Coming to the Wrong Decision Quickly: Why Awareness Tools Must be Matched with Appropriate Tasks

Alberto Espinosa, Jonathan Cadiz, Luis Rico-Gutierrez, Robert Kraut, William Scherlis
Carnegie Mellon University
Pittsburgh, PA 15213 USA
+1 412 268 7694
robert.kraut@cmu.edu

Glenn Lautenbacher
School of Information Science
University of Pittsburgh
Pittsburgh PA 15217 USA
glenn@sis.pitt.edu

ABSTRACT
This paper presents an awareness tool designed to help distributed, asynchronous groups solve problems quickly. Using a lab study, it was found that groups that used the awareness tool tended to converge and agree upon a solution more quickly. However, it was also found that individuals who did not use the awareness tool got closer to the correct solution. Implications for the design of awareness tools are discussed, with particular attention paid to the importance of matching the features of an awareness tool with a workgroup's tasks and goals.

Keywords
Task awareness, workgroups, awareness devices, computer-mediated communication, distributed work, asynchronous work

INTRODUCTION

The group is the fundamental unit of work in organizations [15, 25]. However, recent trends in globalization, downsizing, and outsourcing are changing the look of group work. It used to be a safe assumption that most members of a group worked in the same office at the same time; however, it is not unusual today to find members of groups distributed throughout the world, forcing team members to work in different places and often at different times.

One of the key issues for these distributed, asynchronous groups is how to solve problems quickly and efficiently. Problem-solving often involves exploring and researching several options, weighing the advantages and disadvantages of each, and then choosing the best option based on all the information available. However, when team members are not working in the same place at the same time, important awareness information is often lost. Examples of lost information include who has explored which options, who knows what information, and what people's lines of reasoning are about the best possible solution.

Thus, the goal of this research is to design and test a tool to help distributed, asynchronous groups solve problems by providing task awareness. It is hypothesized that by providing task awareness, group members' lines of reasoning about the best possible solution will converge more quickly, thus helping the group to solve the problem faster.

Groups provided with a task awareness tool did exactly as we predicted: they came to decisions more quickly. However, groups without the tool came closer to the correct solution, necessitating a closer look at how awareness information can affect behavior in asynchronous, distributed group work.

In the following sections, we discuss the research on which this project is based, the design of the awareness tool, the method used to test the awareness tool, and our findings. We will also discuss the implications of our findings, especially in regard to designing awareness tools with features that are appropriate for the nature of a group's task.

PREVIOUS RESEARCH

Group awareness has been defined as "an understanding of the activities of others, which provides a context for your own activity" [7]. The value of providing awareness to teams is suggested in the literature, which indicates that members of workgroups will be more successful if they maintain awareness of the state of the team, task, and environment [4, 20]. It is also suggested that simple awareness of one's colleagues is a strong predictor of success in collaborations, thus highlighting the importance of awareness for team performance [12].

The Task: Distributed Problem-Solving

However, the context of a group's work cannot be ignored when discussing the merits of awareness. The nature of the task makes a difference when studying group performance [18], particularly with computer mediated group interaction [19].

Thus, it is important to note that our research focuses on distributed problem-solving under asynchronous conditions. Schlichter, Koch, and Bürger [23] define distributed problem-solving as the cooperative activity of several decentralized and loosely coupled problem-solvers acting in separated environments. A key point for successful collaboration in distributed problem-solving has to do with how well individual work relates to the objectives of the group as a whole.

We believe one way to help distributed problem-solving groups is by replacing some of the awareness information that is lost when team members do not work in the same location. However, choosing what kind of awareness information to provide can be difficult since the literature defines several types and taxonomies of awareness that could be given to groups [1, 5, 9, 10, 13, 14, 26]. Among these descriptions, the most useful model for this research is the framework of task awareness provided by Chen and Gains [5].

The Focus: Peripheral, Chronological Awareness

Chen and Gaines [5] identify two analogous forms of awareness based on a collective intelligence model:

1) Collective awareness: awareness of the group's collective, long-term memory
2) Awareness of teammates

The first form, collective awareness, exists in groups at three levels of detail: deep awareness (a highly detailed level of information), peripheral awareness (a less detailed, but still substantial level of information), and global awareness (a relatively low level of detail). Choosing the correct level of detail is important since providing too little information could be useless, while providing too much information could overload a person and reduce the amount of time and cognitive resources available to work on the primary task [11, 25, 26]. Taking this into consideration, we believe peripheral awareness is the correct level of detail for asynchronous, distributed, problem-solving groups. Peripheral awareness can help to provide some of the contextual information typically unavailable in the absence of physical proximity and face-to-face interaction.

The second form, awareness of teammates, can be of three types: resource awareness (who has what expertise within the group), task-socio awareness (information about social/political dynamics within the group), and chronological awareness. Chronological awareness is the instantaneous awareness that an individual has regarding the activities of others and knowing that something has changed—i.e., what, when and by whom.

We believe chronological awareness has the most potential to help asynchronous, distributed, problem-solving teams. First, knowing the activities of one's teammates can help the team to divide labor quickly and efficiently. If a person knows that a team member has already explored a certain solution, then a duplication of effort can be avoided. Second, a lesser amount of direct communication is necessary if team members know what each other are doing. Third, knowing the activities of one's teammates can help a team to form a shared mental model of the work, which may trigger new lines of reasoning about how to solve a problem [17].

For these reasons, the awareness tool we developed provides members of asynchronous, distributed, problem-solving teams with peripheral, chronological awareness.

THE AWARENESS TOOL

Figure 1 shows the awareness tool we designed. (This tool was adapted from a similar system developed to deliver awareness information to MBA students at Carnegie Mellon University [11]). Participants in our lab study had to decide how to treat a cancer patient by exploring a set of documents, each with information that may or may not be helpful. Examples of documents include x-rays, results of blood tests, and articles describing types of therapies.

In this study, the set of documents was large enough such that one person could not read and evaluate all the documents in the time allowed. Furthermore, participants could not read any document any time they wished. Sometimes documents were "denied" to participants, meaning that the group had not fulfilled the prerequisites to view a document. For example, an x-ray should not be ordered for a woman before checking her medical history to ensure she is not pregnant, or an expensive test should not be run before reviewing the patient's basic history. Prerequisites are applied at the group level and structured such that using common sense and good judgment should result in fewer denials.

The awareness tool provides information about what

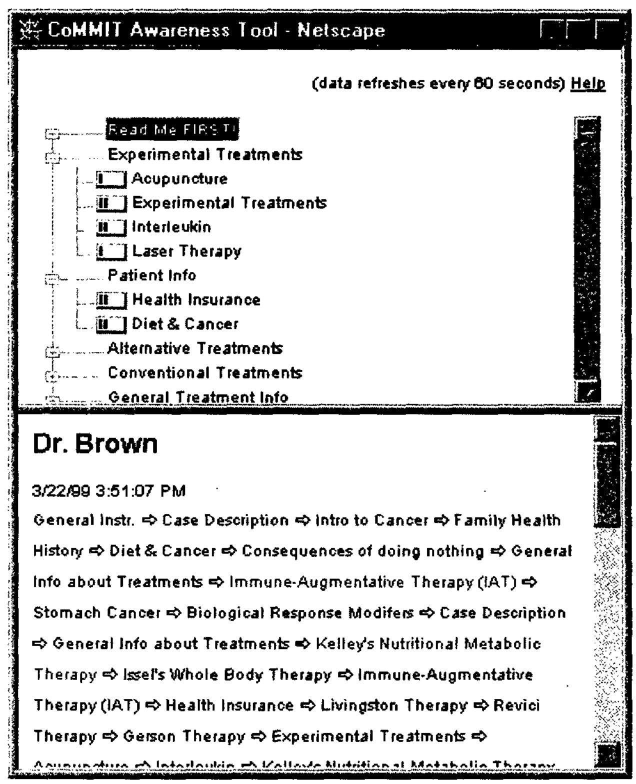

Figure 1: The tool designed to provide peripheral, chronological awareness. Items are color coded so that successful requests are blue while denials are red.

documents have been explored and denied using two displays. The top half of the window shown in Figure 1 contains all the documents in the document universe, structured as a two-level tree. The first level has all the document groups while the second level displays all the document names. An icon with zero to five bars (corresponding to the amount of activity associated with the document) is displayed along with the document name. If the last request for the document was denied, the bars are red. If the last request was successful, the bars are blue. Activity is represented on a five-point scale where each attempt to view a document increases the amount of bars shown. Activity is adjusted for time on a logarithmic scale; thus, recent activity will generate more bars.

The bottom section of Figure 1 contains a chronological listing of all the documents explored by one's team members. As in the top section, successful document requests are blue and denials are red. Such peripheral, chronological awareness information is intended to help individuals infer and follow their team members' lines of reasoning.

The awareness tool was implemented using Microsoft's Visual Basic and Active Server Pages. The data were accessed from a relational database kept by CoMMIT, the system participants used to explore documents and discuss solutions with team members.

The COMMIT System

CoMMIT (Collaborative MultiMedia Instructional Toolkit) is a computer system that facilitates asynchronous, distributed problem-solving. It is a web-based learning system designed to help groups solve diagnostic problems [16]. CoMMIT is used to explore documents with information about medical patients and to discuss possible diagnoses with teammates. As team members move through the documents, they can record notes on the documents and read comments from other team members.

The main window of the CoMMIT system is shown in Figure 2. Documents are structured into groups. To access a document, the user first clicks on one of the groups (labeled area 1 in Figure 2), which causes the documents in that group to display below (area 2). The user then clicks one of the documents in area 2, which causes the contents of the document to appear in area 3. Each document has three portions: the material with information about the problem (area 3), comments from team members (area 4), and a text box where additional comments can be made by the user (area 5).

Note that the group labels in area 1 of Figure 2 correspond to the groups shown in the awareness tool in Figure 1. Similarly, the document names in area 2 of Figure 2 correspond to the document names in the awareness tool.

In addition to the window shown in Figure 2, CoMMIT provides users with an additional "Notepad" window (not shown) that lists all team members' comments on all documents. The notepad window also provides a separate discussion area where team members can talk about strategies, share observations and thoughts, and decide upon the final solution.

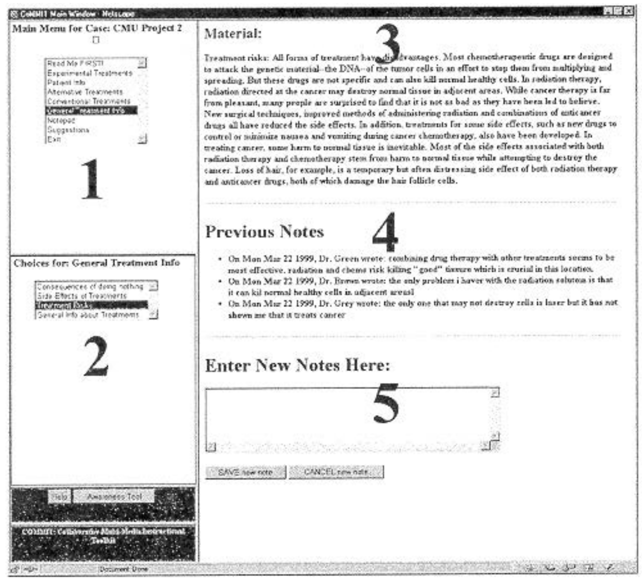

Figure 2: The main window of the CoMMIT system. The "Awareness Tool" button was only visible for participants in the experiment condition.

Even though CoMMIT provides the team with some degree of awareness about each other's thinking about specific documents and the problem in general, it does not help members to know easily who has explored which documents and in what order documents were explored. For this reason, the CoMMIT system was chosen as the focus for this research.

HYPOTHESES

Technology affects social interaction [2, 3, 6]. As such, we speculate that an information technology that provides peripheral and chronological information of team member's document reading behavior will have an effect on team interaction. More specifically, we speculate that peripheral, chronological awareness of team members' information search patterns will help people to infer and follow lines of reasoning of other teammates, thus helping the team converge toward a joint solution.

H1: Awareness of peer document search patterns helps convergence of lines of reasoning of team members, thus helping them achieve consensus on joint team solutions of diagnostic problems.

We anticipate that this more effective team problem-solving process will translate into higher team performance. Furthermore, we also anticipate that peripheral and chronological awareness of information search patterns will help teams develop more effective problem-solving strategies, such as dividing labor (e.g., deciding who needs to read what), developing shared mental models (e.g., following another team member's information search pattern), and reducing communication time (e.g., knowing who read what reduces the need for asking about it).

H2: Tools that provide peripheral, chronological awareness have a positive effect on asynchronous, distributed team performance in diagnostic, problem-solving tasks.

METHOD

To test these hypotheses, 60 participants were recruited from the Pittsburgh community, including students from the University of Pittsburgh and Carnegie Mellon University, as well as some non-student Pittsburgh residents. Participants were randomly grouped into 20 teams of three people each. Half of the teams were provided with the awareness tool.

Team members worked on the diagnostic problem at the same time but were randomly assigned to a lab at either the University of Pittsburgh or Carnegie Mellon University. Participants were identified to each other by pseudonyms and had no way of knowing who their teammates were or where they were located. Participants were trained on how to use the system and tools and then were given one hour to work on the problem.

Even though team members worked on the problem during the same time period, steps were taken to make the interaction asynchronous. These include not guaranteeing that all team members would start or end at the same time, periodically asking participants to stop working and write down their current thoughts about the problem, and having them fill out a survey halfway through the session. In addition, the interaction was made more asynchronous by inherent latencies within the CoMMIT system, as well as the substantial amount of individual text reading and notepad writing required of participants.

Participants were paid $15 for participating in the experiment. As additional motivation, two $150 prizes (one prize for each experimental condition) were awarded to the teams that came up with the best solution to the problem. To make the task more realistic and discourage haphazard document reading, participants were told that minimizing document requests and document denials would improve their chances of winning the $150 prize.

The Problem: How to Treat a Cancer Patient

Participants were told that all of them were medical doctors with the task of determining how to treat a patient with stomach cancer. The problem was created such that no medical knowledge was required to solve the problem. In fact, the problem is a reformulated version of Dunker's radiation problem [8], which is widely used with cognitive science students. The original problem involves the eradication of a malignant tumor without killing the patient, but our reformulated version extended the problem by introducing issues such as medical insurance, side effects, and family health history.

Even though the solution was contained in the information provided to participants, teams had difficulty finding the correct solution within the allotted time. Furthermore, a reading of the participants' chat text reveals that even when a member of a group found the correct solution and suggested it to the group, the other group members often did not agree. Only one group agreed upon the correct solution after substantial discussion.

RESULTS

Where applicable, most statistical tests were conducted at both the individual and group levels of analysis. Aside from usage data gathered by the CoMMIT system, data were also collected via two surveys. The first survey was given after participants had worked on the problem for approximately 30 minutes (session 1). After completing the first survey, participants were given 30 more minutes to work on the problem (session 2). At the conclusion of session 2, a second survey was given.

H1 Results: Convergence of Lines of Reasoning

Convergence of lines of reasoning was defined for our purposes as the degree of team agreement on the final solution. Except where noted otherwise, one-way ANOVA results are reported. The main variable studied captures solution agreement by team members from the second survey. On that survey, after recording the team's solution, participants were asked to rate the degree to which they agreed with the solution on a five point scale (1=strongly agree, 5=strongly disagree).

We did not find a significant difference between conditions with respect to the mean agreement level. The average in both conditions centered around 1.5 indicating a good deal of agreement in both conditions. However, as illustrated in the box plot in Figure 3, an F-ratio test revealed that the teams in the tool condition had significantly less variance in solution agreement ($p<0.01$, individual level; $p=0.25$, team level). This provides some evidence that teams with the awareness tool have a narrower range of solution agreement ratings.

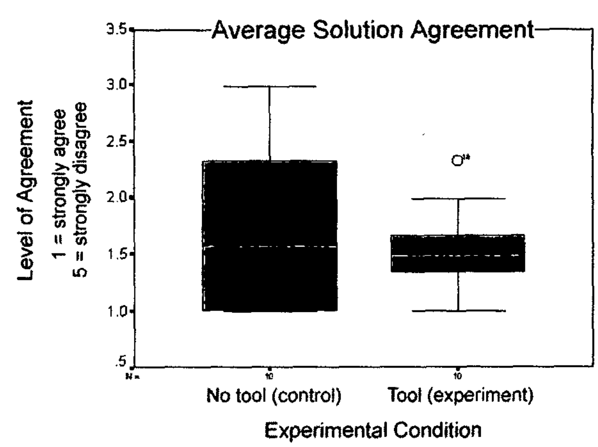

Figure 3: Box plot showing how much individuals agreed with their group's solution to the problem.

Agreement with Solution		Number of People: Without Tool	With Tool	Total
Strongly Agree	1	15	18	33
	2	10	6	16
	3	0	5	5
	4	0	0	0
Strongly Disagree	5	3	0	3
Total		28	29	57

Table 1: Degree of individual agreement with group solution.

In other words, teams in the no-tool condition have more extreme levels of disagreement. This is apparent from the chi-square analysis in Table 1. A chi-square test of independence at the individual level provided evidence that there is a significant difference in the pattern of solution agreement between conditions (χ^2=9.258, p=0.026).

As Table 1 illustrates, more individuals in the tool condition highly agreed with their team solution, while more teams in the no-tool condition highly disagreed with the team solution. No individuals in the tool condition either disagreed or highly disagreed with their team solution. Because we observed that some participants in the tool group did not use the awareness tool frequently, we split individuals in the tool condition at the median into high and low tool users based on how much time they had the tool active on their screen. We found that the high tool usage group had significantly higher levels of agreement (p=0.016) than the low tool users and moderately higher levels of agreement than everyone else (p=0.103). However, these effects were not significant at the team level, providing some evidence that the negative effects of low tool usage by some individuals within the team may offset the benefits of high tool usage by others. Thus, high tool usage by all team members is important when seeking a unified team solution.

In sum, the results above provide adequate support for the first hypothesis: teams that used the awareness tool more had a higher level of convergence toward a team solution. Furthermore, high tool usage accentuates this effect and increases the level of team agreement, provided that most team members are high tool users.

H2 Results: Team Performance

To test the second hypothesis, two types of performance measures were established. The first measure dealt with the correctness of a team's final solution; the second measure dealt with how efficiently the team searched for the solution.

Correctness of Solution

How close a team came to finding the correct solution was measured in two ways. First, at the beginning of the second survey, all participants had to type their group's solution. No team in the tool condition and only one team in the no-tool condition figured out the right solution. This was expected to some extent since the problem had been purposely developed to be difficult to resolve in the time provided. Even the one team that found the right solution had a substantial amount of debate before reaching an agreement. Members of two other teams discussed the right solution but did not select it.

Second, also in the second survey, participants had to rate several possible hypotheses on a five point scale indicating the degree to which they believed the hypotheses were correct. (Most of the analysis for the second hypothesis was done using one-way analysis of variance. Similar results were obtained using two-way analyses of variance and ordinary least squares regression models using a number of control factors. Thus one-way ANOVA results are provided here, except where noted otherwise.) Surprisingly, analyzing these data we found that individuals and teams in the no-tool condition were significantly closer to the right solution (p=0.001, individual level; p=0.17, group level).

We suspected these results could be due to the process constraints imposed by the system and the experimental setting. The system denies documents when not requested in a logical order, while the experimental setting rewards minimal denials and minimal document requests. Thus, an alternative explanation for our results may have been that the awareness tool helps teams cope more effectively with process constraints by providing awareness information about these constraints, at the expense of not focusing as much on the information provided to resolve the case. However, we tested the effect of tool use on denials and document requests and did not find a significant effect.

Search Efficiency

The second type of performance measure was intended to evaluate the team's solution search efficiency in terms of document requests, document denials, and division of labor. Chronbach's α reliability scores were computed for each team on the number of times each document had been requested by each team member. This was done to find the similarity of search patterns and used as a proxy for division of labor within the team. The teams in the no-tool condition had significantly higher reliability scores (p=0.012) indicating that teams in the tool condition had less overlap in their requests for documents, suggesting a higher level of division of labor. Further analysis of document requests, number of document denials, and number of entries into comment notepads revealed no significant differences between the two conditions.

We speculated that these results could be due to the fact that some individuals in the tool condition did not actively use the tool. Thus, once again we split the tool group into high and low tool users. One interesting result found was that high tool users made more entries in the notepads (p=0.02) but used fewer words (p=0.03) than low tool users, providing some evidence that high tool users interact more synchronously. Although not highly significant, a moderate interaction effect between high tool usage and session was found (p=0.151), suggesting that high tool users may become more synchronous than low tool users as the deadline for solution submission approaches.

In sum, these results provide some evidence that teams in the tool condition are somewhat more efficient in their solution search by better dividing labor. Teams with high tool usage were also more efficient by interacting more frequently and synchronously, while using fewer words. Other than these moderate results, the data did not adequately support the second hypothesis. Furthermore, teams in the no-tool condition were more likely to select the right solution as a plausible solution.

Additional Results

We conducted further analysis to investigate the effects of the awareness tool on team processes. Several team process questionnaire items on team satisfaction, strategy strength, team spirit, and team communication were reduced using the Principal Components method of factor analysis. We found no significant tool effects on any of these factors.

Further tests were then conducted to evaluate the effect of the tool on the team interaction process using objective data collected from notepad entries. Two-way analysis of variance models were formulated to test differences between conditions and between the two sessions in the average number of notepad entries made by each team and the average number of words used per notepad entry. The respective differences in variance were also tested for these two variables to evaluate the effect of the awareness tool on homogeneity of interaction within the team.

As Figure 4a (top) illustrates, teams in the tool condition made fewer entries on the notepads, but this difference was not significant (p=0.210). Overall, teams made more entries in the second session (p=0.071). However, as Figure 4b (middle) illustrates, when the same test was done between high tool user groups against all other groups, the difference between groups became moderately significant (p=0.088). Finally, as illustrated in Figure 4c (bottom), teams in the tool condition exhibited more evenness in the use of notepads, but this difference was not significant (p=0.129). As the plot illustrates, the difference in evenness in notepad use between conditions becomes more marked in the second session. In fact, teams in the tool condition tended to become more even, while teams in the no-tool condition tend to become less even towards the end. Interestingly, there is no difference in evenness of notepad use between high tool usage teams and low tool usage teams. Also, the team variance in the number of words per notepad entry tends to be larger for teams in the tool condition than for teams in the no-tool condition, but this difference is not significant (p=0.226). It is interesting to note that teams in the tool condition become more even in the number of entries in notepads, but more variable in the number of words per notepad entry, whereas the opposite happens with teams in the no-tool condition.

In order to gain a deeper understanding of how teams communicated and processed information, we carefully reviewed all chat text entered by teammates in the notepads, which was the only communication channel available. This qualitative review of notepad interaction revealed that teams without the awareness tool did a more thorough job of inspecting and discussing substantive issues regarding document contents, while teams with the awareness tool seemed to be more concerned with discussions of process to help them converge towards a solution. This is consistent

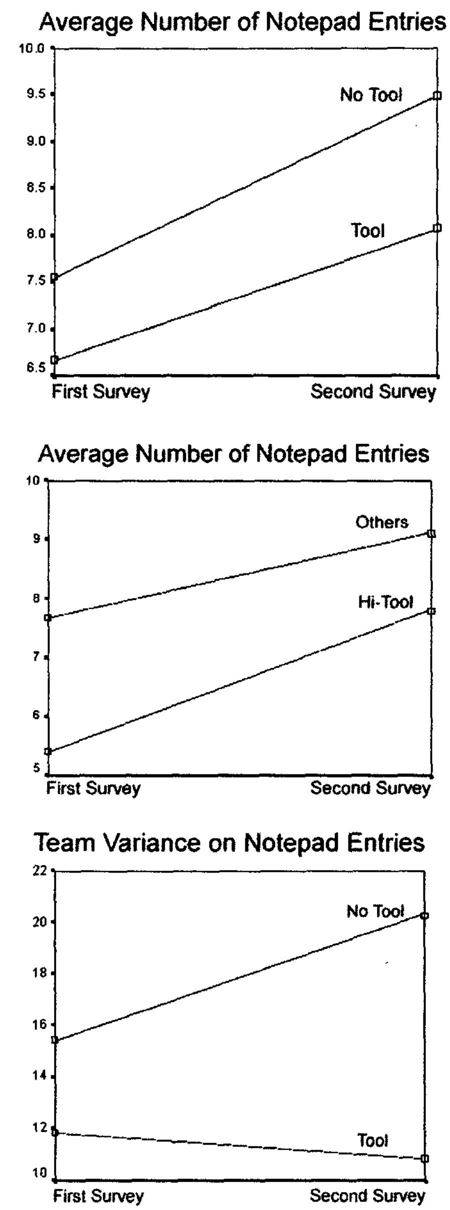

Figure 4: A) Interaction Plot—Notepad Entries Against Tool Condition by Survey (top). B) Interaction Plots—Notepad Entries Against High Tool Usage by Survey (middle). C) Team Variance on Notepad Entries Against Tool by Survey (bottom).

with the fact that teams in the no-tool condition had more overlap in the documents they read, which is perhaps what helped them have more common ground to do a more thorough job of discussing substance. The fact that teams in the tool condition had less overlap in the documents they read means that they were likely to have less common ground for discussions about substance, thus making them focus more on discussions about process.

DISCUSSION

Although the results obtained provide some encouraging evidence about the benefits of awareness tool use, they also make evident how the availability of such tools can be more of a distraction when available but not properly used. While peripheral, chronological awareness seems to help teams converge their lines of reasoning and reach joint team solutions more efficiently, it seems that this comes at the expense of a less thorough review of available information. Similarly, while the awareness tool seems to have contributed to a more efficient division of labor, the resulting reduced overlap in documents read by team members seems to result in a loss of common ground, thus foregoing the benefits of shared mental model formation. This is consistent with the fact that teams in the no-tool condition were less efficient in reaching consensus about the joint solution, but got closer to the right solution than teams in the tool condition.

Similarly, it seems that simply having the awareness tool without using it frequently may make some of these negative effects materialize more often, without capitalizing on some of the benefits of the tool. Furthermore, the negative effects of low tool use may offset some of the benefits of high tool use within a team when only one or two members use the tool actively. Consequently, although it is quite possible that the benefits of high tool usage may be attributed to a self-selection bias, it seems clear that such benefits can only materialize when all members use the tool actively and effectively. This effect is similar to the negative externalities commonly observed in other groupware systems like group schedulers and discussion databases, in which if one or two members do not use a given tool, its benefits are greatly diminished for all [21, 22].

Also, consistent with the literature on groups, it is evident from our results that awareness tools need to be matched to appropriate tasks [18, 19]. The primary focus of our awareness tool was helping teams to solve a problem quickly. This is precisely what the tool did in our experiment, but this benefit came at the expense of people not inspecting enough documents to allow them to come closer to the correct solution. However, not all problems have verifiable, correct solutions. Divergent problems, as discussed by Senge [24], do not have a single correct solution. With divergent problems, coming to a strong consensus about the course of action to take is more important than finding a "correct" solution, since no such solution exists.

Therefore, the features implemented in our awareness tool are adequate for a divergent problem in which there is no apparent right solution, and in which reaching a unified team solution is important. Strategic planning, sports team strategies, surgical teams in the operating room, and economic planning committees are examples of situations in which awareness tools of this type can help. However, in order to provide support for problems in which a correct solution does exist, different types of awareness information would have to be presented to the user. This highlights the all familiar tradeoff between general awareness tools that provide a little help for many types of tasks, and specific awareness tools that significantly help only one type of task. It also highlights the need to find the optimal amount and type of awareness information to make available without creating unnecessary distractions and information overload.

CONCLUSION

Perhaps the most useful lesson learned from this experiment is that although task awareness can be very beneficial to team performance, it may actually be detrimental to the team if the task awareness information provided is not properly matched to the needs of the specific task. Thus, in addition to the obvious conclusion that our tool could have been more helpful had the design matched the task, we believe our tool could also be more helpful in the following conditions:

1) Truly asynchronous conditions: although an asynchronous environment was simulated, we believe our awareness tool would be more effective had the teams interacted asynchronously over a longer period of time. Results of an earlier pilot test where interactions were truly asynchronous suggest that task awareness may yield more benefits under more asynchronous conditions. For example, we noticed the awareness tool was very useful in detecting social loafing, which was not as helpful in the shorter time period of our lab experiment.
2) More dynamic tasks: the literature on situation awareness suggests that general awareness is more critical under dynamic conditions. Our task and tool could be reformulated such that the patient's health deteriorates over time and each piece of information available has a cost to the team. The task used for the experiment was static, thus making situation awareness less important.

ACKNOWLEDGMENTS

This material is based on work supported by the Defense Advanced Research Projects Agency and administered by the Office of Naval Research under agreement number N66001-96-C-8506. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

REFERENCES

1. Adams, M., Tenney, Y., and Pew, R. Situation Awareness and the Cognitive Management of Complex Systems. *Human Factors*, 37(1), 1995.
2. Barley, S. The Alignment of Technology and Structure Through Roles and Networks. *Administrative Science Quarterly*, 1990, 61-103.
3. Barley, S. Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments. *Administrative Science Quarterly*, 1986, 78-108.
4. Cannon-Bowers, J., Salas, E., and Converse, S. Shared mental models in expert decision-making teams. In N. J. Castellan, Jr. (Ed.), *Current Issues in Individual and Group Decision Making*. Erlbaum: Hillsdale, NJ, 1993, 221-246.
5. Chen, L., and Gaines, B. A Cyber-Organism Model for Awareness in Collaborative Communities on the Internet. *International Journal of Intelligent Systems*, 12(1), 1997.
6. DeSanctis, G., and Poole, M. Capturing the Complexity in Advanced Technology Use: Adaptive Structuration Theory. *Organization Science*, May 1994.
7. Dourish, P. and Bly, S. Portholes: Supporting Awareness in a Distributed Work Group. *Proceedings of the ACM Conference on Human Factors in Computing Systems* (INTERCHI '92).
8. Dunker, K. On Problem Solving. *Psychological Monographs*, 58 (Whole No. 270), 1945.
9. Endsley, M. Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors*, 37(1), 1995.
10. Fuchs L., Pankoke-Babatz U., and Prinz W. Supporting Cooperative Awareness with Local Event Mechanisms: The Groupdesk System. *Proceedings of the 4th European Conference on Computer Supported Cooperative Work* (ECSCW '95), 247-262.
11. Fussell, S., Kraut, R., Lerch, F, Scherlis, W., McNally, M., and Cadiz, J. Coordination, Overload and Team Performance: Effects of Team Communication Strategies. *Proceedings of the 1998 Conference on Computer Supported Cooperative Work* (CSCW '98), 275-284.
12. Gaver, W. The Affordance of Media Spaces for Collaboration. *Proceedings of the 1992 ACM Conference on Computer Supported Collaborative Work* (CSCW '92).
13. Greenberg S., Gutwin C., and Cockburn, A. Using Distortion-Oriented Displays to Support Workspace Awareness. Technical report, Dept of Comp. Science, Univ. of Calgary, Canada, Jan. 1996.
14. Gutwin, C., Stark, G., and Greenberg, S. Support for Workspace Awareness in Educational GroupWare. *Proceedings of the 1995 Conference on Computer Supported Cooperative Learning* (CSCL '95).
15. Hackman, R. The Design of Work Teams. In Lorsch, J. (Ed.), *Handbook of Organizational Behavior*, Prentice-Hall, 1987.
16. Lautenbacher, G., Campbell, J., Sorrows, B., and Mahling, D. Supporting Collaborative, Problem-Based Learning Through Information Systems Technology. *IEEE Conference on Frontiers in Education*, 1997.
17. Lautenbacher, G., and Mahling D. Interface Design Issues for Web-Based Collaborative Learning Systems. *Proceedings of the 1997 WebNet Conference*.
18. McGrath, J. *Groups: Interaction and Performance*, Prentice Hall: Englewood Cliffs, N.J., 1984.
19. McGrath, J. and Hollingshead, A. *Groups Interacting With Technology*, Sage Publications: Thousand Oaks, California, 1994.
20. Orasanu, J. and Salas, E. Team Decision Making in Complex Environments. In Klein, G., Orasanu, J., and Calderwood, R. (Eds.), *Decision Making in Action: Models and Methods*, Ablex Publishing Co: Norwood, NJ, 1993.
21. Orlikowski, W. Learning From Notes: Organizational Issues in Groupware Implementation. *Proceedings of the 1992 Conference on Computer-Supported Cooperative Work* (CSCW '92).
22. Orlikowski, W. Improvising Organizational Transformation Over Time: A Situated Change Perspective. *Information Systems Research*, 1996, 63-92.
23. Schlichter, J., Koch, M., and Bürger, M. Workspace Awareness for Distributed Teams. *Proceedings of the Workshop Coordination Technology for Collaborative Applications*, Singapore, 1997.
24. Senge, Peter. *The Fifth Discipline*. Doubleday: New York, 1990.
25. Sproull, L. and Kiesler, S. *Connections: New Ways of Working in the Networked Organization*. MIT Press: Cambridge, Massachusetts, 1991.
26. Wellens, R. Group Situation Awareness and Distributed Decision-Making: From Military to Civilian Applications. In Castellan, J. (Ed.), *Individual and Group Decision Making: Current Issues*, LEA Publishers, 1993.