter 8 can be executed using the same engine and approach. Because we did not have the chance to execute the experiments, we cannot give an answer to our research question. However, the fact that we did realise our prototype opens a lot of possibilities for future research. Not only research concerning sound in CVEs, but also other factors in other types of environments can be examined. For example the influence of excessive colour use in Augmented Reality applications on the performance, and the effects of ambient music in virtual meetings, just to name a few. All these things can be a valuable addition to the knowledge of virtual reality and 3D worlds we have at this moment.

References

- [1] D.A. Norman
Mental models
Lawrence Erlbaum Associates, Inc., Publishers, 1983
- [2] A.M. Glenberg
What memory is for
Behavioral and Brain Sciences 20 (1997) p. 1-55
- [3] T. Schubert, F. Friedmann
The experience of presence: Factor analytic insights
Teleoperators and Virtual Environments 3 (2001) p. 266-281
- [4] B.G. Witmer, M.J. Singer
Measuring presence in virtual environments: A presence questionnaire
Presence: Teleoperators and Virtual Environments 3 (1998) p. 225-240
- [5] M. Slater, M. Usoh
Body centred interaction in immersive virtual environments
Artificial Life and Virtual Reality, 1994
- [6] R. Kalawsky
The science of virtual reality and virtual environments: A technical, scientific and engineering reference on virtual environments
Addison-Wesley Publishing Company, 1993
- [7] E. Sallnäs
Effects of communication mode on social presence, virtual presence, and performance in collaborative virtual environments
Presence: Teleoperators and Virtual Environments 4 (2005) p. 434-449
- [8] S. Dow, M. Mehta, E. Harmon, B. MacIntyre, M. Mateas
Presence and engagement in an interactive drama
Conference on Human Factors in Computing Systems (2007) p. 1475-1484
- [9] D. Hecht, M. Reiner, G. Halevy
Multimodal virtual environments: response times, attention, and presence
Presence: Teleoperators and Virtual Environments 5 (2006) p. 515-523
- [10] N. Ravaja, T. Saari, M. Turpeinen, J. Laarni, M. Salminen, M. Kivikangas
Spatial presence and emotions during video game playing: Does it matter with whom you play?
Presence: Teleoperators and Virtual Environments 4 (2006) p. 381-392
- [11] I. Bretherton
Pretense: The form and function of make-believe play
Academic Press, Inc., 1989

- [12] A.M. Leslie
Pretense and representation: The origins of "Theory of Mind"
American Psychological Association, Inc., 1987
- [13] T.P. Grantcharov, V. Kristiansen, J. Bendix, L. Bardram, J. Rosenberg, P. Funch-Jensen
Randomized clinical trial of virtual reality simulation for laparoscopic skills training
British Journal of Surgery 2 (2003) p. 146-150
- [14] J. Thomson, A. Tolmie, H. Foot, K. Whelan, P. Sarvary, S. Morrison
Influence of virtual reality training on the roadside crossing judgments of child pedestrians
Journal of Experimental Psychology: Applied 3 (2005) p. 175-186
- [15] H.G. Hoffman
Virtual Reality Pain Distraction
APS Bulletin 2 (2005)
- [16] J. Deisinger, A. Kunz
Towards virtual reality for the masses: 10 years of research at Disney's VR studio
ACM International Conference Proceeding Series; Proceedings of the workshop on Virtual environments (2003) p. 11-17
- [17] H.G. Hoffman, T. Richards, B. Coda, A. Richards, S.R. Sharar
The illusion of presence in immersive Virtual Reality during an fMRI brain scan
CyberPsychology & Behavior 6 (2003) p. 127-131
- [18] Y. Yanagida, S. Kawato, H. Noma, A. Tomono, N. Tesutani
Projection based olfactory display with nose tracking
Proceedings of the IEEE Virtual Reality (2004) p. 43- 50
- [19] R.F. Boian, J.E. Deutsch, C.S. Lee, G.C. Burdea, J. Lewis
Haptic effects for Virtual Reality-based post-stroke rehabilitation
11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems/HAPTICS'03 (2003) p. 247
- [20] S. Hotelling et al.
Three-dimensional display system
United States Patent 20080068372, March 20, 2008
- [21] *GazeMaster gaze-corrected videoconferencing project*
<http://research.microsoft.com/BARC/mediapresence/GazeMaster.aspx>
[Accessed at 01-09-2008]
- [22] S. Benford, C. Greenhalgh, J. Bowers, D. Snowdon, L.E. Fahlén
User embodiment in Collaborative Virtual Environments
Conference on Human Factors in Computing Systems (1995) p. 242-249
- [23] DIVE
<http://www.sics.se/dive/>
[Accessed at 01-09-2008]
- [24] MASSIVE
<http://www.crg.cs.nott.ac.uk/research/systems/MASSIVE/>
[Accessed at 01-09-2008]

- [25] Steve Benford, Chris Greenhalgh, Tom Rodden, James Pycock
Collaborative Virtual Environments
Communications of the ACM 7 (2001) p. 79-85
- [26] A. Powell, G. Piccoli, B. Ives
Virtual teams: A review of current literature and directions for future research
ACM SIGMIS Database 1 (2004) p. 6-36
- [27] J.N. Bailenson, A.C. Beall, J. Loomis, J. Blascovich, M. Turk
Transformed social interaction: Decoupling representation from behavior and form in Collaborative Virtual Environments
Presence: Teleoperators and Virtual Environments 4 (2004) p. 428-441
- [28] NSPNET at CiteSeer
<http://citeseer.ist.psu.edu/macedonia94npsnet.html>
[Accessed at 01-09-2008]
- [29] Blaxxun
<http://www.blaxxun.com/>
[Accessed at 01-09-2008]
- [30] Croquet
http://opencroquet.org/index.php/Main_Page
[Accessed at 01-09-2008]
- [31] Project Wonderland
<https://lg3d-wonderland.dev.java.net/>
[Accessed at 01-09-2008]
- [32] P.N. Hamid, A.G. Newport
Effect of colour on physical strength and mood in children
Perceptual & Motor Skills 5 (1989) p. 179-85
- [33] Panda3D
<http://panda3d.org/>
[Accessed at 01-09-2008]
- [34] Smogworks
<http://www.smogworks.com/>
[Accessed at 01-09-2008]
- [35] Multiverse
<http://www.multiverse.net/>
[Accessed at 01-09-2008]
- [36] Second Life
<http://secondlife.com/>
[Accessed at 01-09-2008]
- [37] Nintendo Wii Remote
<http://www.nintendo.com/wii/what/controllers#remote>
[Accessed at 01-09-2008]
- [38] GlovePIE
<http://carl.kenner.googlepages.com/glovepie>
[Accessed at 01-09-2008]