The role of a whiteboard in a distributed cognitive system.

P. Dillenbourg[1] *and D. Traum*[2]

[1] *TECFA, School of Education and Psychology, University of Geneva, Route de Drize 9, 1227 Carouge Switzerland. Phone +41 (22) 705.96.93*
[2] *UMIACS, A. V. Williams Building 9, University of Maryland, College Park, MD 20742 USA Phone: +1 (301) 405-113*

Address for correspondence: Pierre.Dillenbourg@tecfa.unige.ch

Abstract. This contribution reports the results of an empirical study on computer-supported collaborative problem solving. This study aimed to analyze how two users build a shared solution during collaborative computer-mediated problem solving. The collaborative environment included a text-based virtual reality - a MOO environment - and an elementary whiteboard. Our initial hypothesis was that the whiteboard would enable partners to draw schemata to repair the misunderstandings which inevitably occur in dialogue, like we draw a sketch on the napkin during a meal. The results contradicted this hypothesis. The whiteboard does not simply support verbal interactions. Instead, it appears as the central space of interaction, because the information displayed is persistent and because partners can assume shared visibility. The whiteboard reifies the problem states and hence plays an implicit but important role in mutual regulation. The content of information being mediated through the whiteboard varies from pair to pair. The difference between pairs can be expressed as different configurations of a distributed cognitive system, in which the collaborative functions implied by the task are allocated to different tools, but in which the whiteboard generally appears as the central tool.

1. Theoretical framework

The distributed cognition theories (hereafter 'DC' theories) offer an interesting theoretical framework to study collaborative problem solving. The common point of these theories is to consider that cognition is not bound to the processes which occur in our brain, but extends to the social and physical environment in which one acts and reasons. As Salomon (1993), we deliberately use the plural for distributed cognition theories. The broad range of theories can be classified with respect to their main source of influence. Some contributions, such as Hutchins (1995) heavily rely on concepts borrowed from cognitive science (information flow, memories, buffers,...), while other contributions such as Lave (1991) are inscribed in the continuation of socio-cultural theories. The empirical studies conducted on each side differ by their scale: while the former analyze in details the interaction in a small group, solving a task during a short period of time, the latter study the culture of larger groups doing a variety of tasks over a long period of time. This study belongs to the first approach: we look at rather short periods of time (2 hours) between two people who have a clearly defined task to do. We feel not only more comfortable with the conceptual framework, but also prefer its 'constructive' flavor: "*The question is not how individuals become members in a larger cognitive community as they do in apprenticeship studies. Rather the question is how a cognitive community could emerge in the first place*" (Schwartz, 1995, p. 350). We adopt a functional rather than a socio-historical view of culture, i.e. we aim to understand cultural tools as a group adaptation to its environment.[1]

There exist many definitions of collaborations and namely different of understanding of how collaboration differs from cooperation. The definition by Roshelle and Teasley (1995) has become widely accepted: "*Collaboration is a coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a problem* ". The process by which two participants progressively built and maintain a shared conception has been studied in pragmatics under the label 'social grounding'. Grounding is the process of augmenting and maintaining this common ground. It implies communication, diagnosis (to monitor the state of the other collaborator) and feedback (acknowledgment, repair, ...). Clark and Brennan (1991) established that the cost of grounding varies according to the medium. This is especially important in our experiments where the users may communicate through different tools, with distinctive features.

We discriminate different levels of mutuality of knowledge (Dillenbourg, Traum & Schneider, 1996). We transposed Clark's levels (1994), established for spoken conversation, to the peculiarities of virtual workspaces, namely typed communication and spatial metaphor. If agent A want to share information X to agent B, A may know more or less well whether about B knows X :

- A can infer that B can (not) <u>access</u> to X: For instance, in the MOO, A knows that if B is in room 7, where information X can be found, A may infer that B can potentially read this information.
- A can infer that B has (not) <u>perceived</u> X: For instance, if A writes a note on the whiteboard and B moves that note, A can infer that B has read it.
- A can infer that B has (mis-) <u>understood</u> X: For instance, in the MOO, if A says "let's ask *him* a few questions" and B moves to the room where "him" is located, then A can infer that B has well understood what she meant by "him".
- A can infer that B (dis-)<u>agrees</u>. This includes verbal agreement, but also agreement by action. An instance of non-verbal agreement in the MOO is when A says "Let's go to the kitchen" and B moves to the kitchen. An instance of non-verbal disagreement in the whiteboard, A write a note "Hans has a motive to kill" and B puts a red cross on this note or erase it.

This classification enables us to view grounding and agreement as different points in a continuum going from complete mutual ignorance to completely shared understanding. By extending the notion of grounding to the notion conflict resolution, we also relate this research with the socio-cognitive theory (Doise & Mugny, 1984).

In this paper, we frequently refer to the concept of persistence, but this concept covers different things. For instance, the information "It is 5 PM" is less persistent than the information "We close at 5pm". These two information pieces differ by their delay of validity. A note on a board saying "It is 5pm" is more persistent than a clock displaying "5:00 PM". These two pieces of information differ by the duration of their presentation. We respectively refer to these two aspects as **persistence of validity** and **persistence of display**. We will see that pairs selectively communicate different types of information through different media in such a way that the persistence of display is adapted to the persistence of validity. The persistence of validity is not only a function of time, but also of what makes the information obsolete. The validity of an information such as "It is 5 PM" is inherently limited in time, while the information "Jacques was

---

[1]. Let us however mention that the MOO environment used in this study would also be a useful tool to study the development of culture, such as in Reid (1994) and Bruckman (1992).

alone at 9PM" remains true until some evidence is brought that Jacques was not alone at that time. In other words, persistence of validity can be internally or externally limited.

2. Experimental setting

In the 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 text-based virtual world (a MOO environment) 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, they find various objects which help them to discover the murderer. They are told that they have to find the single suspect who (1) as 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 network, using a MOO system and a Whiteboard. The MOO environment is a standard MOO. The subjects move in different rooms and talk to each other via two commands: "say...", to communicate with anybody in the same room, and "page John..." to communicate with John where ever he is (See Example 1). The detectives carry a notebook which automatically records the answer to all the questions that they ask. They can merge the content of their notebooks or exchange them. The subjects use a MOO client - TKMOOlight - which runs on UNIX stations and includes an elementary whiteboard: both users draw on a 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.

Join sherlock

Auberge du Bout de Nappe: Lower Corridor
Obvious Exits: Lobby (to Lobby), UC (to Upper Corridor), B (to Bar), P (to Private Residence), R1 (to 1), R2 (to 2), R3 (to 3), and R4 (to 4).
Auberge Guest Room: 1
You see Rolf Loretan and Claire Loretan here.
Sherlock is here.
Obvious Exits: Out (to Lower Corridor).
Sherlock asks Claire Loretan about last night
Claire Loretan answers "I was in the restaurant with my husband and the Vesuvios. When the restaurant closed, I briefly went to my room and then joined the others in the bar."
Sherlock asks "Do you know when the bar has closed?"
wisper Did you notice that he is an insurance agent?
I don't understand that.
"what are doing?
You ask, "what are doing?"
ask rolf about the gun
hercule asks Rolf Loretan about the gun
Rolf Loretan answers "i looks like a military issue gun. Why don't you ask that Colonel?"
Sherlock says "Forget it. I thought it could help if we make a tab with the informations about where were th people at what time."
"Actually sounds a good idea.
You say, "Actually sounds a good idea. "
"I think we should find more information about the gun
You say, "I think we should find more information about the gun"

Example 1: An excerpt from the MOO window. Bold lines are entered by the user (Hercule).

3. Observations

We ran the experiments with 20 pairs. The subjects were not used to work together. The average time to identify the murderer was two hours. The complete protocols from MOO dialogues and the content of whiteboards are available on WWW[2]. A more complete analysis can be found in Dillenbourg and Traum (1997). We present here only some results which are important with respect to the role of the whiteboard in collaborative problem solving..

3.1 Envisioning the solution

The first surprise when one looks at the whiteboards drawn by the pairs is that they are often made of a large collection of text notes. We did not encounter many elaborated graphics. On the 20 pairs, four drew a timeline (e.g. figure 1), four drew a map (e.g. figure 2) and three drew a graph, mainly indicating social relations among suspects (e.g. figure 3). Timelines can be helpful since its is difficult to reason about intervals without visualizing them. However only one complete time line has been produced. The maps did not really help to solve the task since the solution of the enigma did not require any spatial reasoning such as "Hans could not go from room 1 to the bar without being seen by Rolf." The pairs who drew a map where probably influenced by the spatial metaphor very salient in the task. Finally, the graphs were generally abandoned before to be completed.

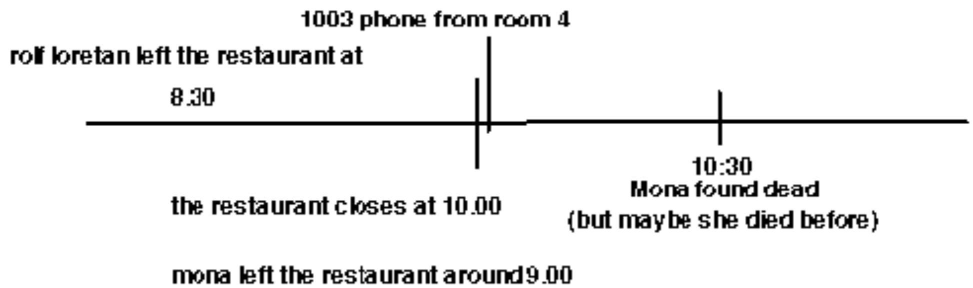

Figure 1: Uncompleted timeline in Pair 22.

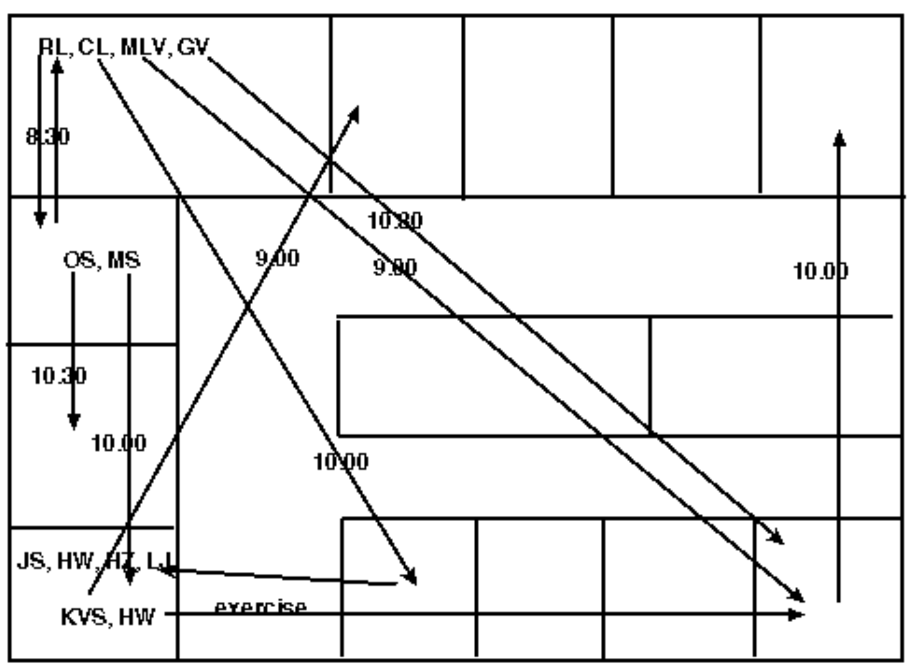

2 http://tecfa.unige.ch/tecfa/tecfa-research/cscps/Experiments/key.html

Figure 2: Representation of suspects' moves on a map (from Pair 16).

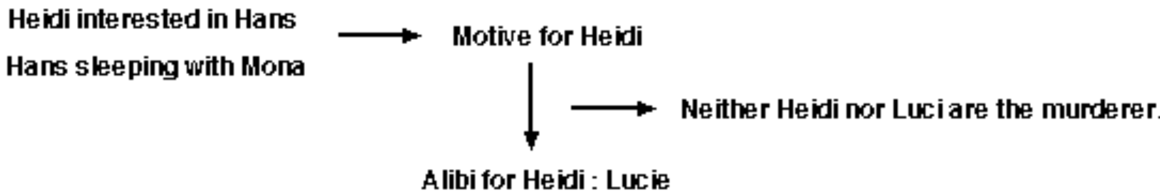

Figure 3. A graph of inter-suspect relationship (from Pair 18)

In brief, the expressivity of graphics is only partly exploited by the subjects. Our interpretation is simply that the task did not require such explanatory schemata: it did not conveyed unknown concepts to be illustrated or complex causal structures to be articulated. The main difficulty of this task was the information load, which explains that the whiteboard was mainly use as an information manager (see section 3.3). Nevertheless, some graphical features of the whiteboard were used to structure the collection of notes. Three pairs use color codes for indicating the author of the note. Many pairs used the 2-D space to structure the information, either systematically in a two-entry table (2 pairs), more often, by grouping the notes concerning the same suspect.

3.2 Disambiguating dialogues

The absence of elaborated schemata, pointed out in the previous section, does not only affect individual reasoning, but also communication. Schemata are useful for **clarifying an idea** which is difficult to turn into words. Since we observe few schemata on the whiteboard, we may reject our initial hypothesis regarding the role of the whiteboard: in these experiments, it was not used to prevent and repair misunderstandings in dialogue. Again, this observation is bound to the task which did not involve the type of misunderstanding which require schemata.

Nevertheless, the whiteboard can also support mutual understanding by helping partners **solve references in dialogue**. An utterance such as "*Put this one there*" makes only sense if it accompanied by gestures indicating to which object refers "*this one*" and to which location refers "*there*"[3]. In a related study, we observed 8 pairs of subjects sitting in front of a same machine. They had to solve a problem in a learning environment called Memolab (Dillenbourg et al., 1994). All pairs together, these subjects accomplished 878 gestures in front of the screen and 87% of them were simple deictic gestures (Roiron, 1996). At the opposite, in the experiment we report here, deictic gestures were not observed for two reasons. First, the users could not see each other's cursor. Second, even if cursors were visible (it was possible to put a small mark instead or to move the object being referred to) MOO-dialogues are written and it is difficult for the speaker to simultaneously type "here" and move the cursor wherever "here" means (and similarly, it would be difficult for his partner to read the message in the MOO window and to follow simultaneously cursor moves on the whiteboard).

Hence, we could conclude that the whiteboard and the MOO form two rather separated spaces for dialogue. We even see cases where the subjects discuss about one thing in the MOO and interact about another point on the whiteboard. However, we do also observe cases (about 6% of all acknowledgment) of **cross-modality interactions**: an utterance in the MOO is acknowledged by an action on the whiteboard and vice-versa.

---

3 In some specific dialogues contexts, these deictic gestures are not 100% necessary.

- In talk-whiteboard acknowledgment, we have two types of interaction patterns: either one subject invites (through dialogue) his partner to do something on the whiteboard (example 2), either the whiteboard is used to archive some information which has previously been agreed upon in dialogue (example 3).

64.4	r1	H	And what about Rolf?
64.8	r1	H	Are you going to put him on the whiteboard?
64.9	r1	S	draw black text 499 362 "ROLF"

Example 2: Acknowledging a request for action on the whiteboard (from Pair 18)[4].

73	Lobby	H	say do we have someone with a good alibi?
73.7	Lobby	S	yes, Lucie and Heidi
74	Lobby	H	say so we can make a red rectangle around them...
74.3	Lobby	S	ok
74.3	Lobby	H	draw red rectangle 10 167 346 209
74.5	Lobby	H	draw red rectangle 8 105 416 138

Example 3: The whiteboard for achieving shared inferences (from Pair 14)

- In whiteboard-talk interactions, the dialogue generally serves to ground the information which has been put on the whiteboard, like in examples 4 and 5. This leads us to our third point: when there is a multimodal interaction, involving MOO dialogues and whiteboard interactions, it is often the whiteboard which is the central space, and the MOO dialogue which helps to clarify it, and not the opposite as we initially hypothesized it. Our interpretation of the importance of the whiteboard is that the whiteboard information is more persistent than the MOO information.

138.1	r4	S	draw red rectangle 19 721 150 804
138.9	r4	H	say what is the red square for ?
139.2	r4	S	Nothing, my screen was frozen

Example 4: Clarifying the meaning of a whiteboard object (from Pair 22)

48.3	K	S	draw black text 233 216 "Oscar saleve is a liar "
49.5	Bar	H	'sherlock How do you know Oscar is a liar?

Example 5: Request for justifying a whiteboard note (from Pair 21).

3.3 Private and shared memory

As we said earlier, most whiteboards include a large set of notes. These notes can be viewed as a support for both individual and group memory. We hence analyzed the content of these notes. Actually, we coded all interactions, both MOO utterances and whiteboard items, with respect to their content. We discriminate 3 main content categories: task-level information (which suspect was where, who has a motive to kill, ...), meta-level information (discussing the problem-solving strategy, e.g. where should I go, ...) and meta-communication (discussing how to interact, e.g. turn taking, delay, ...)[5].

---

4 The first column indicates the time; the second the subject's location in the virtual space (room1); the third column is the subject (H=Hercule, S=Sherlock); the fourth column show what they in dthe MOO or *draw* on the whiteboard.

5 A fourth category as been added: when partners discuss about technical problems.

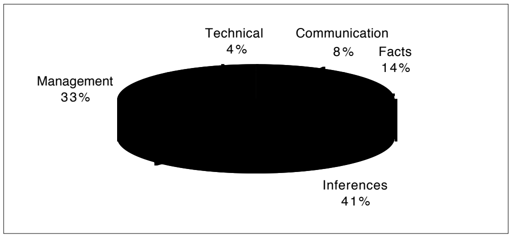

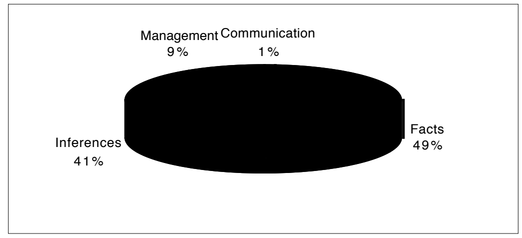

Figure 1: Categories of content in MOO (left) and whiteboard (right) interactions

Figure 4 compares the content of interactions in MOO dialogues (4a) and through the whiteboard (4b). The main difference between MOO and whiteboard interactions is that the subjects do not explicitly interact about their strategy on the whiteboard. Our interpretation is the following: the display persistence on whiteboard is high, while the validity persistence of strategical knowledge is low. Utterances classified as "strategy" includes statements such as: "I am in the lobby", this information being valid for a few seconds (until the subject moves), "I am going to ask questions to Heidi", this information being valid for a few minutes, or "I take the upper corridor and you take the lower", this information being valid for 30 minutes. If such an information is put on the whiteboard, it will be displayed until a subject erases it. There is a high probability that it will still be displayed while it is not true anymore. In brief, if the display persistence is longer than the validity persistence, one comes to display obsolete information unless a systematic effort of maintenance is carried out. Conversely, task-level information is more persistent.

Within the task-level knowledge, we discriminate facts, which are information obtained directly from the environment, by asking questions to the suspects ("Lucie went out at 9") and inferences, which is new information derived by the subjects (e.g. "Giuzeppe was jealous"). Facts are fully persistent: an information found at time t remains true at time t'. Inferences are internally persistent but externally refutable, i.e. they remain true as long as they is not refuted. We might then expect that inferences are less displayed on the whiteboard than facts. This is true for some pairs who prefer to negotiate inferences in the MOO (especially the pairs which filled their whiteboard with notes). However, persistent displays do actually also support information refutation, since they enable the refuter to come back on any information in order to discard it under the light of some new information. This explains that both facts and inferences are negotiated on the whiteboard.

3.4 Supporting regulation

The fact that the whiteboard does not include explicit notes regarding the strategy does not imply that it does not help to elaborate a joint strategy. To illustrate this, we first describe the problem solving strategy. It generally includes two processes: **data acquisition** - the subjects ask questions and examine objects - , and **data synthesis** - the subjects relate pieces of information in order to find the murderer. For forward chaining pairs, these two processes constitute to successive stages: they first collected all possible facts and then draw inferences. For backward chaining pairs, these two processes are intertwined: the first collected facts lead to some suspicions on the basis of which new facts are collected and so forth.

The whiteboard help to regulate the pair actions:

- During data acquisition, several pairs split spatially. In general, they agree initially - implicitly or explicitly - that one detective will visit the rooms of the upper corridor while the other will take the lower corridor. Very often this division of labor does not hold, because one detective is faster than the other and/or because this division does not specify who will explore the rooms which do not belong to any of these corridors (the bar, the restaurant, ...). This insufficient specification of the strategy should lead to complementary negotiation of who does what. This is not the case. It could be explained by the fact that the subjects used MOO commands such as "who" to trace their partner itinerary in the MOO.

This is also not the case, we observe very few command such as MOO or messages such as "where are you"[6]. Another explanation of this efficient coordination is that each agent can 'trace' his partner by looking at the information he puts on the whiteboard: If Hercule sees that Sherlock puts on the whiteboard the information collected from Lucie Saleve, he may infer that Sherlock is in that room. The whiteboard gives a permanent display of how much information has been collected and, from there, an estimation of how much has still to be collected.

- During data synthesis, a frequent method is to start with all the suspects and to discard one by one any suspect who had either no motive to kill, either no opportunity to get the weapon, either no opportunity to kill (cf. Figure 5). This process takes generally a concrete form on the whiteboard: the detectives cross one by one the notes regarding any suspect which is discarded. The whiteboard provides them permanently with the set of remaining suspects.

In brief, both during data acquisition an data synthesis, the whiteboard reifies the problem state. And the problem state is a key information to coordinate action.

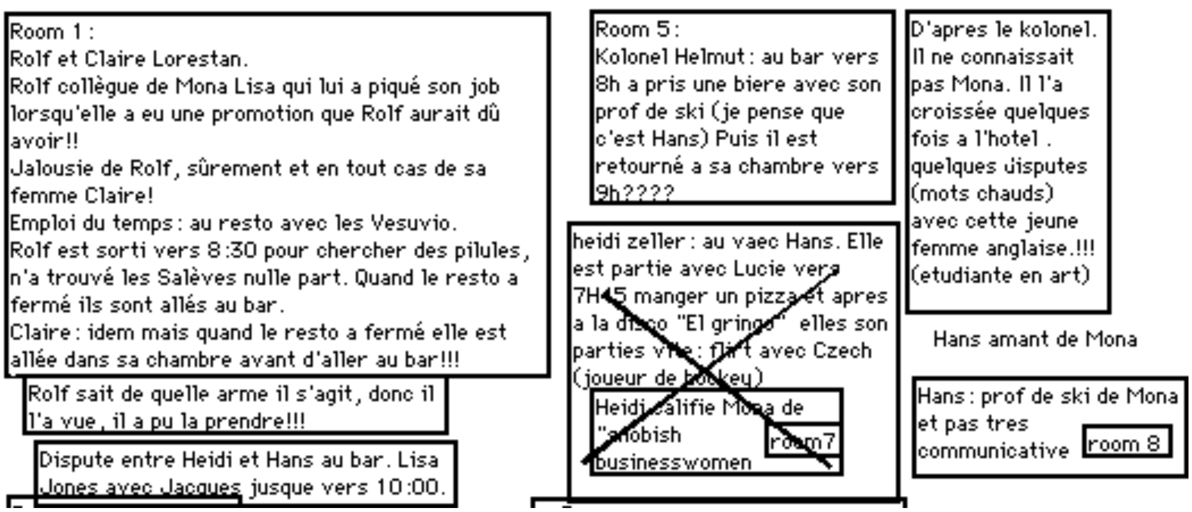

Figure 5: Associating elements on spatial basis (from Pair 7)

3.5 Shared visibility

In this study, we set up the size of whiteboard in a way that no scrolling was possible. Hence, partners could postulate that they both see exactly the same information. This is of course the raison d'être of a whiteboard. The shared visibility of information can also exist in the MOO but only when the two partners are in the same room and are mutually aware of it. On the whiteboard, shared visibility is permanent. This impacts on the way subjects interact. We computed the rate of acknowledgment, i.e. the ratio of MOO utterance or whiteboard notes which are acknowledged by the partner. The average rate of acknowledgment for inferences (46 %) is higher than for facts (26 %), which can be interpreted by the fact that there is nothing much to disagree or to misunderstand about facts. More interestingly, there is a significant interaction effect on the acknowledgment rate

---

[6] Actually, the subjects can infer where there partner is when this partner use a 'page' command. When Hercule pages Sherlock from room 5, Sherlock sees a message such as "You sense that Hercule is looking for you in the library.". For a longer discussion of spatial coordination problems see Dillenbourg, Montandon and Mendelsohn (in preparation).

between the knowledge category and the mode of negotiation (F=6.09; df=2; p=.001): Facts are rarely acknowledged on the whiteboard (Fig.6). Our interpretation is the following. Since there is a low probability of misunderstanding or disagreeing about facts, their acknowledgment in MOO conversation (37%) basically means "ok, I read your message". Acknowledgment simply aims to inform one's partner about shared perception (level 2). Conversely, on the whiteboard, mutual visibility is the default assumption, which makes acknowledgment unnecessary.

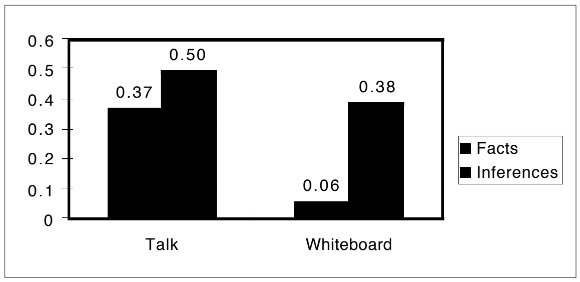

Figure 6: Interaction effect on the acknowledgment rate between the mode of interaction and the content of interaction.

4. Discussion

These results lead us to a conclusion which is the reverse of what we expected: the whiteboard is not used to support MOO dialogues, instead it is the central space of interaction, the place for sharing any important information. Our hypothesis is that it fulfills this role because of two features: persistence of display and shared visibility. Hence, our observation cannot be generalized to any system labeled 'whiteboard' by its designer, but instead to any system including a space of shared and persistent information. We could for instance design a MOO client which includes a text-based area with these two features. Moreover, this generalization is limited by the task we selected: the main difficulty of the task was information management, hence it is not surprising that persistence showed up to be a key feature.

It would be however incorrect to talk about the 'role of the whiteboard' as it was used by all the pairs in the same way. Basically the task involves 6 functions (collecting facts, sharing facts, sharing inferences, storing facts or inferences and coordinating action) and 4 tools (MOO dialogue, MOO action, whiteboard and merging/reading notebooks[7]). Table 1 indicates how these different tools support different functions. The matrix in table 1 is theoretical. A matrix corresponding to an actual pair is less redundant: For instance, if a pair communicates all facts through dialogues, the whiteboard will be globally more available for inferences; if a pair exchanges many facts through notebooks, it will communicate fewer facts through dialogues, and so forth. The actual [function X tool] matrix varies from one pair to another. It may also vary within a pair as the collaboration progresses, one function being for instance progressively abandoned because the detectives become familiar with another one. This **plasticity**, this ability to self-organize along different configurations justifies the descriptions of a pair as single cognitive system.

Tool / *Function*	MOO dialogue	MOO action	Whiteboard	Notebooks

[7] From a DC perspective, the two users should also be included in this matrix, but this would make it less readable since they are involved in each function,

Collecting facts		X		
Sharing facts	X		X	X
Sharing inferences	X		X	
Storing facts			X	X
Storing inferences			X	
Coordinating action	X	X	X	

Table16: Configuration of a cognitive system: matrix of possible allocations of functions to tools.

The CMC system is not just a neutral communication channel. It carries out some functions of the problem solving process. Hence, if we want to abstract observations beyond the particular technological environment which has been used, one has to reason from a distributed cognition perspective, i.e. to consider the software tools and the users as different components of a single cognitive system. Without this systemic view, we would continue to ask why a cyclist moves faster than a runner, despite a similar rate of leg movements per minute.

5. References

Clark, H.H.(1994) Managing problems in speaking. Speech Communication, 15:243 – 250.

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.

Dillenbourg, P. , Traum, D. & Schneider, D. (1996) Grounding in multi-modal task-oriented collaboration. In P. Brna, A. Paiva & J. Self (Eds), *Proceedings of the European Conference on Artificial Intelligence in Education*. Lisbon, Portugal, Sept. 20 - Oc. 2, pp. 401-407.

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., Mendelsohn, P. & Schneider, D. (1994) The distribution of pedagogical roles in a multi-agent learning environment. In R. Lewis and P. Mendelsohn. Lessons from Learning (pp.199-216) Amsterdam: North-Holland.

Doise, W. & Mugny, G. (1984) The social development of the intellect. Oxford: Pergamon Press.

Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.

Frohlich, D.M. (1993) The history and future of direct manipulation, Behaviour & Information Technology, 12 (6), 315-29

Gaffie, J. (1996) Etude des stratéges de recueil de données lors d'une résolution de problème en collaboration dans le MOO. Rapport de recherche non-publié. TECFA, Faculté de Psychologie et des Sciences de l'Education, Université de Genève.

Lave J. (1991) Situating learning in communities of practice. In L. Resnick, J. Levine & S. Teasley (Eds.), Perspectives on Socially Shared Cognition (63 - 84). Hyattsville, MD: American Psychological Association.

Roiron, C. (1996) Expérimentation d'un logiciel éducatif utilisant des techniques d'intelligence artificielle. Rapport de recherche 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

Salomon, G. (1993) No distribution without individual's cognition: a dynamic interactional view. In G. Salomon. (Ed). *Distributed cognitions. Psychological and educational considerations* (pp. 111-138) Cambridge, USA: Cambridge University Press.

Schwartz, D.L. (1995). The emergence of abstract dyad representations in dyad problem solving. The Journal of the Learning Sciences, 4 (3), pp. 321-354.