

Evaluation of Collaborative Construction in Mixed Reality

Breght R. Boschker and Jurriaan D. Mulder

Centre for Mathematics and Computer Science CWI, Amsterdam, The Netherlands

Abstract

Collaborative virtual and augmented reality are an active area of research and many systems supporting collaboration have been presented. Just like there are many different systems for VR and AR, there are many different types of collaboration. In some cases, virtual reality is used to enhance an existing collaborative process. In other cases, it enables new types of collaboration that previously were not possible (e.g. distributed VR). Other systems support tasks that can be performed either individually as well as collaboratively. While these tasks may allow to be performed collaboratively, little has been said on what the added benefit is in doing so. We present a user study of a collaborative construction task in a shared physical workspace virtual reality environment under various degrees of interaction in collaboration. Our results show that, for this type of task, a pair of subjects concurrently interacting can be significantly more effective, even though individual user performance decreases. Our results further show that there is no significant benefit in giving only verbal and non-verbal assistance over a single user performing the task.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.6.3 [Simulation and Modeling]: Applications

1. Introduction

Collaborative augmented and virtual reality (CVR) is an active area of research which has resulted in a number of systems supporting collaboration in VR or AR over the years.

The CVR-systems developed over the years by the research community differ in size, functionality and (display) technology. Systems range from large and immersive environments such as the CAVE to small, hand-held devices such as used in the Invisible Train [WPLS05].

Next to differences in technology for CVR systems, collaboration in itself is not a rigid concept with fixed parameters. For example, the number of participants, the role they play, the degree of interaction each can perform, and what communication channels are available all play a role.

One use of collaborative AR/VR is to enhance existing collaborative processes. For example, people from different disciplines can combine their knowledge in new ways, e.g. a radiologist and a surgeon both having an image specific to their discipline overlaid on a patient. Additionally, CVR enables new types of interaction between users that were previously not possible. For instance, users at different locations can employ virtual reality to concurrently work together over

(large) distances. This is referred to as *distributed* collaborative VR.

In the aforementioned situations, the benefits of collaborative AR/VR are clear. However, many tasks can be performed both individually, as well as collaboratively. In these cases the question arises what the benefit is of performing such a task collaboratively.

In this paper, we present a study on collaboration in a desktop VR/AR system with a shared physical workspace. We address face-to-face collaboration through an experiment where people either work individually, assist each other by communicating, or actively participate in the task at hand. This task comprised a simple 3D assembly application in which the subjects had to construct a virtual pipe out of single elements. The goal of this experiment was to assess the effectiveness of the various degrees of participation in collaboration. We expect that, with the increase of participation, user and task efficiency increase.

This paper is organized as follows: in Section 2, we address related research on collaborative virtual and augmented reality with a shared workspace as well as related research on collaboration. Then, Section 3 briefly describes

our CVR system, the Joint Space Station (JSS, [MB04]), the construction application, and the types of collaboration in the JSS. In Section 4, we describe the experiment, and in Section 5 we discuss the results of our experiment. Finally, Section 6 gives the conclusions and directions for future research.

2. Related Work

Over the years, the research community has produced a number of virtual environments. In this work, we focus on collaborative virtual and augmented reality with a shared physical workspace, a number of which we will address in Section 2.1. In Section 2.2, we address related work in the area of research on collaboration in VR and AR.

2.1. Systems for AR/VR with a Shared Physical Workspace

Systems for augmented and/or virtual reality that allow users to collaborate in the same physical environment come in different flavors. Based on display type, these systems can be grouped either under stationary display systems, or non-stationary display systems.

One of the benefits of stationary display systems is their high resolution (compared to most non-stationary display systems). Various approaches have been taken to be able to share a single display by multiple users. One approach taken to share one display with multiple users, is to partition the screen, so that each user can see a part of the screen. This approach is taken by Kitamura *et al.* [KKYK01].

A more common approach to share one display between multiple users is to interleave the (stereoscopic) images each user sees. Examples of such so-called time-sequential systems are the multi-user CAVE [CNSD*92], the Two-User Responsive Workbench [ABM*97], the Virtual Showcase [BFSc01].

Finally, there is a group of stationary single-display systems based on an array of (small) lenses placed in front of a display. Such *lenticular* displays refract light from the display for each eye so that stereoscopic images can be formed. One example of a collaborative lenticular display is the one made by QinetiQ [QHPVL04].

One system that takes a different approach, is to provide each user with his or her own stationary display. The Joint Space Station [MB04], in use at our institute, is such a system. Another example of such a system is the PIT [APT*98], in which each user has his or her own back-projected screen. A more recent approach is the system described by Hirakawa *et al.* [HK04], comprising of a large (top-/front) projection screen, in front of which users can stand and interact with system and other users by gesturing and props.

Another common approach is to provide each user with

his or her own non-stationary display, such as for example, a Head Mounted Display (HMD). The benefits of these systems is that users can be mobile and therefore co-location of visual and interaction space is easy, but on the other hand, there are issues of ergonomics and resolution HMD displays (although resolution is rapidly increasing). Examples of collaborative systems based on HMDs is the StudierStube Augmented Reality Project [FP01], the SeamlessDesign system [KTY99], and the Virtual Round Table [BMS00]. An approach that combines head mounted displays, projection-based graphics, PDAs and CRT/LCD-screens (e.g. on a laptop) is MagicMeeting [RWB02].

2.2. Research on Collaboration in VR and AR

As early as 1991, Tang [Tan91] researched work practices during collaborative design on a table. Tang observed the importance of hand gestures as significant resource for communicating information. Although Tang's work is focused on a groupware drawing tool, he mentions the importance of sharing a view on the task at hand, as well as the need for a sense of awareness of other participants for effective collaboration.

In the area of collaborative augmented reality, Billinghurst *et al.* published their findings from experiments with face-to-face collaboration in co-located AR environments [BBGK02, BKK*02]. In their work, Billinghurst *et al.* found users in a collaborative AR environment exhibited many of the same behaviors that are present in unmediated face-to-face collaboration. They conclude that interaction with content in AR should be through the tangible interfaces metaphor. They further conclude that a mediated view on the real world affects a user's ability to perceive their collaborator and his/her non-verbal cues. As a result, speech patterns may change to compensate for the limited visual perception. They advise to use employ a display system that provides an unmediated view on the real world for collaborative AR.

Recently, Scott *et al.* [SSCI04] have described issues of *territoriality* in collaborative tabletop workspaces. For this, they performed two observational studies in which subjects play various tabletop games. Their results show that subjects use different areas on the table ('territories') to coordinate their interaction. In their work, they give arguments for each user to have an individual workspace in addition to the workspace that is shared by the collaborators.

In their work on Distributed Collaborative Virtual Environments (DCVEs), Otto and Roberts [OR03] stress the importance real world communication in closely coupled collaboration. The results from a user study they performed, show that when subjects were aware of the communication methods they had (i.e. a microphone), usage of those methods increased. Furthermore, they reported that this awareness resulted in the team working together more successfully.





Figure 1: Two users working in prototype JSS (in duo-configuration) and schematic view of set-up.

Also based on research in a distributed system, Fussell [FKS00] describes the effect of sharing a visual context in coordination of communication and task performance. In their work, they show the added benefit of sharing visual context but raise questions on the adequacy of video communication technology.

In their work, Gutwin and Greenberg [GG00] describe the mechanics of collaboration in a Computer Supported Collaborative Work (CSCW) setting. They describe low level actions and interactions that must be carried out to successfully perform a task in a shared manner. Although not strictly VR-related, their findings that communication and coordination, and monitoring and assistance are important for successful completion of collaborative tasks, are nonetheless interesting for our work.

3. The Joint Space Station (JSS)

The Joint Space Station (JSS, see Figure 1) is a tool designed for collaborative desktop virtual and augmented reality with a shared physical workspace. In Section 3.1, we describe the architecture of the JSS, followed by a description of the construction application in Section 3.2. Finally, in Section 3.3, we address the types of collaboration that can occur in the JSS.

3.1. Architecture

The JSS (see Figure 1) is constructed out of multiple modular desktop VR/AR stations (called Personal Space Stations or PSSs [MvL02]). The PSS is a mirror-based system equipped with a standard CRT monitor. The user looks at the monitor via a (half or fully silvered) mirror. Because the reflected image of the CRT is perceived behind the mirror, visual and interaction space are co-located. Interaction with the system is performed by reaching under the mirror and manipulating physical objects articulated with retro-reflective fiducials. The objects are tracked using a custom optical tracking system. This system consists of two camera's equipped with a ring of infra-red (IR) LEDs and an




Figure 2: Screenshot of example run of construction application. In this run, a number of pieces of pipe have been placed by a user, and the water has begun to flow (blue piece of pipe). On the right of the screen (two pieces of pipe, coupling on top, elbow below), the part selector is shown. In the foreground, a currently selected piece (red wireframe) that still needs to be attached to pipe is shown.

IR-pass filter in front of the lens. Additionally, each PSS has three pedals which double as mouse buttons.

By cascading PSSs, the individual workspaces of each PSS coincide and form one shared physical workspace. Users can see one another, thus allowing natural communication as normally present in face-to-face collaboration.

3.2. Application: *Pipe Dream 3D*

The construction application is a collaborative game called *Pipe Dream 3D* (see Figure 2) based on the classic arcade game 'Pipe Mania'. The goal of this game is to build a pipe. After an initial delay, water starts flowing through the pipe. Once the water reaches the end of the pipe, the game is over.

Our three-dimensional implementation is based on a reg-




Figure 3: Invalid moves. Left: an invalid move is indicated by a red wire frame cube around a candidate. Right: trying to build a pipe into an obstacle (opaque cube) is disallowed.

ular volume of cells in which users can place pieces of pipe. New pieces need to be connected to the end of the pipe already constructed. A starting point is provided by a piece sitting on the lowest level of the volume. Additionally, the volume contains obstacles that the user is to avoid (shown as opaque red cubes in Figure 2). To aid users in interaction, the size of the volume is shown by a grid on one side and the bottom of the volume. Furthermore, the current position of the input device is highlighted on the grid.

At the opposite side of the volume, two pieces of pipe are depicted from which the user can select the piece to play next. In the game, there are two types of pipe: a straight piece ('coupling') and a corner piece ('elbow'). To select a piece, the user holds the input device at the position of one of these pieces and presses the pedal. The selected piece is attached to the input device such that the user can move and orient it. To connect a piece to the end of the pipe, the user positions the selected piece at the correct position and presses the pedal to release the piece. When released, the pose is 'snapped' to a discretized 90 degree orientation. A selected piece can only be released when held in a valid pose; if a chosen pose is invalid, this is indicated by a red wire-frame box (Figure 3). A pose is invalid when it does not connect to the existing pipe, if it would cause the pipe to run into an obstacle (Figure 3, right), or would cause the pipe to leave the volume.

3.3. Collaboration

In collaboration, we distinguish two aspects: inter-user communication, and participation in manipulation. Communication in a natural, face-to-face setting does not solely consist of speech: people gesture, use body language, and use aids such as a piece of paper and a pen when explaining etc. Research has shown the importance of such non-verbal communication methods (e.g. [GG00, FKS02]). Systems supporting collaboration should therefore maintain or support the natural communication channels as much as possible.

When performing a task collaboratively, the degree of participation of each user in the actual manipulation can vary. For example, users can take the role of passive spectator, as-

sist verbally and non-verbally, or actively participate in manipulation in the task at hand.

The JSS is a system designed with these aspects in mind; its physical construction allows users to see each other. If a half-silvered mirror is used the hands of other users can still be seen next to the virtual content. Furthermore, the design of the JSS does not impose rules on users concerning the degree of participation in manipulation: users can either have their own manipulation devices, share them or be passive observers.

In this study, we focus on one aspect of collaboration: the degree of participation in a collaborative construction task. We studied four configurations:

SU: Single-user task — as a reference, we include a single-user condition. In this condition, only the currently 'active' user has virtual content displayed.

SUA: Single-user interaction task with assistance — in this condition, both subjects have virtual content, but only *one* of the subjects may perform interaction. The other subject was instructed to assist by giving hints (e.g. gesturing or speech).

C1D: Collaborative interaction, 1 device — this condition differs from the former in that both subjects were allowed to interact with the virtual content. However, there was only one interaction device. Therefore, subjects had to pass on the device to allow the other subject to interact.

C2D: Collaborative interaction, 2 devices — a 'fully collaborative' condition, in which each subject was given his or her own interaction device, i.e. both subjects could concurrently interact.

We hypothesized that, with increasing degree of participation, user and task efficiency would increase. As the main performance metric, we used the number pieces of pipe a subject could place per second. Furthermore, we looked at total pipe lengths as a measure for task performance, and the ratio of the types pipe pieces. The latter can be seen as an indication of pipe complexity, since straight pipe pieces are easier to place than elbow pieces.

4. Experiment Description

In the following Sections, we detail the experiment parameters, subject population and measurements taken during the experiment.

4.1. Test System Details

In our experiment, we used the dual-system configuration of the JSS. The system employs two half-silvered surface-reflecting mirrors, thus creating an augmented reality environment. The total size of the shared workspace is about 60cm wide, 50cm deep and 50cm high.

Each PSS was equipped with a 22" CRT monitor running

at 120Hz in stereo. Head tracking was provided using a custom optical tracking system described in [MJvR03]. Interaction tracking is done by a single set of progressive-scan camera's running at 60Hz and a dedicated PC equipped with frame grabber boards by Leutron Vision. This system broadcasts the tracker data over a standard 100Mbit Ethernet link to both other PCs.

4.2. Application Details

In the experiment, subjects were asked to construct a 3D virtual pipe using the *Pipe Dream 3D* application introduced in Section 3.2. Figure 2 shows the application in action and Figure 4 shows the view for the first user.

The two water flow parameters, initial delay and speed of the running water were tuned experimentally. The initial delay (the time between connecting of the first piece to the start piece and the start of the water flow) was set to 30 seconds. The the water flowed at a speed of 1.5 seconds per piece of pipe. When the water flow had reached the end of the pipe, subjects had one period (i.e. 1.5 seconds) of 'slack time', in which a new piece could be connected. If this period had expired and no new piece had been connected, the run was over.

The construction volume was centered between both subjects, and placed about 15 centimeters above the table top. A physical helper table was provided to serve as a frame of reference. Serving as a black background, the helper table also aided users in fusing the stereoscopic images. The size of the volume was chosen to be 10 cells wide (left-right), 7 cells deep (seen from the user) and 5 cells high. Each cell was a cube of 3x3x3cm, bringing the volume size at 30x15x21cm.

The grid was filled with a total of 18 virtual obstacles randomly distributed throughout the volume (see Figure 2). The obstacles were checked for impossible situations (e.g. an obstacle directly in front of the start piece) and distribution. To prevent learning effects between runs, 8 different configurations of obstacles were used.

Because an elbow can be positioned in more poses (eight poses) than a coupling (which has only three unique poses), it costs more time to get an elbow in a correct pose. To make sure this discrepancy in difficulty would not affect the choice of type of piece of pipe too much, each part selected was given an initial (discretized on 90-degree angles) random orientation.

4.3. Task Description and Conditions

Subjects were asked to construct a pipe that was as long as they could, until the water reached the end of the pipe.

Before (a pair of) subjects started the actual experiment, they were given time to familiarize themselves with the system and the application. During this test run, the workings




Figure 4: View for first user (application content has been overlaid for clarity). In this figure, the first subject is holding the input device with a virtual elbow attached. The user can see his own hands as well as the other user's hands and actions.

of the optical tracking system, the head tracking system and the foot pedal were explained. Furthermore, the application and task were explained. Just like in the real experiment, the test application contained obstacles. In this first test run, the water flow was disabled.

In the instruction phase, the meaning of the objects in the workspace was explained, as well as the size of the volume and what types of moves were allowed. Users were instructed on the 'snapping' of a selected piece of pipe and that they therefore did not have to be too precise in aligning the piece with the pipe: as soon as the red wireframe would disappear, a pose was valid and they could attach it. Finally, subjects were asked to gesture and speak as much as they liked, both between one another and with respect to the experimenter. When subjects felt comfortable with the system and application, a second test run was started, this time with the water flow enabled. When subjects familiarized themselves with the application and the concept, the experiment was started.

Each of the conditions in Section 3.3 was performed three times in a row. To gain insight in single user performance, three runs of the single user (SU) and single user with assistance (SUA) conditions were performed by both subjects. Both collaborative conditions were performed three times per pair. To prevent influence of learning effects, the order of conditions was permuted (the four sets single user conditions were counted as two conditions, bringing the total to six 'conditions').

	mean	std	min	max
SU	27.75	13.95	2	60
SUA	28.18	16.64	5	74
C1D	24.39	13.50	3	62
C2D	53.61	24.43	19	138

Table 1: Pipe length (in number of pieces of pipe placed by a subject)

Between different conditions, subjects were given time to relax their arms, shoulders and neck, as well as relax their eyes. After this period, subjects were instructed on the next condition.

4.4. Subjects

A total of 24 subjects participated, ranging in age from 24 to 54 with an average of 30 years old. All subjects were either scientific or supporting staff. 12 of the subjects had participated in prior user tests with the PSS and thus were more or less familiar with the system. In our subject population, there were no left-handed people. In total, there were 5 female subjects. 19 subjects indicated that they wear glasses or lenses.

4.5. Measurements

Throughout the experiment, we recorded head tracker and interaction tracker data as well as place and select actions, foot pedal events and other application events (e.g. time steps in the water flow) for analysis or play-back of the session at a later. Furthermore, each session was registered on video using a standard digital video camera.

5. Results and Analysis

To find significant differences in the measurements, we performed an analysis of variance (ANOVA, $\alpha = 0.01$) on the data recorded. The first measure we looked at, were the lengths of the pipes constructed. In Table 1, we have included statistics on pipe lengths for all four conditions. These results are displayed graphically in Figure 5.

Furthermore, we subjected our main metric for individual user performance, the number of pieces placed per second, to an analysis of variance. The results of this test are given in Table 2 and shown graphically in Figure 6.

Table 1 shows a significant increase in average pipe lengths in the fully collaborative C2D condition; pipe length was on average nearly twice as large as in the other three conditions. Interestingly, pipe lengths under the other three conditions are not significantly different.

For our main metric, the results are similar: Table 2 shows no significant difference in performance between the first




Figure 5: Box-Whisker plot for pipe lengths under all four conditions. See also Table 1.

	SU	SUA	C1D	C2D	mean
SU	—	0.344	0.179	0.000	0.338
SUA	0.344	—	0.598	0.000	0.323
C1D	0.179	0.598	—	0.000	0.313
C2D	0.000	0.000	0.000	—	0.434

Table 2: Results of ANOVA (p -values) for pipe pieces placed per second (pps). Two groups of samples differ significantly if their p -value is < 0.01 . The rightmost column denotes the average number of parts per second for each condition.

three conditions, and a significant difference between all first three conditions and the last (fully collaborative) condition. The number of pieces placed per second is significantly higher in the C2D condition.

Interesting is that, although the number of pieces of pipe is *higher* under this condition, this number is combined for two users. This means that per user, the average number of pieces of pipe placed per second is only half as large. From this, it can be concluded that, although the task performance *increases*, per-user performance *decreases*.

There are a number of possible explanations for the decreased per-user performance. For one, it could be that, because of the on average greater pipe length, users were hindered more by the obstacles. If such should be the case, the number of elbows used would increase. To make a valid

	mean	std	min	max
SU	0.3939	0.1948	0.0000	1.5000
SUA	0.4225	0.1861	0.1667	1.0909
C1D	0.4512	0.1806	0.2222	0.9091
C2D	0.4633	0.1662	0.2927	1.1000

Table 3: Elbow-coupling ratio




Figure 6: Box-whisker plot of the average number of pieces placed per second for all four conditions.

analysis between conditions, we define the ratio between the number of elbows and the number of couplings used in a pipe as a measure for pipe complexity. In Table 3, we list the mean ratios for all four conditions. An analysis of variance on the ratio shows no significant difference in these numbers. Therefore, we have to conclude increased pipe complexity is not the cause of decreased user performance.

A different reason for the decreased user performance might be found in the actions performed by one subject influence the actions of the other. For example, subjects having to wait for the other user to complete their action, or subjects taking more time to anticipate the other user's actions. To analyze these effects, we defined two more metrics:

Average placing time — we defined *placing time* as the time between the selection of a piece of pipe from the part selector and the corresponding connecting of that piece to the end of the existing pipe. Thus, for each run, for a pipe of length N , there is a vector of 'placing time' of N samples.

Average linger time — *linger time* complements *placing time*. It is defined as the time between the connecting of a piece to the end of the pipe and the selection of a new piece in the part selector. For a pipe of length N , there is a vector of 'linger time' of at most $N - 1$ samples.

An analysis of the variance in *placing time* (Table 4) and *linger time* (Table 5) showed that in the fully collaborative case C2D both significantly differed from the other three conditions: in either case, the time average time subjects needed increased. From this, we can conclude that decreased user performance was not solely caused by waiting for the other user to finish placing their piece before deciding which type of piece to play next. In that case, linger time would significantly increase, while placing time would not significantly increase.

	mean	std	min	max
SU	1.9546	0.9939	1.1337	6.4697
SUA	2.1263	1.1027	1.1230	7.2214
C1D	2.1407	1.1372	1.1983	6.4879
C2D	2.5286	0.6194	1.8439	4.4956

Table 4: Average placing time for all four conditions

	mean	std	min	max
SU	1.5632	1.4460	0.7957	11.3962
SUA	1.5349	0.8593	0.8043	5.0381
C1D	1.6475	1.1338	0.9236	7.4406
C2D	2.1892	0.4443	1.4889	3.0717

Table 5: Average linger time for all four conditions

Similarly, the decrease in user performance is not solely caused by a user having to wait before a piece can be attached, because the other user is currently attaching a piece. In that case, only placing time would significantly increase. The increased linger time and placing time indicate probably both is the case; the process of users anticipating on what the other user is doing caused the increased times. However, because these the actions of both users are interleaved, the resulting pipe can, on average, be longer.

5.1. Observations and Subject Remarks

During the instruction phase of the experiments, subjects were asked to make any remarks with respect to the test whenever they liked. A summary of observations made during the experiment and relevant remarks made by subjects follows.

Even though subjects were allowed to pass the input device in the collaborative, 1 device condition C1D, the number of device passings was, on average, low.

One interesting observation is that in the fully collaborative condition, a lot more verbal communication was used than in the other two (semi-) collaborative conditions. Even in the single user with assistance condition SUA, there was a relatively low number of utterances. Some subjects gestured heavily during the SUA condition, trying to mimic a desired pose with their hands.

Some subjects reported that during the test, when they were working at full speed, they sometimes would not notice when they had 'missed' the part selector and therefore did not have a piece of pipe attached to their input device until they tried to attach the piece to the pipe. One subject commented that this could be because he found the part selector too small. This particular subject was myopic and had trouble estimating depth in the application and commented on the projection of the 'cursor' position on the grid was of great help to him.

A few of subjects remarked during the test that building near the part selector was easier for them, because of the smaller distance they needed to travel with the input device. Some subjects remarked that building near the part selector was their new 'strategy'. One pair of subjects adopted a strategy of analyzing the new field (with obstacles) when first shown and making a plan before attaching the first piece of pipe. In some cases, this discussion phase took a few minutes (this subject pair built on average the longest pipes in the C2D condition).

At the end of all runs, subjects were interviewed shortly on how they had experienced the test and if they had any final remarks. The larger part of the subjects commented on the foot pedal as being strenuous and that they rather had used a hand-operated button (e.g. keyboard) for picking and dropping, had such a button been available. Most users commented on the experiment as being fun.

6. Conclusions and Future Work

The results from the user study we presented show that there is a significant added benefit in a user that can actively participate in a construction task. On the other hand, our results show that there is no significant benefit over a single user performing the task, when the other user is only allowed to speak, gesture, or incidentally help.

Furthermore, our results show that while a pair of subjects concurrently constructing a pipe performs significantly better than a single user that is 'aided', individual performance decreases. The cause of the decreased user performance in the fully collaborative case is likely caused by users basing their actions on what the other user is doing. To further identify the source and nature of these interactions and their implications, more research is needed.

Although the task we employed in our study has a number of common VR actions and interactions (picking, dragging etc.), in future work, we would like to perform a study that assesses effectiveness of collaboration under different tasks. One example of such a task could be a more knowledge-intensive task, such as collaboratively solving a puzzle. We expect to see more influence of verbal and non-verbal assistance in such a task.

An interesting alternative to using AR for a future study, is to use a head tracked VR condition. In this case, only verbal communication (at least non gesture-based communication) can be used in addition to the visual feedback provided by the system.

References

- [ABM*97] AGRAWALA M., BEERS A. C., MCDOWALL I., FRÖHLICH B., BOLAS M., HANRAHAN P.: The two-user responsive workbench: support for collaboration through individual views of a shared space. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 327–332.
- [APT*98] ARTHUR K., PRESTON T., TAYLOR R., BROOKS F., WHITTON M., WRIGHT W.: *Designing and Building the PIT: a Head-Tracked Stereo Workspace for Two Users*. Tech. rep., Department of Computer Science, University of North Carolina - Chapel Hill, ftp://ftp.cs.unc.edu/pub/publications/techreports/98-015.pdf, April 1998.
- [BBGK02] BILLINGHURST M., BELCHER D., GUPTA A., KIYOKAWA K.: Communication behaviors in co-located collaborative AR interfaces. In *International Journal of Human Computer Interaction (IJHCI)* (2002), Lawrence Erlbaum Associates.
- [BFS01] BIMBER O., FRÖHLICH B., SCHMALSTIEG D., CAO L. M. E.: The virtual showcase. *IEEE Computer Graphics and Applications Virtual Reality November/December 2001* (2001), 48–55.
- [BKK*02] BILLINGHURST M., KATO H., KIYOKAWA K., BELCHER D., POUPYREV I.: Experiments with face-to-face collaborative AR interfaces. In *Virtual Reality Journal* (2002).
- [BMS00] BROLL W., MEIER E., SCHARDT T.: The virtual round table- a collaborative augmented multi-user environment. In *Proceedings of the third international conference on Collaborative virtual environments* (2000), ACM Press, pp. 39–45.
- [CNSD*92] CRUZ-NEIRA C., SANDIN D. J., DEFANTI T. A., KENYON R. V., HART J. C.: The cave: audio visual experience automatic virtual environment. *Communications of the ACM* 35, 6 (1992), 64–72.
- [FKS00] FUSSELL S. R., KRAUT R. E., SIEGEL J.: Coordination of communication: effects of shared visual context on collaborative work. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work* (2000), ACM Press, pp. 21–30.
- [FKSB02] FUSSELL S. R., KRAUT R. E., SIEGEL J., BRENNAN S. E.: Relationships among speech, vision, and action in collaborative physical tasks. In *CHI '02 extended abstracts on Human factors in computing systems* (2002), ACM Press, pp. 916–917.
- [FP01] FUHRMANN A., PURGATHOFER W.: *Studierstube: An Application Environment for Multi-User Games in Virtual Reality*. Tech. Rep. 011, VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH, 2001.
- [GG00] GUTWIN C., GREENBERG S.: The mechanics of collaboration: Developing low cost usability evaluation methods for shared workspaces. In *IEEE WETICE 2000* (2000).
- [HK04] HIRAKAWA M., KOIKE S.: A collaborative augmented reality system using transparent display. In *ISMSE '04: Proceedings of the IEEE Sixth International Symposium on Multimedia Software Engineering (ISMSE'04)* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 410–416.
- [KKYK01] KITAMURA Y., KONISHI T., YAMAMOTO S., KISHINO F.: Interactive stereoscopic display for three or more users. In *ACM Siggraph 2001* (August 2001).
- [KTY99] KIYOKAWA K., TAKEMURA H., YOKOYA N.: Seamlessdesign: A face-to-face collaborative virtual/augmented environment for rapid prototyping of geometrically constrained 3-d objects. In *ICMCS* (1999), pp. 447–453.

- [MB04] MULDER J. D., BOSCHKER B. R.: Collaborative desktop virtual and augmented reality with a shared workspace. In *Proceedings of IEEE VR'04* (2004).
- [MJvR03] MULDER J. D., JANSEN J., VAN RHIJN A.: An affordable optical head tracking system for desktop VR/AR systems. In *Proceedings of the workshop on Virtual environments 2003* (2003), ACM Press, pp. 215–223.
- [MvL02] MULDER J. D., VAN LIERE R.: The personal space station: Bringing interaction within reach. In *Proceedings of VRIC 2002; 4th Virtual Reality International Conference* (2002), pp. 73–81.
- [OR03] OTTO O., ROBERTS D.: Importance of communication influences on a highly collaborative task. In *Proceedings of the Seventh IEEE International Symposium on Distributed Simulation and Real Time Applications, DS-RT 2003* (Delft, The Netherlands, 2003), pp. 195–201.
- [QHPVL04] QINETIQ HIGH PERFORMANCE VISUALIZATION LTD. HAMPSHIRE U.: Collaborative autostereo 3d display system. <http://www.qinetiq.com/hpv>, 2004.
- [RWB02] REGENBRECHT H., WAGNER M., BARATOFF G.: Magicmeeting: A collaborative tangible augmented reality system. *Virtual Reality* 6, 3 (2002), 151–166.
- [SSCI04] SCOTT S. D., SHEELAGH M., CARPENDALE T., INKPEN K. M.: Territoriality in collaborative tabletop workspaces. In *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work* (New York, NY, USA, 2004), ACM Press, pp. 294–303.
- [Tan91] TANG J.: Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies* 34, 2 (1991), 143–160.
- [WPLS05] WAGNER D., PINTARIC T., LEDERMANN F., SCHMALSTIEG D.: Towards massively multi-user augmented reality on handheld devices. In *Proceedings of the Third International Conference on Pervasive Computing (Pervasive 2005)* (Munich, Germany (to appear), 2005).