

Does text-based virtual space influence collaboration?

P. Dillenbourg¹, L. Montandon² and D. Traum³

¹ TECFA, School of Education and Psychology, University of Geneva, Route de Drize 9, 1227 Carouge Switzerland. Phone +41 (22) 705.96.93. Pierre.Dillenbourg@tecfa.unige.ch

² lydia.montandon@tecfa.unige.ch

³ UMIACS, A. V. Williams Building, University of Maryland, College Park, MD 20742 USA Phone: +1 (301) 405-1139

Abstract *This contribution reports the results of two experiments in MOO environments. MOO are text-based virtual realities. In some sense, they are more 'virtual' than 3D graphical systems since the user does not see rooms, objects or characters but only verbal descriptions of them. One may hence wonder whether the spatial metaphor of MOOs does impact on the users behavior. The first study reveals that users modify their communication behavior when they meet their partner in a virtual room: the change communication verb, acknowledge more often and more quickly their partner message. Surprisingly, despite this sensitivity to virtual co-presence, the subjects do not often explicitly coordinate their movements. The second study indicates that the rarity of explicit acts of spatial coordination is (at least partly) explained by the fact that the MOO automatically provides spatial information through side-messages. We conclude that the spatial metaphors in MOOs does impact on users behavior, but that space is here perceived as a social concept rather than a physical or geographical concept.*

1. Introduction

MUDs and MOOs (Curtis, 1993) are often described as text-based virtual realities. The users are located in a virtual space consisting of various rooms. They can move from one room to another, talk together inside a same room or across different rooms. The rooms, the objects in these rooms and the other users in the environment are described to users with short pieces of text, which are usually pre-stored. A MUD is basically a chat box (multi-user synchronous written communication) enriched with a spatial metaphor and the possibility to interact with objects. Moreover, objects maintain state which can persist from session to session, allowing extended collaboration. MOOs are a kind of MUD with an object-oriented programming language, allowing greater flexibility for customizing part of the environment.

In some sense, MOOs are more 'virtual' than 3-D virtual realities in which rooms, objects and users are graphically represented, such as DIVE (Hagsand, 1996) since the user perception of space and objects rely on his or her mental imagery. Usually the scene is very briefly described and the envisioning process relies heavily on naming the rooms by analogy with familiar spaces (offices, bars, meeting rooms, ...). At the opposite, in experiments on 3-D virtual realities (Witmer et al, 1996) revealed very detailed perceptive issues arise such as the user's ability to resolve detail, the breadth of the field of view, the lack of proprioceptive feedback, and so forth. This need for high definition graphics raise high technical requirements and one may wonder whether such a visual realism is necessary. For instance, Matsubara et al (1997) discuss how the 3D modelling can be simplified by replacing small objects (switches in an electric power plant) out of visual focus by textures in order to spare computing resources. This contribution addresses this issue under another angle: **does a text-based spatial metaphor significantly impacts on the user's behavior.**

We address these two issues from a **social perspective**: to which extent do spatial positions impact on communication and how do users maintain mutual knowledge regarding their respective positions. We report two laboratory studies in which pairs of users had to solve a well-defined task. These narrow scale studies enabled a detailed analysis of spatial coordination, as one aspect of the process of building a shared solution. Basically, spatial positions impact on the communication process through which partners elaborate a shared understanding of the task and its solution. Roschelle and Teasley (1995) defined collaboration by the process of elaborating and maintaining a shared understanding of the solution. We use MOO environments as tools for observing the construction of shared knowledge in order to develop new computational accounts of collaborative processes (Dillenbourg, 1996).

1.1 The function of space in a MOO environment

In the real world, the space impacts on our behavior in some functional ways (e.g. we cannot cross walls, we cannot see across a mountain, we cannot ski in a lift) and also by social uses (e.g. one should not sleep in one's office). A MOO environment reproduces some of these functional constraints (e.g. users must use doors to leave a room), but not all of them (e.g. tele-transportation is possible). We can express these functions with Clark and Brennan's (1991) terminology:

- Space controls **access**: Most commands can only be performed on the objects in the same room as the user (e.g. look, read, get, take, ...)
- Space controls **visibility**: Objects and partners can only be viewed (i.e. described by a piece of text) if they are in the same room; the action of another user is often only visible if (s)he is in the same room
- Space controls **audibility**: What user-A can hear depends on where user-B is and which communication command he uses. The command "Say hello" displays the "hello" message to all users in the same room, while the command "Page Kaspar hello" sends the message "hello" only to Kaspar, but wherever he is. Actually audibility is just a special case of visibility since messages are typed text.

In other words, by its very design, a MOO imparts two functional aspects to space: because access to information and visibility are space-bound, the spatial path reifies the information search process; and because visibility and audibility are space-bound, space contributes to differentiate private and public conversations. In addition to these functions, intrinsic to the design of standard MOOs, space can have task-specific functions. Typically, the importance of space is different if two users have to write a text collaboratively in the MOO or if they have to escape collaboratively from a labyrinth. The tasks selected for the empirical studies presented here differ by the extent to which spatial aspects play a role in the solution process. The user's behavior may be influenced by the virtual space beyond these constraints. For instance, if two users prefer to discuss their work when they meet in their virtual office while they would rather talk about their private lives when they meet in the virtual bar, this difference of behavior cannot be explained by the functional factors mentioned above, but only by the transfer of social habits from the real world into the virtual environment. In this contribution, we measure the impact of space **by detecting cases where space influences the users behavior beyond strict functional necessity**.

1.2 Types of spatial information

Users maintain different types of spatial information which we characterize along two axes, its complexity and the depth of mutuality.

1.2.1 Complexity of spatial information

The complexity of information is higher if the user maintains a representation of the global space versus his local space. The audibility functions only requires local space information: for instance, to choose between "say" and "page", one must simply know whether one's partner is in the same room or not. It is not necessary to know where he is in the global virtual space. Conversely, the other function of space, i.e. reifying the information search process, requires knowing where one's partner is exactly and who or which objects are located in the same room. This implies reasoning about topology and distances:

- **Topological reasoning**. If the user knows that the upper corridor has 4 aligned rooms and that his partner has visited the three first ones, he will certainly expect the partner to visit the fourth room. Actually, the topology is only partly determined by MOO information: knowing that a room has two exits does generally indicate whether the adjacent rooms are located on the right or left or even above, etc. For instance, figure 1 shows two maps collaboratively drawn by a pair of subjects visiting a sub-space of TecfaMOO. Both maps correspond to the description provided in the rooms (i.e. the number of exits and the connecting rooms, the position of exits was not provided).
- **Reasoning on distance**. The user may infer that his partner will take control of room-X since it is much closer to his partner's current position than to his own position. Distances can be estimated as the number of commands

necessary to move between two points (i.e. the number of intermediary rooms). The notion of distance is of course only meaningful if tele-transportation is not allowed.

As it is often the case, reasoning about space is not independent from reasoning about time. Regarding the global space, the user often has to memorize not only his partner's current position but also the (recent) previous positions. In terms of cognitive load we hence have three levels of complexity:

1. current position (in same room/out)
2. current global position (in room X)
3. path described as a sequence (e.g. room X,Y,...) or summarized by a direction (e.g. follow-corridor)

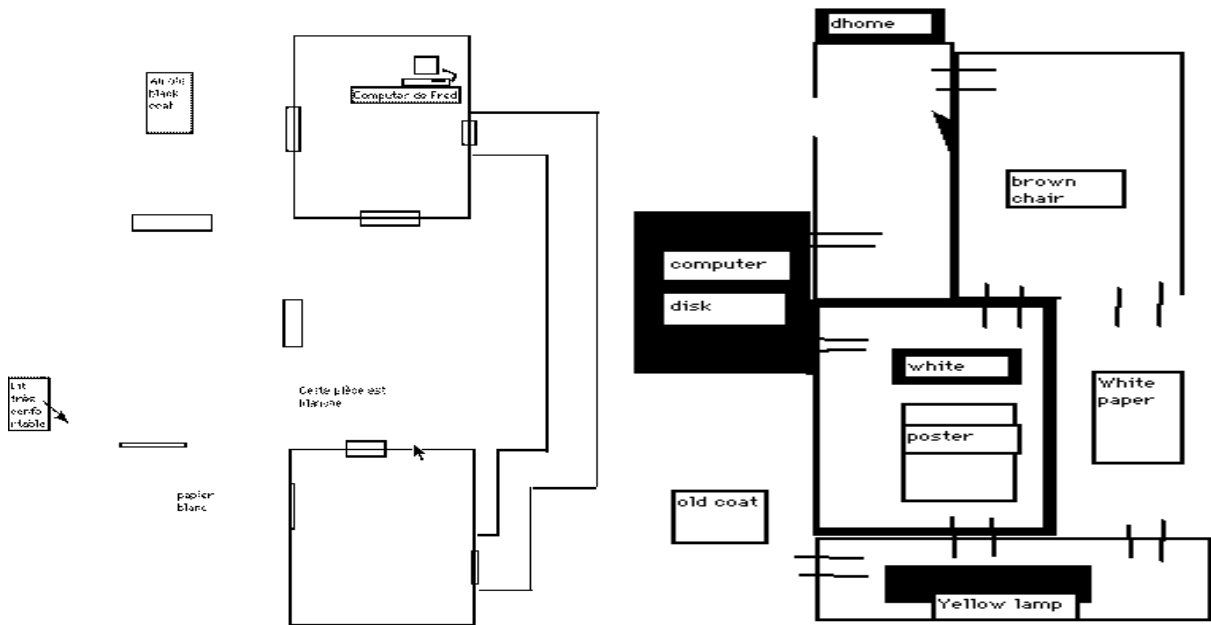


Figure 1: Two correct representations of a subspace of TecfaMOO, drawn collaboratively by pairs of subjects in a whiteboard (this exercise was part of the warm-up task to familiarize subjects with the environment).

The difference of cognitive load between local and global positions is not only in the complexity of the information itself, but also on the process for maintaining this information. When a user leaves or arrives in the room where I am, a MOO message is automatically generated. On the other hand, if I am in the bar and if my partner moves from the library to his office, I will not notice it unless I explicitly check position (e.g. using the "Who" command). One may hence expect that users almost always maintain information regarding the local space, but verify global space positions only when they have a good reason to need it.

1.2.2 Depth of mutuality

In collaboration, the elaboration of shared knowledge implies some mutual modelling, each partner having to model the knowledge of his partner: does he know X, does he know that I know X, does he know that I know that he knows X, etc. This knowledge is necessary to verify if one's partner has understood what was meant and repair eventual misunderstandings (Clark, 1994). When one considers the knowledge about spatial positions, one can also have multiple levels of mutuality of knowledge:

1. A knows where A is located;
2. A knows where B is located;
3. B knows that A knows where B is located;
4. A knows that B knows that A knows where B is located;
5. ...

The previous discussion on complexity of information was illustrated with level 2 (knowing one's partner position). However, it also applies at level 3 (does my partner know where I am. However, this level 3 requires a good knowledge of MOO commands. For instance, if Colin in the bar types "Page Kaspar hello", Kaspar will see a message such as "You sense that Colin is looking for you in the bar." followed by the actual message page by Kaspar. Hence, when Colin pages to Kaspar, Colin may infer that Kaspar has seen this side-message and hence knows where Colin is. Therefore, reasoning on mutual position is related with the level of MOO expertise.

1.3 Social Uses of Space

As well as the functional necessities of space, there are also some social aspects due to the way people are comfortable in collaborating with each other. While many of these aspects have functional utility in a general sense, their use is often primarily due to social and interactional reasons rather than strict functional necessity. Thus, even at times when it is not necessary for task performance to know where one's partner is, subjects will often try to keep in touch. Likewise, partners will often have a preference for a feeling of co-presence when having detailed discussions that does not require their presence in separate places. Standard MOOs subtly encourage this preference for virtual co-presence in the different styles of reporting of say and page messages.

2. First empirical study (Bootnap)

The first study was not designed specifically to address the issue of spatial coordination. The task was inscribed in a spatial context (a small Alpine Inn), but the solution process was not intrinsically spatial. The experiments mainly aimed to study how a whiteboard supports mutual understanding (Dillenbourg & Traum, 1997). However, the data collected provide some very interesting information with respect to spatial aspects.

2.1 Experimental setting

In these experiments, the two subjects play detectives in a mystery solving game: Mona-Lisa has been killed and they have to find the killer. They walk in a virtual Auberge (a subspace of TecfaMOO) where they meet suspects, ask questions about relations with the victim, what they have done the night of the murder, and so forth. Suspects are programmed robots. When exploring rooms, detectives also find various objects which help them to discover the murderer. They are told that they have to find the single suspect who (1) has a motive to kill, (2) had access to the murder weapon and (3) had the opportunity to kill the victim when she was alone.

The subjects communicate through the internet, each using a MOO client. As well as the standard MOO facilities, the detectives each carry a special "detective's notebook" which automatically records the answers to all the questions that they ask of suspects. They can also merge the contents of their two notebooks or exchange them, to find out directly what the other has been told. The MOO client uses is TKMOO-light (which runs on Unix/X workstations, Windows, and Macintosh computers) which also contains a shared Whiteboard: both users draw on the same page, can see and edit the objects drawn by their partner, but they do not see each other's cursor. All actions in the MOO and in the whiteboard are recorded. The subjects are provided with a map of the auberge (see figure 2) printed on a paper sheet. It must be emphasized however that the solution to the murder was not intrinsically spatial, i.e. it did not require inferences such as "Helmut cannot go from the bar to his room at 9PM without meeting Hans who left the bar at the same time, hence he lies.".




Figure 2. The map of the "Auberge du Bout de Nappe", the space that subjects had to explore to ask questions to suspects and inspect interesting objects.

We ran the experiments with 18 pairs (Dillenbourg & Traum, 1997). The subject were at different levels of MOO experience, ranging from novices to experts. The subjects were not used to working together with their partner. The average time to identify the murderer was two hours. The system automatically records all actions and interactions in the MOO and in the whiteboard. The complete protocols from MOO dialogues and the content of whiteboards are available on the WWW (<http://tecfa.unige.ch/tecfa/tecfa-research/cscps/Experiments/key.html>). We present here only some results which are relevant with respect to spatial aspects..

2.2 Space sensitivity in dialogue

When there are only two subjects in the global space, there is no more longer any functional difference between private and public communication. In this case, there is no functional necessity to differentiate the "say" and "page" commands. There could be a technical reason to do so, due to the syntax of these commands: typing "page Kaspar hello" is longer than typing "say hello". However, we controlled that parameter by allowing abbreviations for both verbs, respectively ' hello and "hello. Hence the production cost for communicating via 'page' or 'say' was identical. There is one other difference, still, between say and page -- this is in reception costs of messages. "Page" produces two lines to the recepiet - the message itself, and a header line saying where the pager is. It also produces an automatic acknowledgement (as long as the recipient is connected), rather than repeating the actual message to the sender. Since using "say" creates the risk that the partner does not receive the message (he was not there or left the room while one types the message), the optimal communication strategy, considering receipt of messages in these experiments would be to always use the "page" command. We observed that this not the case, on average the pairs use 'say' twice as often as 'page' for local talk. We do however observe a very large variety among pairs, as illustrated by figure 3, ranging form pairs using 'say' exclusively for talking within the room to pairs always using 'page', wherever they are. Pairs on the right hand side of figure 3 are those who are more sensitive to space, i.e. who will systematically use 'say' when they meet their interlocutor.

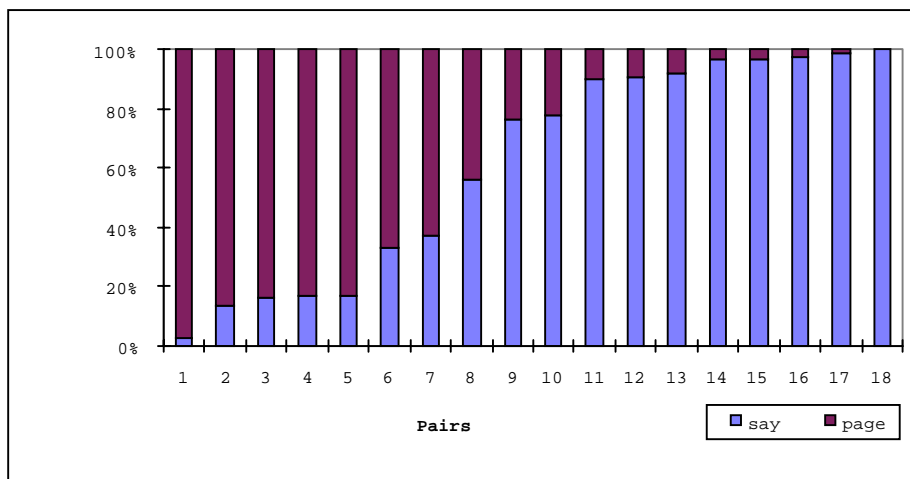


Figure 3. Average number of "say" and "page" commands used respectively in local talk (same room).

How do we interpret these results. First there is a strong correlation (0.69) between the say/page balance for each member of the pair. If one looks more closely at the data, we see that 7 subjects never used 'say' in local talk. The average percentage of their respective partners is 37%. At the opposite extreme, 6 subjects always use say for local talk, and the average percentage of their respective partners is 91%. We interpret this phenomena as social influence: one subject always/never uses the verb 'say' for local talk and influences his partner.

This spatial sensitivity is strongly related to MOO expertise. All 7 subjects who never used 'say' are MOO novices, while 9 out of the 10 subjects who almost always use 'say' (98%-100%) were experienced users. We also computed a 'space sensitivity' factor, which takes into account the cases where they use 'say' for distant talk (hence the message being not received). It is computed as the sum of 'say' commands performed when the subjects were in the same room

plus the number of 'page' commands when detectives were in different rooms, divided by the total number of messages. The average value for experienced MOO users is 75% for novices versus 87% for more experienced users ($F=4.39$; $df=1$; $N= 18$, $p=0.05$). This relationship between spatial sensitivity and MOO expertise is somewhat surprising. One could have expected that users with a good understanding of MOO mechanisms would notice that using 'say' is not relevant in this context. This is not the case. Hence, our explanation is that frequent users transfer the habits they acquire in previous experience, where using 'say' is relevant.

The conclusion is that for this experienced users, the choice of communicative verb goes beyond the level which could be explained by strictly functional criteria. This observation constitutes the first element of the stream of factors which convince us of the salience of the spatial metaphor.

2.3 Co-presence and acknowledgment

We coded all protocols and counted the rate of acknowledgment as the percentage of utterances being answered by the partner. Acknowledgment goes from elementary back-channel messages (such as "uh huh") up to elaborated answers, rephrasing, counter-arguments and so forth. The average rate of acknowledgment across all situations was 41%. We studied how this rate varies according to the spatial positions (same room or not). When the subjects are not in the same room, they acknowledge in average 34% of utterances. When they are in the same room, they acknowledge 50% of utterances. This consistent with Cherny's findings (1995): she observed that back channels are significantly absent from long distance conversation (page) versus co-present interactions (say). Once again this is surprising: a 'page' message is preceded by a side-message ("You sense that X is looking for you in room Y"), hence causes a more important visual change MOO window (scrolling).

This does not imply that subjects significantly meet for acknowledgment, this would be too expensive. However, it often occurs in these experiments that the detectives explicitly decide meet in the same room when they have long discussions, especially at the end of the task, when they have got all information and want to synthesize it. In table 1, Sherlock accepts (78.7) such a proposition, and verifies that they actually are together (79.1) before to resume discussion.

76.9	K	H	' so shall we meet to discuss our solutions?
78.6	K	H	who
78.7	r5	S	' Yes, let meet in the bar
78.8	r5	S	who
78.8	r5	S	go out
78.9	K	H	walk to bar
78.9	UC	S	go LC
79	LC	S	go B
79.1	Bar	S	who
79.2	Bar	H	ok, what's your guess?

Table 1: Subjects often meet for long discussions (from Pair 12) (Columns display respectively the time, location, subject and actions)

We do also observe that the delay of acknowledgment is shorter when subjects are in the same room. Independently from the use of 'say' or 'page', the subjects answer to each other after 39 seconds when they are in the same versus 59 seconds when they are in different rooms ($F=6.56$, $N = 18$, $df 1$, $p= .015$). A shorter delay may indicate a greater attention from the subject who answers and/or a tendency to give shorter answers (The delay includes the time necessary to type the answer because a message is processed only when the user hits the return key). This however not the case: the average length of 'say' and 'page' messages is almost identical (respectively 46.8 and 48.8 characters per message). We may hence interpret shorter response delay as a sign of a greater attention from the subjects who are in the same room. This confirm the previous observation that the space influences the users behavior beyond what is strictly functional.

2.4 Spatial coordination

Given the salience of the spatial metaphor, revealed in the previous section, we expected that the subject would maintain knowledge about their respective positions. However, we do not observe many cases where the subjects explicitly interact to maintain mutual knowledge about their current position ("I am in lobby", "Where are you ?", ...) or uses the command "who" (which indicates who is where). Our interpretation is not that the subjects do not know where their partner is, but rather that they maintain this knowledge without explicit acts. There are several ways in which the subject may know where his or her partner is:

- The Moo automatically provides information about mutual positions: every time a page is received, every time users meet or separate, plus every time where one sees the consequence of an action performed by one's partner, and so forth. It may be that the information so provided is sufficient for the task. The second study explores this hypothesis.
- The whiteboard indirectly provided information on the visited rooms. Most of the pairs put on the whiteboard a collection of small notes containing the information collected. This information was often structured by room, either implicitly or explicitly as in figure 4. It is hence straight forward to infer where one's partner is or has been. The partner's position could also be inferred from the information reported by the partner in MOO dialogues. However, in these experiments, the raw information was often displayed on the whiteboard. More elaborated inferences were shared through MOO dialogues, but these inferences do not necessarily match with the position of the detective when he emits it.
- Two pairs of subjects used the whiteboard more directly for maintaining mutual position. They both reproduced the provided auberge map on the screen. In one pair, the users pasted their initial ("S" for Sherlock, or "H", for Hercule) on the map and started to move it when they moved in the MOO. In this case, the cost of spatial coordination is very high, since subjects must update the graph every time they move. Position information does not remain true long enough to be advantageously displayed manually on a persistent medium. Each subject may doubt whether his partner has updated his position or not on the graph. This approach is only viable if it would be carried out automatically by the system.
- In another pair, the subjects pasted a "done" label on each visited room. This information is more suitable for a whiteboard since it is more persistent: The fact that Hercule has visited room 4, remains true. There is necessity here for subject to select a medium such as the persistence of information validity corresponds to the persistence of its display on this medium (Dillenbourg & Traum, 1997).

We also observe that subjects often talk about future positions ("I am going to the restaurant", "Where are you going?", ...). These utterances contribute to spatial coordination, although a detective can never be sure whether his partner will actually go to the agreed location nor how long (s)he will stay there. These utterances fulfill a more important function with respect to the problem solving strategy. Since, information is distributed over the virtual space (suspects and objects are in different rooms), negotiation where to go next is a concrete way of discussing how to collect information. We hence can not count these interactions as explicit acts of spatial coordination. After observing the spatial behaviours in this study, we designed a second study in order to observe specifically the mechanisms of spatial coordination, in a context in which it was more important to task success to coordinate on awareness of mutual position.

in the same room or when one leaves. The richness will only be related to the side-message which come with every 'page' communication. Moreover, the 'poor' condition is not extremely poor, since it still includes the 'who' command. The contrast between the two conditions is hence moderate.

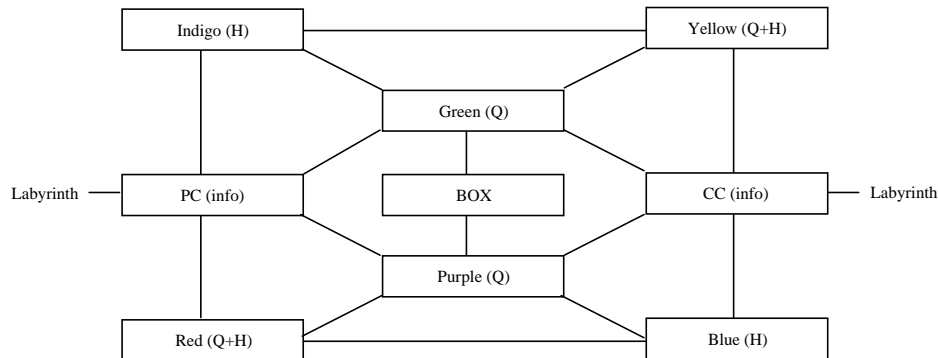


Figure 5. The spatial structure of task 1. (PC and CC are the starting rooms for each subject, Q indicate where questions were located and H where hints were located). The labyrinth structure is not represented here.

The experiment was run with 20 pairs, each pair passing through the two conditions- In order to counterbalance the order effect, 10 pairs did task 1 in the rich condition and then task 2 in the poor condition, while the 10 other pairs did task 1 in the poor condition and then task 2 in the rich condition (the balancing hence concerning the condition and not the tasks, considered as equivalent). The subjects were familiar with the MOO. Some of them did it from a machine in our lab, while others were connected from elsewhere in the world, via the Internet. A few subjects in remote places stopped in the middle of the experiment, without giving any explanation. They are not counted here.

Of the 20 pairs, 4 found the solution to both tasks, 8 solved only one task. Among the 8 pairs which did not solve either problem, 4 were close to the goal (3 letters identified out of 4), while 4 were far from the solution. Six pairs never went to the labyrinth, 10 went once, the remaining ones went between 2 and 6 times. On average they spent 37 minutes for the whole experiment.

3.2 Results

We counted three types of explicit acts of spatial coordination:

- EASC-current: Communication regarding the current position (e.g. "page castor I am in blue room" or "where are you?")
- EASC-future: Communication regarding a future position (e.g. "page castor I go to blue room" or "where do you go?")
- EASC-who: Using the 'who' command (which indicates who is where).

The third category differs from the previous ones with respect to the mutuality of knowledge. If Heidi pages "I am in the lobby", Kaspar knows where she is, but in addition Heidi knows that Kaspar knows where she is. If Heidi uses "who", she knows where Kaspar is, but Kaspar remains ignorant that she has verified his position. A type-1 or type-2 question followed by an answer was counted as one EASC. Table 2 gives the results:

# of EASC per minute	Condition "rich"	Condition "poor"
EASC-current	mean = 0.08, SD = 0.09	mean = 0.17, SD = 2.9
EASC-future	mean = 0.17, SD = 0.16	mean = 0.2, SD = 0.18
EASC-who	mean = 0.30, SD = 0.31	mean = 0.34, SD = 0.35

Table 2. Comparison of EASC per minute in the two conditions

Since the subjects spent significantly more time in the rich condition than in the poor condition, we counted the number of EASC per minute. Given the high variance of these data, the following results are presented with caution.

As expected by our hypothesis, the number of EASC-current is significantly higher in the 'poor' condition than in the 'rich condition' ($T(19) = 3.28, p < .05$). The number of EASC-future is also higher in the 'poor' condition, but the difference is not significant. The fact that our hypothesis is confirmed for EASC-current and not for EASC-future is logical: the side-messages provided by the MOO convey information regarding current position, but does not inform about future positions. The number of EASC-who is slightly higher in the poor condition than in the rich condition, but the difference is not significant. If we do not take the time into account, the subject even perform slightly fewer "who" commands in the poor condition. There is some redundancy between EASC-current and EASC-who, which may explain that those who perform more EASC-current do not need more EASC-who.

On average, the pairs perform more EASC in the second task, probably because they become more familiar with the task or the environment ($T(19) = 2.29, p < .05$). However, we observe an interaction between the order effect and the condition effect. The rich/poor difference is more important for the second task than for the first task. Our interpretation is the following. The pairs which pass from the poor to the rich condition continue to produce EASC, since they learned that it was important. Hence, the number of EASCs does not decrease in the second task. At the opposite, the pairs which pass from the rich to the poor condition, were used to receive spatial information from the MOO and hence perform many EASCs when this information is suppressed.

Surprisingly, the pairs which produce more EASCs are not more efficient in avoiding meetings. Some pairs do not coordinate action, they simply take risks. Other pairs establish from the beginning a distribution of rooms into territories in such a way that they should never meet (8 pairs did it for the first task, 4 for the second - 3 of them using these territories for both tasks).

4. Conclusions

If 3D-virtual spaces do not reproduce all the kinesthetic perception as a real space, they provide some visual and auditive stimuli which approximate a real movement in space. In MOO environments however, the virtual space is not described by images but by text. Virtuality is not a Boolean. One may expect that the feeling of immersion resulting from text description is lower than with 3D imagery. Most experienced MOO users would however disagree with this statement. Our results confirm this intuition in the case of peer problem solving. The first study reveals that users modify their communication behavior when they meet their partner in a virtual room: they change their communication verb, acknowledge more often and more quickly their partner message. However, despite this sensitivity to virtual co-presence, the subjects do not often explicitly coordinate their movements. The second study indicates that the rarity of explicit acts of spatial coordination is (at least partly) explained by the fact that the MOO automatically provides spatial information through side-messages. The design of MOO environments results from a culture which has grown within the MUDs/MOOs communities for many years. From an ecological perspective, these side-messages reflect the functional adaptation of these communities to the needs of the new virtual environments.

This study only covers one aspect of space, the social space, and ignores more intrinsically spatial aspects which imply the users mental imagery. Within the social aspects of space, we only consider the perception and effect of co-presence. Other social space phenomena could be observed in more ecological experimental settings, involving many users with various tasks, such as the emergence of group territories, the difference of behavior between rooms digged by the user and pre-existing rooms, etc.

However, within this narrow scope, our study establishes the influence of the spatial metaphor on the user behavior. Another issue is whether this metaphor makes collaboration more efficient. Our study does not answer this question. The arguments for defending the spatial metaphor generally rely on the hypothetical transfer of familiar schemes. Designing a virtual world on the basis of spatial aspects which shaped the development of our culture over thousands of years may give a feeling of familiarity, enable users to transfer social habits and hence facilitate orientation. However, if the cyberspace simply breaks one rule of real space, e.g. if one can open a room in the basement and arrive directly to the third floor, it may generate disorientation. The problem of hyperspace has been documented in empirical evaluations of hypertexts (Stanton, 1991). Alternatively, we may wonder why the cyberspace should reproduce most of constraints of the material world while humanity has been waiting thousands of years to escape from them.

5. Acknowledgments

The first empirical study was funded by the Swiss National Science Foundation (grant #11-40711.94). We would like to thank all the subjects in Geneva and elsewhere who participated to the experiments. Special thanks to Daniel Schneider, Richard Godard, Patrick Mendelsohn and Andrew Wilson.

6. References

- Cherny, L. (1995). The situated behaviour of MUD back channels. Internal Report. Stanford University, Linguistics Department.
- Clark, H.H., & Brennan S.E. (1991) Grounding in Communication. In L. Resnick, J. Levine & S. Teasley (Eds.), *Perspectives on Socially Shared Cognition* (127-149). Hyattsville, MD: American Psychological Association.
- Clark, H.H. & Wilkes-Gibbs, D. (1986) Referring as a collaborative process. *Cognition*, 22:1-39.
- Curtis, P (1993). *LambdaMOO Programmer's Manual*, Xerox Parc.
- Dillenbourg, P. (1996) Some technical implications of distributed cognition on the design of interactive learning environments. *Journal of Artificial Intelligence in Education*, 7 (2), pp. 161-179.
- Dillenbourg, P. & Traum, D. (1996) Grounding in computer-supported collaborative problem solving. FNRS #11-40711.94 research report. (soon on <http://tecfa.unige.ch>)
- Dillenbourg, P. & Traum D. (1997) The role of a whiteboard in a distributed cognitive system. Swiss workshop on collaborative and distributed systems. Lausanne, May 1st 1997.
- Hagsand, O. (1996) Interactive MultiUser VEs in the DIVE System, *IEEE Multimedia Magazine*, Vol 3, Number 1.
- Matsubara, Y., Toihara S., Tsukinari, Y. & Nagamachi, M. (1997) Virtual learning environments for discovery learning and its application on operator training. *IEICE Transactions on Information and Systems*, 80 (2), pp. 176-188.
- Montandon, L. (1996) Etude des mécanismes de coordination spatiale dans un environnement virtuel de collaboration. Mémoire de Diplôme d'Etudes Supérieures en Sciences et Technologies de l'Apprentissage. Non publié. TECFA, Faculté de Psychologie et des Sciences de l'Education, Université de Genève.
- Roschelle, J. & Teasley S.D. (1995) The construction of shared knowledge in collaborative problem solving. In C.E. O'Malley (Ed), *Computer-Supported Collaborative Learning*. (pp. 69-197). Berlin: Springer-Verlag
- Stanton, N.A. (1994) Explorations in hypertext: spatial metaphor considered harmful. *International Journal of Educational Technology*, 31 (4), pp. 276-294.
- Witmer, B.G., Bailey, J.H., Knerr, B.W and Parsons, K.C (1996) Virtual spaces and real world places: transfer of route knowledge. *International Journal of Human-computer Studies*, 45, pp. 413-128.