

Using Teleporting, Awareness and Multiple Views to Improve Teamwork in Collaborative Virtual Environments

Trevor J. Dodds[†] and Roy A. Ruddle[‡]

Visualization and Virtual Reality Research Group
School of Computing, University of Leeds, UK

Abstract

Mobile Group Dynamics (MGDs) are a suite of techniques that help people work together in large-scale collaborative virtual environments (CVEs). The present paper describes the implementation and evaluation of three additional MGDs techniques (teleporting, awareness and multiple views) which, when combined, produced a 4 times increase in the amount that participants communicated in a CVE and also significantly increased the extent to which participants communicated over extended distances in the CVE. The MGDs were evaluated using an urban planning scenario using groups of either seven (teleporting + awareness) or eight (teleporting + awareness + multiple views) participants. The study has implications for CVE designers, because it provides quantitative and qualitative data about how teleporting, awareness and multiple views improve groupwork in CVEs.

Categories and Subject Descriptors (according to ACM CCS): C.2.4 [Computer-Communication Networks]: Distributed Systems – Distributed applications; H.1.2 [Models and Principles]: User/Machine Systems – Human factors; Software psychology; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems – Artificial, augmented and virtual realities; H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – Collaborative computing; Computer-supported cooperative work; Synchronous interaction; I.3.7 [Computer Graphics]: Three Dimensional Graphics and Realism – Virtual Reality

1. Introduction

Previous research developed techniques called Mobile Group Dynamics (MGDs), which helped groups of people work together while they traveled around large-scale collaborative virtual environments (CVEs) [DR08]. Compared to a conventional CVE, these techniques led to a seven-fold increase in the amount of communication that took place between participants. However, two major areas for improvement were also identified.

First, participants tended to spatially regroup to discuss their findings, even though MGDs allowed communication over an infinite distance (there was no distance attenuation for audio communication between group members). This meant that unnecessary amounts of time were spent traveling to meeting places.

Second, if participants wanted to see what others were looking at (e.g., a point of interest that was being discussed) then they had to ‘walk’ to the appropriate location.

The present paper describes how these shortcomings were tackled by adding new functionality to MGDs. We were reminded that ‘CVEs...do not necessarily need to reflect or embody the characteristics of conventional environments to enable them to support particular forms of activity or interaction’ [FGV*00], and ‘This is the trap VR often falls into; VR tries to imitate reality...’ [Pek02]. Therefore, the new functionality took advantage of the fact that CVEs do not need to be limited to real world constraints. Three types of functionality were implemented.

First, our hypothesis was that participants collocated to communicate with group members because, given that the CVE used directional sound, participants assumed that the audio was also distance attenuated. In fact, distance attenuation was disabled for intra-group communication, as was explicitly stated in participants’ verbal and written instruc-

[†] e-mail: trev@comp.leeds.ac.uk

[‡] e-mail: royr@comp.leeds.ac.uk

tions. To overcome this, we developed ‘awareness’ functionality that provided visual feedback about who was receiving audio at a given moment in time, and who was speaking.

Second, it is well known that in CVEs participants often find it difficult to understand what each other is looking at [HFH*00]. To overcome this, a participant’s own (main) view was supplemented by small viewports that showed the views of fellow group members.

Third, ‘walking’ is time consuming, so teleporting functionality was added. This allowed participants to move directly to any point in the environment by clicking on the appropriate location on the display, or move directly to another participant’s position by clicking on ‘their’ viewport.

This following sections describe further background research, the implementation of the new MGDs functionality, and its evaluation using an urban planning scenario.

2. Background

This section briefly reviews research relevant to teleporting, promoting awareness of others and the use of multiple views in virtual environments. People often wish to spend large amounts of time exploring an environment, be it virtual or real. To minimize the amount of time spent traveling in the real world people may choose to drive rather than wait for a bus, or run instead of walk. In electronic environments, however, we can make use of non-naturalistic forms of movement. For example, using a search engine we can generate hyperlinks that jump straight to task-related pages.

An evaluation of the use of hyperlinks in virtual environments showed that there was a speed-accuracy trade-off when compared to conventional navigation (walking) [RHPJ00]. Put simply, participants found their target locations faster with hyperlinks, but visited more locations in the process. This suggests that the efficiency of teleporting is in the speed, so CVE designers might infer from this that *instantaneous* teleportation should be implemented to get the most out of the speed increase. However, [RHPJ00] also highlights the importance visual continuity. This is something that instant teleportation can lose, e.g. teleporting between self-contained areas in chat-based VEs such as There[†] and Second Life[‡].

Elvins et al. [ENSK01] helped overcome visual discontinuity by providing ‘worldlets’, or small 3D thumbnails, of landmarks which the user could teleport to. This was more useful to users than a 2D image or textual description of the destinations. These worldlets overcame the visual discontinuity that can occur when a user teleports from one place to another, but other techniques are needed to provide information about the spatial relationship of the two places and the

rapid controlled movement used to implement teleporting in the present study is one method of achieving this.

Participants in desktop VEs experience two kinds of problems understanding the actions of others. 1) ‘Fragmented views’, where another participant refers to an object or point of interest in the environment, but their avatar and the point of interest are not simultaneously visible in the viewport [HFH*00]. 2) What you see is *not* what I see, which makes it difficult to understand another’s perspective.

Problem 1 is likely to happen because of a narrow field of view. This is inherent in desktop VEs wanting to minimize distortion to keep realism [FGV*00]. Problem 2 is due to a removal of real world sensory data, such as eye movements and depth perception. This problem tends to be compensated for by a large increase in the amount of verbal communication that takes place [RSJ02].

A combination of these two problems occurs if two users wish to meet at a point of interest. This is a ‘Come here! Look at this’ scenario (see [YO02], p. 136), where the respondent needs to know the location of the user who is talking (they are unlikely to be within the viewport, see problem 1), and what they are referring to (problem 2).

To overcome these problems, Wössner et al. [WSWL02] provided a ‘what you see is what I see’ (WYSIWIS) view in their CVE, which would eradicate problem 2. They designed two CVE interfaces, one of which provided a master/slave style view (where one participant had complete control), and the other which provided a more flexible approach where participants still had some independence (they could change orientation). However, it was found that users preferred the independent viewpoint, so they didn’t interfere with the other participant. Sonnenwald et al. [SWM03] found that users saw a benefit in both independent views and shared perspectives – users liked to be able to figure things out on their own and then discuss them collaboratively. Therefore, our hypothesis was that by providing multiple viewports to the user, MGDs would provide the best of both worlds.


Real time views of other users’ perspectives have been used many times in multiplayer games, such as console games that are designed to be played together on a single display. In this case, however, the views are aimed at different people (e.g. looking at another person’s view could be considered cheating). The ‘split screen’ is used to provide a cheap alternative to the players, instead of requiring that they have multiple visual display units. The present study investigated the provision of multiple views to each user, via a main window and thumbnail views of other participants. Participants were allowed to click on these thumbnails to teleport to the appropriate person.

[†] <http://www.there.com> (Accessed 23/1/2008)


[‡] <http://www.secondlife.com> (Accessed 23/1/2008)

3. Implementing Mobile Group Dynamics

The basic MGDs techniques incorporated an explicit hierarchical grouping system, represented using a ‘group graph’ metaphor, and methods to assist movement as a group. For this study, we added awareness of who was talking and who would receive audio (i.e. participants within hearing range; this included all participants in one’s own group because there was no distance attenuation for intra-group audio communication), teleporting functionality and multiple views (including the ability to teleport to a fellow group member by clicking on their view). This new functionality, along with the group graph, can be seen in Figure 1.



(a) Teleport condition, shown in the over-the-shoulder perspective



(b) Multiple views condition, shown in the bird's-eye view

Figure 1: Screenshots of the environment in the two conditions: teleport and multiple views. The graph metaphor, speech icon, teleporting arrow and participants within hearing range can be seen in both figures. The views of fellow group members can be seen in (b)

The faces of all participants who were within hearing range were displayed on the Head-Up Display (HUD). These faces were photographs of the participants (extracted from

their photographic avatars), so they could be easily recognized. When a participant spoke, the faces were highlighted. This was designed to make the participant aware that they could be heard by all the participants shown on their HUD, even if some of them were fellow group members whose avatars were a considerable distance away. When another person was talking, their face was highlighted on the HUD, with a speech icon next to it. This gave participants additional information as to who was speaking, which was particularly useful if the associated avatar was out of sight.

The teleporting functionality took place as rapid controlled movement, to help prevent disorientation associated with an instantaneous change of location (see Section 2). It utilized the same algorithm as the automatic following functionality from the original MGDs which could also be used to rapidly move to another’s location (even if the other was a moving target). Inspiration for the algorithm was taken from [MCR90], with the addition of gradual acceleration as well as deceleration. To avoid breaks in visual continuity caused by teleporting through walls, our implementation raised a participant to a birds-eye view during teleportation, so the participant could clearly see where they were being taken.

4. Experiment

The experiment used the context of urban planning, with participants asked to use a CVE to review the design of a new housing estate. Participants were run in two batches. In the first of these (the teleporting condition), participants had all the basic MGD functionality from [DR08], and new MGD functionality to provide awareness of who was talking, who was within hearing range and teleporting. Participants in the second batch were provided with multiple views (the multiple views condition), in addition to all the MGD functionality that was provided to the other batch of participants.

4.1. Method

The experiment took place in an undergraduate computing laboratory. Each participant was provided with a headset, and they were spread out across the laboratory so they could only communicate using audio and text communication from within the environment. Participants used two adjacent computers, one for the CVE and the other for the urban planning report write-up. The CVE application, environment and experimental procedure were the same as in [DR08].

4.1.1. Participants

All participants were undergraduate students from the School of Computing, who had not taken part in the previous study. Eight participants were recruited for each run, but one participant in the teleporting condition was unavailable on the day of the experiment. The remaining seven participants in the teleporting condition (6 men and 1 woman) had

a mean age of 21.7 ($SD = 5.2$). The eight participants in the multiple views condition (5 men and 3 women) had a mean age of 21.8 ($SD = 4.1$).

All the participants volunteered for the experiment, gave informed consent and were paid an honorarium for their participation.

4.1.2. CVE application

The software application and 3D sound model are described in the previous study [DR08].

Distance attenuation was turned off for communications between members of the same group. This was clarified by displaying photographs of the faces of participants who would receive any transmitted audio. These faces were displayed on the HUD, and were added and removed appropriately as participants changed their position in the environment and switched groups. In addition, an icon was placed above a participant's avatar, and by the side of their face on the HUD, when they were talking.

4.1.3. Environment

The environment was a residential estate that was based on a real estate in Leeds. An annotated map of the estate is shown in Figure 2.

All participants were represented in the environment with a photographic avatar (using four photos: front, back, left and right). Participants were given an over-the-shoulder perspective, with the option of switching to and from a bird's-eye view. An over-the-shoulder perspective meant that participants could see each other relative to their avatar, and be more aware of how others perceived them [CFS02].

4.1.4. User Interface

The participants used desktop workstations, and a two-handed control method, with one hand on the keyboard and the other hand on a 3-button mouse. By holding down appropriate arrow keys a participant could move forward/backward/left/right at 6 m/s, and heading and pitch could be changed by moving the mouse. This is a common gaming control method (e.g. [BB04]).

The 'Insert' key was used to take screenshots, the 'Home' key to toggle between over-the-shoulder and bird's-eye views, and holding down the 'Page Down' key allowed the participant to use voice communication.

Text communication was achieved by simply typing letters or numbers, which were transmitted the moment each was typed, appearing in a speech bubble above the participant's avatar. The text expired after approximately ten seconds from the moment the enter key was pressed. Each participant was provided with a stereo headset for audio communication. The default recording and playback volumes were automatically set using a shell script.

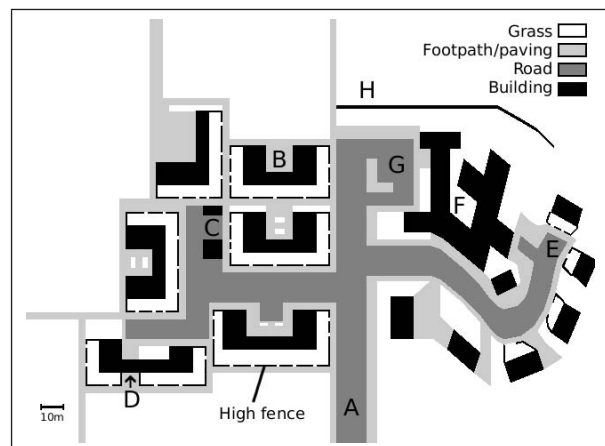


Figure 2: A map of the estate. The estate had an entrance road in the middle (point A), which acted as a dividing line between two styles of building. On the left-hand side of the entrance road, there were brown-bricked terraced houses, which were mostly horse-shoe shapes creating partially enclosed private space (e.g. point B). The front gardens were bordered by high fences, and there were six garages in the road (C). There was an archway under one of the terraces (D). On the right-hand side of the entrance road there were red-bricked bungalows (single story buildings) along the edge of the curved road, with gardens bordered by low brick walls (e.g. E). There was a single-story care home for elderly people (F), with a car park to the left with space for six cars (G), and a hedge-row above it partly separating private land around the care home from public parkland (H).

The basic MGD functionality used three mouse buttons, and the 'Delete' key to move up one level in the group hierarchy. The display had a crosshair in the middle used for selection. Selecting an avatar with the left mouse button formed/joined a group. Selecting the avatar of a fellow group member with the right mouse button rapidly moved to their location and automatically followed them. Pressing the middle mouse button anywhere moved to the mean location of the group.

Holding down the numpad zero key released the mouse from controlling heading and pitch, and allowed it to control the position of the red teleporting arrow. Once the arrow was positioned in the desired location, a left mouse click teleported the participant there.

Participants in the multiple views condition could position the teleporting arrow over one of their group members views, and clicking the left mouse button would teleport them to that group member's location. By default, the participant's subsequent movements were tethered to that group member (the automatic following functionality in basic MGDs) but the participant could 'free' themselves simply by pressing a movement key.

The multiple views took up the bottom quarter of the screen. A limit was imposed of three views, each taking up a quarter of the horizontal space, with the remaining quarter reserved for displaying the faces of any other group members. These could be selected using the numpad zero key to release the mouse pointer. Selecting them showed their view in one of the existing viewports, swapping out the member whose view had been replaced.

4.1.5. Procedure

A 10 minute meeting was held with participants a few days before the experiment. They received a verbal explanation of the experiment, a single-sided A4 sheet containing extracts from UK urban planning guidelines and a consent form. They also had photos taken for their avatar during this time.

The experiment itself lasted one hour. At the start participants were provided with another copy of the urban planning guidelines sheet, an instruction sheet for using the CVE, an experiment schedule, and an electronic copy of an urban planning report which they had to complete during the experiment. The report contained the following questions, which participants were asked to illustrate using screenshots:

- **Question 1, Permeability:** (a) How many entrance and exit points are there around the estate? What are these for (i.e. cars or pedestrians)? (b) What reduces the speed/volume of traffic? (c) Are there suitable pedestrian routes around the environment? (d) Are the blocks small enough or do you have to walk too far before you reach a choice of direction?
- **Question 2, Character:** (a) Which parts of the environment follow the same pattern/building structure? (b) Find a part of the environment that is not consistent with the layout of the estate. (c) Is this acceptable or should it be changed? (d) Does the estate have character?
- **Question 3, Safety & Security:** (a) Comment on the safety and security of the estate based on your own thoughts, the information in the guidelines and your discussion with other participants. (b) Find examples of where public and private space is clearly distinguished and where it isn't. (c) Discuss which part(s) of the estate you think are least safe. (d) Can you find any blank walls that you think should be overlooked to improve the feeling of safety and help prevent graffiti? (e) Try to suggest some improvements with regard to the safety and security of the estate.

5. Results

There were two types of work that took place in the experiment: taskwork and teamwork [BGG02]. Taskwork refers to the answers given in participants' reports, whereas data about teamwork were provided by the server's log of the movements, communication and groups that participants formed.

The urban planning reports were marked like an exam. Participants names were on the reports, marking wasn't blind. An independent samples t-test showed no significant difference between the teleporting and multiple views conditions, $t(13) = 1.49, p = .16$. Participants in the teleport condition had a mean mark of 18.7 ($SD = 3.3$) out of 24, and 16.3 ($SD = 3.1$) in the multiple views condition. Our focus, however, was on how participants went about doing the task (i.e. the teamwork), and how different MGD functionality affected participants' behavior. This was analyzed both quantitatively and qualitatively.

5.1. Quantitative Analysis

For each batch of participants, the spoken and text communication was transcribed and analyzed using a communication coding approach [BJSB98] to classify each utterance as one of the following:

- (a) **Greetings** (e.g. 'Hey M!', 'Hi G!')
- (b) **Functionality** – communication regarding the system and the groups (e.g. 'Think we need smaller groups than all of us!', 'You do realize that if you just press 'Home' you get a bird's-eye view and it's a lot easier to see!')
- (c) **Environment** – discussion about the 3D world, but not in relation to the task (e.g. 'I swear you should be able to see uni from here.', 'I kind of might have figured out where the pictures were taken of, you know the Leeds skyscrapers ones.')
- (d) **Task related** (e.g. 'Which part's the least safe?', 'I'd say where we're stood now, J.')
- (e) **Idle chat** (e.g. 'Party at my flat. Come on, let's go!')

These data were analyzed in terms of the quantity of communication that took place, and where participants were relative to each other when they communicated. For comparison, data are provided from a previous study [DR08] when other participants had performed the same urban planning task either in a conventional CVE ('control' in Figure 3) or with basic MGDs functionality (see Figures 3 and 4). Note that the average group size in the basic MGDs, teleport and multiple views conditions was 3.5, 2.5 and 3.0 respectively.

The total number of utterances made by participants in the basic MGDs (data from [DR08]), teleport and multiple views conditions (data from the present study) was analyzed using a univariate analysis of variance (ANOVA). This showed that there was a significant difference between the conditions, $F(2, 20) = 3.91, p = .04$. Tukey HSD posthoc tests showed that the difference between basic MGDs and multiple views was significant ($p = .03$) but the other pairwise comparisons were not. The mean amount of communication increased by 226% from the basic MGDs to the teleport condition, and by another 27% from the teleport to the multiple views condition. Within this, task related communication increased by a factor of two from basic MGDs to the teleport and multiple views conditions, but this was not

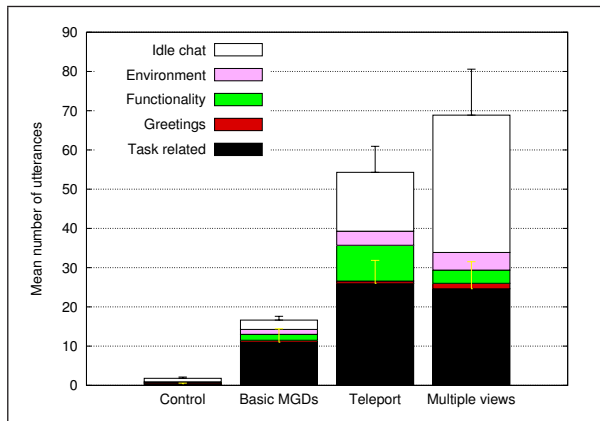


Figure 3: Mean number of utterances made by the participants in each condition. The control and basic MGDs conditions are from [DR08]. The error bars are shown for task related utterances and idle chat.

significant. Idle chat more than doubled from the teleport to the multiple views condition (see Figure 3).

One of the limitations identified in our previous research was that participants tended to assemble in one place in the CVE before communicating, even though this was unnecessary with the basic MGDs functionality that was provided (see Introduction). To determine whether the new functionality provided in the present study overcame this limitation, each time a participant made an utterance the distance to their nearest group member was calculated, and the mean for each participant in the basic MGDs, teleport and multiple views conditions was analyzed using a univariate ANOVA. The two participants who didn't speak at all during the experiment were excluded from the analysis, one was from the basic MGDs condition and the other was from the multiple views condition. The ANOVA showed that there was a significant difference between the conditions, $F(2, 18) = 3.56, p = .05$. Tukey HSD posthoc tests showed that the difference between basic MGDs and multiple views was significant ($p = .04$) but the other pairwise comparisons were not (see Figure 4).

5.2. Qualitative Analysis

The quantitative analyses show that teleporting and multiple views increased both the quantity of communication that took place and the distance over which participants communicated. The purpose of the qualitative analysis was to understand the underlying behavioral changes that cause these increases, and how teleporting was used in general.

The server log allowed the distances participants traveled while teleporting and walking to be calculated and showed that, overall, 16% of travel was by teleporting. Further investigation showed that there were two distinct uses of teleport-




Figure 4: Mean distance to the nearest group member at the time of each participant's utterances. The basic MGDs condition was from [DR08].

ing. First, teleporting was used to speed up exploration of the environment, particularly when participants first entered the environment (see Figure 5). Second, teleporting was used to reach points of interest. For example, at one point during the experiment the some participants' conversation was about blank walls, which was relevant to one of the questions in the task. The blank walls were at the ends of the horseshoe-shaped buildings. Participant *O*, represented by a green line in Figure 6, teleported across the building on the left to view the blank walls (timestamp [26:20]). *O* then teleported up to the top of the map to see the walls that *I* and *R* were talking about.




Figure 5: Path showing the first 5 minutes of movement of a participant who used teleporting to speed up their exploration of the environment. A solid line represents walking, and a dotted line represents teleportation.

In order to teleport to a point of interest, a participant must

first know its location within the environment. This is sometimes difficult, as the following conversation extract from the teleport condition shows:

[09:25] O: Look, show me, M, show me the two entry points then, the road ones at least.
 [09:29] M: Alright well, are you where I am now?
 [09:35] O: Where are you?

The multiple views condition helps with this by allowing participants to teleport to the location of a group member by clicking on their viewport:

[39:42] C: That's useful!
 [39:43] S: Yeah! Where are you? I'll show you!
 [39:49] C: I'll teleport to you! Hang on!

Each component of the new MGDs functionality (awareness of who could hear one's communication, multiple views and teleporting) that was provided in the present study had the potential to increase the distance over which participants communicated. The data indicate that multiple views made the greatest contribution (see Figure 4). To identify whether awareness or teleporting was the most important secondary cause a detailed analysis was made of the communication and movement of the two participants (*I* and *O*) who spoke the most in the most persistent group in the teleport condition.

I and *O* both spoke in 18 conversation blocks, but used teleporting in only four of these blocks. On all four of these occasions, *I* and *O* used teleporting to collocate within the environment. In the other blocks *I* and *O* either were together (5 blocks), remained separated (2 blocks), separated without teleporting (3 blocks) or collocated without teleporting (4 blocks). This suggests that the awareness functionality was more important than teleporting for increasing the distance over which participants communicated.

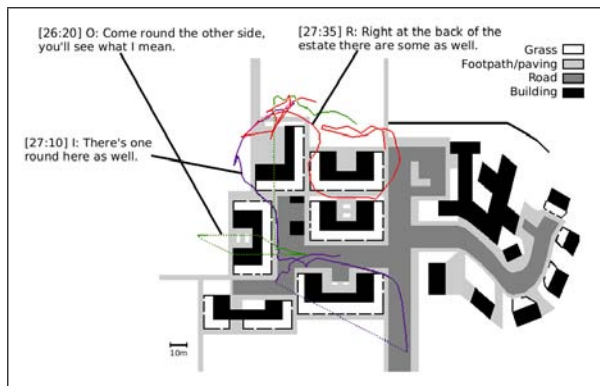


Figure 6: Paths taken by participants *O* (green line), *I* (purple line) and *R* (red line) when talking about blank walls (a point of interest). The solid lines representing walking, and dotted lines represent teleporting (participant *O* teleports to point of interests at the ends of the horseshoe-shaped buildings). To place participants' movement in context, the paths are labeled with timestamps and conversation utterances.

6. Discussion

We identified problems in conventional CVEs of finding other participants in relation to points of interest and understanding their perspective (see problems 1 & 2 in Section 2). The group graph metaphor could help with finding others, since the graph 'tracked' participants and the nodes corresponded to avatars, with edges denoting group membership (see Figure 1). One could find a group member by following a line from their avatar until they reached a node. However, the teleporting and multiple views functionality took this a step further. It allowed participants to teleport directly to a group member of their choice by selecting the appropriate viewport, and the qualitative data gave an example of how this helped participants (see the conversation extract in Section 5.2). Furthermore, providing participants with multiple views specifically tackled the problem of understanding another's perspective.

The qualitative data showed that teleporting was used in two ways: increasing the speed of movement (in particular, an initial speed search) and movement to points of interest. This is a simple method of time-saving functionality, however potential drawbacks must not be overlooked. Firstly, as with any new functionality, a potential problem might be making the system over complicated. Other features may be forgotten about and may not be used to their full potential. Secondly, and perhaps more subtly, teleporting could mean people lose the feel for distance. In an urban planning context, it is important that participants in the CVE get a feel for the scale of the environment, in particular the size of buildings and proposed developments. One of the questions for the urban planning report was asking participants if they thought the blocks of houses were the right size. Financial savings are made by building the houses joined together in blocks, but a large block size decreases permeability of the estate, making it bad for transport and pedestrians (they have to go further before they can change direction). Teleportation may mean participants lose a sense of scale and large blocks could go unnoticed.

One of the places where the original MGDs techniques fell short of their goals was in facilitating communication when participants were spatially separated within the environment. The fact that participants tended to collocate to communicate in the basic MGDs condition was a sign of inefficient groupwork – participants were either taking time to collocate when they wanted to communicate, or they were waiting until they were coincidentally collocated before they said anything. Providing functionality to communicate with group members from a distance, and informing the participants of this in the instructions, was not enough to make their behavior more efficient. This study indicates that by providing feedback to the participants, they became more aware of how the system works. The quantitative data showed that in both the teleporting and multiple views conditions, participants communicated across greater distances than in the ba-

sic MGDs condition (see Section 5.1). The interesting thing about this feedback from the system is it's not specifically new functionality in the sense of a new tool at the users' disposal, like teleporting and multiple views are. Instead it provides awareness of *existing* functionality: the ability to communicate with group members from a distance. As Schroeder et al. reflect, do we improve usability 'by means of improving the systems and features of the environment, or by improving the users' awareness of their activities and settings?' [SHT06] (p. 666).

Finally, in previous research, participants communicated a great deal to overcome the lack of sensory information that CVEs provided [HFH*00] [RSJ02]. However, in the present study participants communicated much more than in conventional CVEs because they were provided with more sensory information (e.g. awareness of who could hear you and who was speaking, and multiple views providing an 'extra pair of eyes'). The quantitative data showed that the amount of conversation increased 4 times from the basic MGDs condition to the multiple views condition. This increase in communication was indicative of more teamwork taking place.

Acknowledgments

The first author is funded by a Doctoral Training Grant Studentship from the School of Computing, University of Leeds.

References

- [BB04] BROWN B., BELL M.: CSCW at play: 'There' as a collaborative virtual environment. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)* (Chicago, Illinois, USA, November 2004), ACM, pp. 350–359.
- [BGG02] BAKER K., GREENBERG S., GUTWIN C.: Empirical development of a heuristic evaluation methodology for shared workspace groupware. In *Proceedings of the 2002 ACM conference on Computer Supported Cooperative Work (CSCW '02)* (New Orleans, Louisiana, USA, November 16–20 2002), ACM, pp. 96–105.
- [BJSB98] BOWERS C. A., JENTSCH F., SALAS E., BRAUN C. C.: Analyzing communication sequences for team training needs assessment. *Human Factors* 40, 4 (1998), 672–679.
- [CFS02] CHENG L., FARNHAM S., STONE L.: Lessons learned: Building and deploying shared virtual environments. In *The Social Life of Avatars: Presence and Interaction in Shared Virtual Environments*, Schroeder R., (Ed.). Springer, London, 2002, pp. 90–111.
- [DR08] DODDS T. J., RUDDLE R. A.: Mobile group dynamics in large-scale collaborative virtual environments. In *Proceedings of the IEEE Virtual Reality (VR '08)* (Reno, Nevada, USA, March 8–12 2008). In press.
- [ENSK01] ELVINS T. T., NADEAU D. R., SCHUL R., KIRSH D.: Worldlets: 3-D thumbnails for wayfinding in large virtual worlds. *Presence: Teleoperators and Virtual Environments* 10, 6 (2001), 565–582.
- [FGV*00] FRASER M., GLOVER T., VAGHI I., BENFORD S., GREENHALGH C.: Revealing the realities of collaborative virtual reality. In *Proceedings of the 3rd International Conference on Collaborative Virtual Environments (CVE '00)* (San Francisco, California, September 2000), ACM, pp. 29–37.
- [HFH*00] HINDMARSH J., FRASER M., HEATH C., BENFORD S., GREENHALGH C.: Object-focused interaction in collaborative virtual environments. *ACM Transactions on Computer-Human Interaction* 7, 4 (2000), 477–509.
- [MCR90] MACKINLAY J. D., CARD S. K., ROBERTSON G. G.: Rapid controlled movement through a virtual 3D workspace. *Computer Graphics* 24, 4 (1990), 171–176.
- [Pek02] PEKKOLA S.: Critical approach to 3D virtual realities for group work. *Proceedings of the Second Nordic Conference on Human-Computer Interaction (NordCHI '02)* (October 2002), 129–138.
- [RHPJ00] RUDDLE R. A., HOWES A., PAYNE S. J., JONES D. M.: The effects of hyperlinks on navigation in virtual environments. *International Journal of Human-Computer Studies* 53, 4 (2000), 551–581.
- [RSJ02] RUDDLE R. A., SAVAGE J. C. D., JONES D. M.: Verbal communication during cooperative object manipulation. In *Proceedings of the 4th International Conference on Collaborative Virtual Environments (CVE '02)* (Bonn, Germany, September 30–October 2 2002), ACM, pp. 120–127.
- [SHT06] SCHROEDER R., HELDAL I., TROMP J.: The usability of collaborative virtual environments and methods for the analysis of interaction. *Presence: Teleoperators and Virtual Environments* 15, 6 (2006), 655–667.
- [SWM03] SONNENWALD D. H., WHITTON M. C., MAGLAUGHLIN K. L.: Evaluating a scientific laboratory: Results of a controlled experiment. *ACM Transactions on Computer-Human Interaction* 10, 2 (2003), 150–176.
- [WSWL02] WÖSSNER U., SCHULZE J. P., WALZ S. P., LANG U.: Evaluation of a collaborative volume rendering application in a distributed virtual environment. In *Proceedings of the workshop on Virtual Environments 2002 (EGVE '02)* (Barcelona, Spain, May 30–31 2002), The Eurographics Association, pp. 113–221.
- [YO02] YANG H., OLSON G. M.: Exploring collaborative navigation: the effect of perspectives on group performance. In *Proceedings of the 4th International Conference on Collaborative Virtual Environments (CVE '02)* (Bonn, Germany, September 30–October 2 2002), ACM, pp. 135–142.