Improving Group Attention: An Experiment with Synchronous Brainstorming

Antonio Ferreira[1], **Pedro Antunes** [1], **Valeria Herskovic**[2]

[1] University of Lisboa, Bloco C6, Piso 3, Campo Grande, 1700 Lisboa, Portugal

[2] Universidad de Chile, Avenida Blanco Encalada 2120, Santiago, CP 837-0459, Chile

Abstract

In this paper we address the problem of information overload in synchronous group work: the large quantity of information, multiple information sources, and the need to sustain reciprocal interdependence have a negative impact on the capacity to attend to the group. We propose a group attention model characterizing the dynamic coupling between the group members and the mediating technology. Based on that model, we developed a compensation mechanism capable to estimate the most adequate time to raise the users' attention to the group. We describe how this compensation mechanism was applied to synchronous brainstorming and present results from a laboratory experiment. The obtained results indicate that groups using the compensation mechanism produced 9.6% more ideas when compared to the control groups. A detailed post-hoc analysis of the data obtained in the experiment also indicates that users using the compensation mechanism had 7.5 seconds of extra uninterrupted time to think about and type an idea, which they began to write 6.4 seconds sooner, and completed in 4.2 seconds less time.

1. Introduction

Information overload is a challenging problem in our information-rich world. In 2003 it was estimated the production of 23 Exabytes of *new* information, of which about 530 Petabytes were accounted to Internet contents, including Web pages, e-mail messages and persistent conversations (Lyman and Varian 2003). Such an immense quantity of information, of which only a very small fraction is ever needed, makes heavy demands on human and organizational processing capabilities.

We find accumulated evidence that information overload causes multiple problems to individuals, including the buildup of stress and confusion, and a greater propensity for human error (Reason 1990; Dekker 2006, Eppler and Mengis 2006). And we also find evidence of the same problem occurring at the organizational level (Simon 1971; Reason 1997). For instance, sense-making (Weick 1995) and situation awareness (Endsley 2003) are two examples of organizational-level cognitive functions that may be negatively influenced by information overload.

As it always happens with wicked problems, information overload may at the same time be a cause and an effect. And it also may be associated with many other phenomena, according to the stance one assumes regarding the problem. We may for instance regard

Ferreira, A., P. Antunes and V. Herskovic (2011) "Improving Group Attention: An Experiment with Synchronous Brainstorming." Group Decision and Negotiation, online first. doi:10.1007/s10726-011-9233-y.

information overload as being caused by the rise of information quantity, ambiguity and uncertainty, task complexity and interdependencies, new technologies, shortened life-cycles, too many standards, simultaneous inputs of information, in diverse scopes such as optimality of decision-making, search strategies and extent of information-retrieval, sense of control in communication processes, information coherence in knowledge management, and disposition to learning, to name just a few (Eppler and Mengis 2006).

Researchers have pointed out that technology and organizations have been evolving towards increasing complexity (Perrow 1999). The trend shows that organizations are moving from mechanization, to automation, and then to joint cognitive systems (Hollnagel and Woods 2005), accentuating the importance of groups and collaboration. We expect this trend to continue in the next years, which will certainly emphasize the importance of studying the impact of information overload in group work.

Dealing with information overload at the group level is challenging. On the one hand, we have to consider the group members may establish and maintain multiple communication channels. The number of communication channels and overall information exchange will therefore increase with the group size, which may lead to information overload. This perspective has been used to explain the well-known "mythical man month" paradox, where bringing more people to a delayed project will inevitably delay it further, since people will spend more time communicating (Brooks 1975).

Furthermore, discussing group work also leads to a discussion on task interdependence, that is, the extent the group members depend on each other to accomplish the task. According to Van de Ven and Delbecq (1976), we may find three increasing levels of interdependence: pooled, sequential and reciprocal. In the pooled situation, the participants work in parallel tasks and do not actually depend on each other. In the sequential case, the tasks are done sequentially and thus some group members rely on the others to timely finish their tasks. The reciprocal case is the most demanding one, as work becomes a joint effort where all the members depend mutually on each other to accomplish the task. Regarding information exchange, we observe that reciprocal interdependence is the most demanding type of group work and the one that most clearly leads to information overload: the group participants have to attend to multiple information flows in a timely, cohesive and often projected way.

The discussion on interdependence brings the distinctions between synchronous and asynchronous collaboration (Tung and Turban 1998). In asynchronous collaboration the group members are not dependent on time to interact with the group. In this context information overload is mostly related with the cognitive capability to process and make sense of information, finding cues, analyzing meaningful interpretations and enacting consequent activities (Weick 2001). On the contrary, synchronous collaboration is sustained on timely interaction. The outcome is that information overload is more related with time-related phenomena such as attention failures, interruptions and distractions, time pressure and mental workload (Reason 2008).

This paper studies information overload (1) at the reciprocal task interdependence level and (2) concerning synchronous collaboration. So we are addressing the most complex and demanding working conditions. Furthermore, within all possible time-related phenomena, we will only consider the phenomenon of (3) attending to the group.

In Figure 1 we present a roadmap of our research. It is based on the assumption that a large quantity of information, multiple information sources and the need to sustain reciprocal interdependence will lead to an overload of synchronous information and consequently to decreased capacity to attend to the group and loss of group performance. The study presented in this paper seeks to understand if active compensation mechanisms are capable to recover the capacity to attend the group and therefore increase group performance.

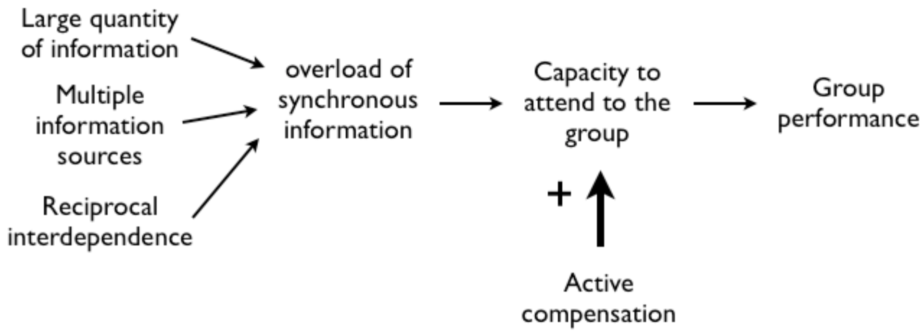

Figure 1: Research framework

The paper is organized as follows. We start by reviewing the related work, focusing on human attention and the use of computers to improve human attention. In Section 3 we propose our group attention model. Based on the proposed model, in Section 4 we discuss how to improve group attention in synchronous brainstorming using a software component named opportunity seeker. This section also describes the tool developed to try out the opportunity seeker. Section 5 is devoted to characterize the laboratory experiment. The obtained results are detailed in Section 6 and discussed in Section 7. Finally, in Section 8 we present our conclusions and point out future work.

2. Related work

Group attention is the main phenomenon studied in this paper. Considering our particular research context, we define group attention as: the capacity to attend to synchronous information delivered by the group members with the purpose of maintaining reciprocal interdependence. In the following we start by introducing attention in a general scope, moving then to an overview of technology support to group attention.

2.1. Human attention

Human attention is associated with the selection of relevant information and the attenuation or discard of non-relevant data. It is a cognitive process that optimizes the use of our cognitive resources so that we may perceive or act accurately and quickly (Sternberg 2003; Anderson 2005).

Over the decades, psychologists have been identifying the goals and limitations of human attention. The goals of attention are usually defined in terms of accuracy and speed responding, both contributing to decrease task execution times (LaBerge 1999).

Accuracy occurs when we successfully remove or attenuate the influences of extraneous and confusing information. For example, the "cocktail party" phenomenon, that is, the ability to keep track of a conversation in a crowded room (Conway, Cowan et al. 2001), is related with this goal. Technology may contribute to accuracy by filtering irrelevant or less important information. For instance, some popular electronic discussion forums display contributions according to the relevance attributed by the participants, graying out the less relevant ones.

Speed responding manifests when we are able to respond faster to anticipated events and almost always involves a clear expectation of when to initiate the response. Technology may improve speed responding by detecting and interpreting the user's activities. For instance, popular Internet instant messengers provide anticipatory cues when users start typing a message, preparing the other users to the imminent arrival of new messages.

Human attention is limited. In fact, some limitations are known to occur even after tremendous training and strict safety rules (Perrow 1999). One such limitation is the "attentional blink" phenomenon: a delay between paying attention to one object or task and attending to the next stimulus (Dehaene, Changeux et al. 2006). Still, there is evidence that the response time to a stimulus may be reduced if the time between attention switches is made longer and, in particular, constant (Roda and Thomas 2006). Groups may have difficulties dealing with attentional blinks, since parallel synchronous tasks may generate a high rate of unexpected cues.

Another attentional phenomenon is "change blindness," which manifests itself when we fail to notice changes, even dramatic ones (Dehaene, Changeux et al. 2006). As long as the change matches the context, for instance, swapping persons during a brief conversation, we may simply miss the change (Anderson 2005). If we do want to check if anything has changed, then we may have to engage in a very slow process of scanning the surrounding context. This happens because, although we can attend to four or five objects simultaneously, we may only detect one change at a time (Simons and Rensink 2005).

The effect of change blindness in groups is quite pertinent, since groups share a common context; and the presence of several contributors stimulates scenarios in which multiple changes may occur simultaneously. This creates the conditions for people not noticing changes, which may affect the ability to catch up with the group.

To compensate for change blindness, technology may highlight changes so they occur less frequently and are easier to notice. One example of this approach is the use of a bold typeface for new messages displayed in e-mail inboxes.

2.2. Attention support

The idea of using computers to improve human attention is the core research stream known as Attentive User Interfaces, or AUI (Vertegaal 2003; Roda and Thomas 2006). A prime motivation for AUI is the recognition that as the needs for information and communication rise so do the costs of not paying attention and being interrupted. So, instead of assuming that users take the full responsibility for managing their own attention, AUI afford to establish, optimize and negotiate the priorities for presenting information.

Most research on AUI has been experimenting with various devices capable of detecting human attention and informing the computer about human behavior. This includes using eye-trackers and body orientation, physiological, pupil dilatation, and heart rate sensors (Vertegaal 1999; Chen and Vertegaal 2004; Vertegaal, Shell et al. 2006).

Another research line concerns enhancing the human-computer interaction. The major goal is allowing the users to remain focused on their primary task without getting too much distracted by secondary, typically unrelated and unexpected tasks. For instance, Horvitz et al (2003) applied statistical models to estimate and balance the value of information against the cost of interrupting the user, based on cues about the working environment, including past activities, work patterns and even appointments in the personal calendar. Fogarty et al (2005) developed a method to estimate the best time to interrupt users using software development environments, based on salient tasks such as coding, scrolling code, debugging, and others. And Iqbal et al (2005) combined cues about mental workload with hierarchical task models to predict the most adequate time to interrupt the users.

2.3. Attention support in group work

Research on attention support to group work has mainly focused on the use of eye-trackers to facilitate the detection of who is talking in remote meetings. The GAZE system does this automatically with an eye-tracker placed in front of each user and uses photos to give cues about who is talking and who is listening (Vertegaal 1999).

In GAZE-2, the users' photos are substituted by live video feeds. In this system, each user's face is captured by three video cameras at different angles. An automated camera director chooses the video feeds based upon the amount of parallax error as determined by the eye-tracker (Vertegaal, Weevers et al. 2003). As in GAZE, the users' virtual representations are rotated to reflect their focus of attention, typically directed towards the current speaker. Moreover, as the angle of rotation increases, the quality of the video stream is intentionally reduced. In other words, full video quality is granted only to the user who is being fixated upon by the other users. According to the authors, this technique is based upon our natural limitations concerning peripheral vision.

Another GAZE-2 feature is automatically filtering voices when multiple conversations are being held at the same time. Depending upon the focus of attention, the audio streams are amplified and attenuated, but not completely eliminated.

The concepts studied in GAZE-2 were further explored in eyeView, which supports large remote meetings by manipulating the sizes of video windows and the voice volumes of each group member according to the their focus of attention (Jenkin, McGeachie et al. 2005).

So far the available research studies seem to primarily focus on evaluating the enhanced input/output devices. The study of their impact on the actual group work is yet to be done. Furthermore, these research studies continue to conceptualize work as a primary task often interrupted by events triggering secondary tasks necessary to attend to the other group members or to maintain awareness about what they are doing.

However, is has been found that groups, instead of doing a single extensive task with few interruptions, often execute a series of intertwined and constantly interrupted tasks (McDaniel, Olson et al. 1996). In this perspective, the distinctions between primary and secondary tasks are less defined and therefore their impact on attention should be further analyzed.

3. Group attention model

The proposed model may be considered a specialization of the Model Human Processor, or MHP (Card, Moran et al. 1983), offering a selective view over the cognitive processes necessary to pay attention to the group. The MHP has been widely adopted to describe human-computer interaction at the cognitive level. It represents human information processing based upon perceptual, motor, and cognitive processors, as well as working and long-term memories. The MHP assumes the user interacts with the computer according to the following cycle of events:

- The user acts upon the user-interface (UI), say, by moving the mouse with the hand;
- The UI responds with feedback information, for example, moving the mouse pointer (visual feedback);
- The user captures the information conveyed by the UI, for instance, recognizing where the mouse pointer is;
- The user interprets the information supplied by the UI and decides what to do next.

Feedback serves the fundamental purpose to dynamically couple human action with the computer response (Douglas and Kirkpatrick 1999). Feedback occurs almost constantly during human-computer interaction, to the point that users find it disturbing when the user-interface does not give immediate feedback.

Actually, feedback is not the only dynamically coupling element we find in human-computer interaction. Feedforward is another key element, having one important difference to feedback: it does not emerge as a response to the user's actions, but instead serves to stimulate human response (Wensveen, Djajadiningrat et al. 2004). The following cycle of events illustrates the differences to feedback:

- The user-interface delivers feedforward to the user, for example, showing that two buttons may be pressed;
- The user captures the information conveyed by the UI, for instance, recognizing the associated functions;
- The user interprets the information supplied by the UI and decides what to do next;
- The user acts upon the UI, for example, by pressing one button.

Regarding group work, we should extend the MHP characterization to accommodate the dynamic coupling between individual actions, computer mediation and group response. This may be described with the following cycle of events:

- One user acts upon the user-interface;
- The UI responds to the user with feedback information;
- Furthermore, the UI delivers feedthrough information to the other users, making them aware of the actions upon the UI done by the user;

- The group members interpret the information supplied by the UI and decide what to do next.

Feedthrough is the technology-mediated transmitting of cues to multiple users reporting on actions executed by one user. Feedthrough is necessary to provide group awareness and to construct meaningful contexts for collaboration (Hill and Gutwin 2003). Without feedthrough, users would be able to work in parallel and perhaps coordinate their activities, but would be unable to sustain the synchronous information flows necessary to build reciprocal interdependence.

A simple way to support feedthrough is by multiplexing feedback information to several users. But more interestingly, feedthrough does not have to exactly match feedback, since the computer may support various communication channels and control the timing, granularity and other attributes of the information being delivered. Various levels of feedthrough information may therefore be supported:

- Media feeds: includes the synchronous transmission of audio, video, gestures and body language, such as supported by audio and video conferencing systems;
- Low-level user-interface feeds: includes, for example, multiplexing keyboard and mouse movements. In the former case, a user would be able to see what others were typing, say, in a shared text document; in the latter, multiple mouse pointers would appear on the computer display, one per user (these are called telepointers);
- High-level user-interface feeds: delivering information about the manipulation of UI elements, such as buttons, menus, input data fields, and others, which augment the information conveyed by low-level user-interface feeds. For instance, instead of only seeing telepointers, users would be able to notice that a user is choosing a particular menu option;
- Application-level feeds: delivering only the events relevant to the application using filtering and aggregation mechanisms, for instance, by only delivering information that matches some user-defined criteria or by only showing complete sentences rather than individual keystrokes (as in popular instant messengers).

Based on the above analysis of the main elements dedicated to dynamically couple human action with the computer, we may finally propose a group attention model. See Figure 2 for details. First, we shall consider the model concerns a group of users interacting with and through a user-interface. All the interaction may be conceived according to an information processing lifecycle where users act upon the UI, receive a set of cues and continue acting upon the UI. According to this model, group attention concerns the perceptual response and interpretation of the cues delivered by the UI.

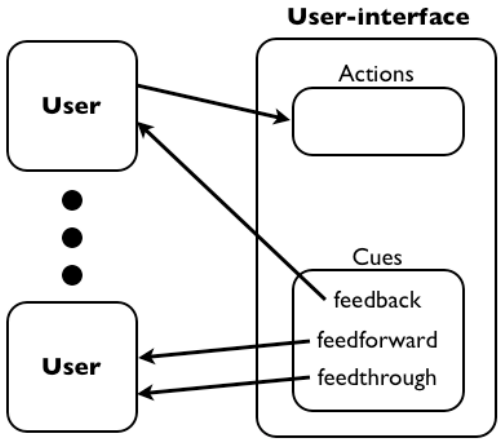

Figure 2: Group attention model

The set of cues that we have identified are:

- Feedback: cues about the actions invoked by the user in the user-interface;
- Feedforward: cues about action possibilities set by the UI;
- Feedthrough: cues about the actions invoked by other users in the UI, including media, user-interface and application feeds.

4. Improving group attention

Improving group attention deals with conciliating the needs to keep the users mindful of the group and reducing information overload. Based on the model described in the previous section, we may rephrase the main goal of improving group attention as: reducing feedthrough information while at the same time keeping the users mindful of the actions performed by the group.

This goal is accomplished with a compensation mechanism developed to control how the feedthrough information is delivered to the users. This compensation mechanism was implemented in a synchronous brainstorming tool specifically developed with the purpose to conduct laboratory experiments. We will start by discussing the basic assumptions behind the functionality of the compensation mechanism and then will discuss in detail the developed tool.

The compensation mechanism is designated *opportunity seeker*. The opportunity seeker handles exclusively the user-interface and application feeds. We do not currently address media feeds. The opportunity seeker collects the events associated to feedthrough in a buffer and automatically manages the timing and quantity of information delivered to each user based upon criteria that try to optimize the opportunity for raising attention.

Regarding this criteria, we note there is a trade-off in managing the timing of deliveries, in that too few may give the wrong impression about what the group is doing, damaging synchronous work and reciprocal interdependence, while too many deliveries may become distracting, reducing the capacity to effectively attend to the group. We address this trade-off by leveraging the alternation between two main activities: acting upon the user-interface (UI) and attending to feedthrough (see Figure 3).

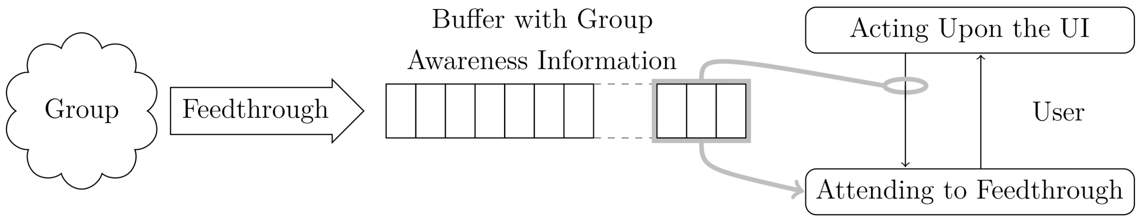

Figure 3: Conceptual view of the opportunity seeker, showing the feedthrough flows being directed to an internal buffer, whose contents are gradually delivered in batches to the user as s/he switches tasks. The opportunity seeker manages one buffer per user

Based on research done by Bailey and Konstan (2006), we consider the best opportunities for raising attention occur at the boundaries between consecutive tasks, such as editing a piece of text, moving an object, giving a comment, and others. As we mentioned earlier, group work tends to be fragmented in various small tasks rather than one longer task, thus the opportunity seeker should have plenty of occasions to deliver feedthrough information immediately after it detects the end of one of those small tasks.

This choice of providing feedthrough between consecutive tasks is well aligned with the results of a study about dual-task interference in group support systems, which showed that participants in a text discussion had more difficulties in processing new information because of the need to simultaneously contribute to the discussion (Henninger, Dennis et al 2006). Consequently, the authors proposed the introduction of formal stages so that the users stayed focused on one task at a time. We note, however, that this somewhat artificial adaptation of the initial task is not needed with the opportunity seeker as long as it is possible to automatically determine task boundaries.

Concerning the optimal quantity of feedthrough information delivered at once by the opportunity seeker, we could not find guidance in the research literature. The purpose is to avoid overloading the user with too much information and, therefore, a small batch should be defined. But a small batch also means that feedthrough information may stay longer in the buffer, which should be avoided. As we describe later, we had to empirically define the batch size.

Of course, we also have to consider what to do when the users pause doing work. In these situations, if there is undelivered feedthrough in the buffer and after a predefined timeout, then the opportunity seeker should deliver feedthrough (considering the batch size mentioned above). Our assumption is that raising the users' attention in these circumstances may contribute to sustain reciprocal interdependence.

4.1. Synchronous brainstorming

A synchronous brainstorming tool allows users producing as many ideas as possible in a short time period, using computer technology to support parallel activities in remote and local places. Brainstorming tools have been the subject of much experimental research (Gallupe, Dennis et al. 1992; Dennis and Valacich 1993; Gallupe, Cooper et al. 1994; Dennis and Williams 2003; Dennis and Reinicke 2004), which gives a good opportunity for benchmarking our own experiment.

The rules of brainstorming (Osborn 1963) encourage the participants to accomplish two cognitive tasks: production of as many ideas as possible, because quantity is wanted;

and cognitive stimulation by glimpsing the others' ideas. Regarding the production task, one of the positive effects of synchronous brainstorming is supporting parallel work. The users may develop ideas in parallel, which reduces production blocking and improves the group's productivity expressed by the raw quantity of ideas generated by the group (Gallupe, Bastianutti et al. 1991; Pinsonneault, Barki et al. 1999).

Cognitive stimulation is more challenging however: as the number of ideas increases, for example, because the group is inspired or the group size is large, the users may become distracted by the interruptions and ultimately be unable to divide their attention between producing ideas and glimpsing the others' ideas. If incorrectly implemented, the synchronous brainstorming tools may indeed cause information overload. This effect may explain why some experiments with brainstorming have been equivocal (Pinsonneault, Barki et al. 1999).

We developed the opportunity seeker exactly to compensate this type of information overload: as the number of ideas from others increases, the opportunity seeker will store them in a buffer until determining the best moment to raise the user's attention with feedthrough. The Attentive Brainstorming Tool (ABTool) was developed and instrumented to research the functionality of the opportunity seeker.

The major practical challenge regarding the ABTool development was adapting the theoretical concepts behind the opportunity seeker to the concrete aspects of synchronous brainstorming. As previously discussed, the opportunity seeker leverages the alternation between two main activities: (1) acting upon the user-interface, which we could now designate production; and (2) attending to feedthrough, which corresponds to stimulation in our brainstorming context. In order to alternate between production and stimulation, the opportunity seeker must determine the boundaries between two consecutive production tasks.

We adopted an empirical approach to determine these boundaries. We asked several groups of five volunteers to participate in synchronous brainstorming sessions using the ABTool with the opportunity seeker *inactive* (supplying immediate feedthrough). Beyond the interaction with ABTool, no other communication was allowed.

We then recorded three types of events: (a) user key presses while typing ideas; (b) the moments when the users submitted ideas to the group; and (c) the instants when feedthrough was delivered to the user's computer screens.

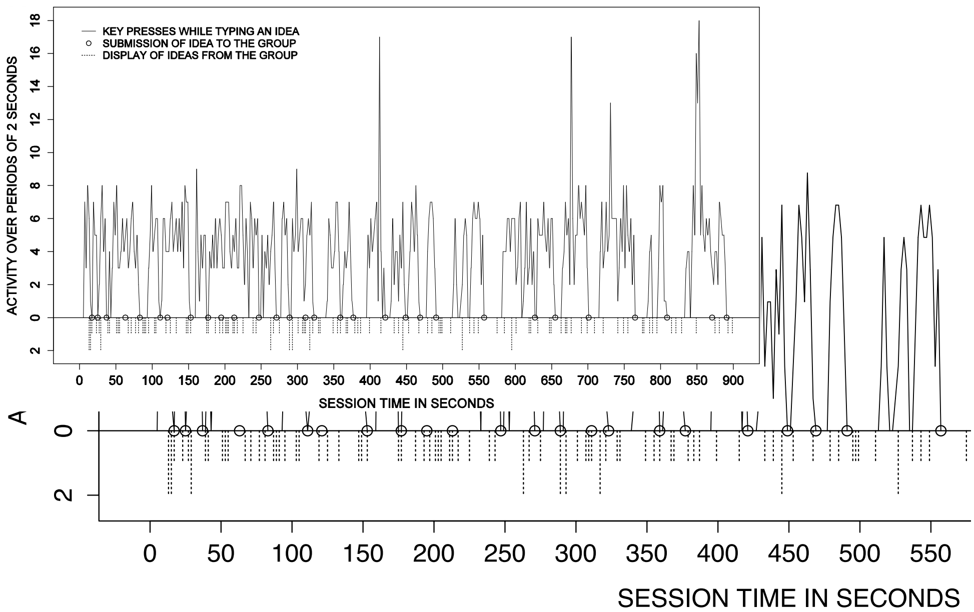

Figure 4: User and group activity during a brainstorming session

Figure 4 shows a sample of the data we obtained and illustrates the results from an entire 15-minutes session in which 152 ideas were produced. From the evidence we collected three patterns of user activity emerged: (1) users usually did not stop typing when they received feedthrough, thus, we assume they continued focused on the production task; (2) users typically paused after putting forward an idea, presumably to keep up with the group; and (3) there were numerous periods of time with no typing activity (not shown in Figure 4).

Based upon these patterns, we hypothesize that a task boundary occurs when the user submits an idea. The opportunity to get the user's attention to the group would therefore occur immediately after the user submits an idea.

In addition, since we observed several periods of typing inactivity, we confirmed the need to incorporate a timeout in the opportunity seeker, necessary to trigger the user's attention even if no task boundary was detected. Based on our empirical observations, the set period of inactivity was 10 seconds. Figure 5 shows the state transition diagram that models the user as assumed by the opportunity seeker in ABTool.

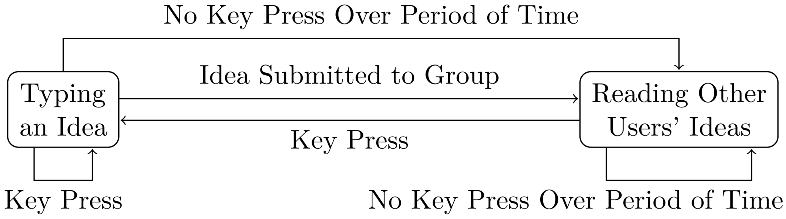

Figure 5: Model of user behavior adopted by the opportunity seeker in ABTool

Another feature of the opportunity seeker is that it imposes a limit on the number of ideas from others that can be displayed at once. Again based on our empirical observations, the limit was set to 10 ideas. This limit avoids overloading the user with too many stimuli. The empirical assumptions defined for the opportunity seeker will be further discussed with the presentation of our experimental results.

4.2. Software architecture and design

Technically, ABTool is characterized by a client-server architecture, in which the server mediates the information flows. The server also collects performance data, which is stored in an XML log. The purpose of the clients, one per user, is to receive input from the users and pass it on to the server, and to display new ideas as they become available from the server.

ABTool is written in C# and is based upon the Microsoft .NET Framework 2.0. Communication between the clients and the server is done via TCP/IP sockets and all messages (ideas, key presses, users joining or retiring the group, sessions starting or ending) are automatically serialized and de-serialized using BinaryFormatter objects attached to NetworkStream instances.

Within the client and server applications, messages are propagated using events, to which consumer objects will subscribe. Given that almost all data classes in ABTool handle message events, including the user-interface, opportunity seeker and other classes responsible for exchanging network messages, we specified the IHandlesMessages interface and the default implementation for it, DefaultHandlesMessages, which relies on reflection to allow those classes to delegate the determination of the method to run as a function of the type of message associated with the event.

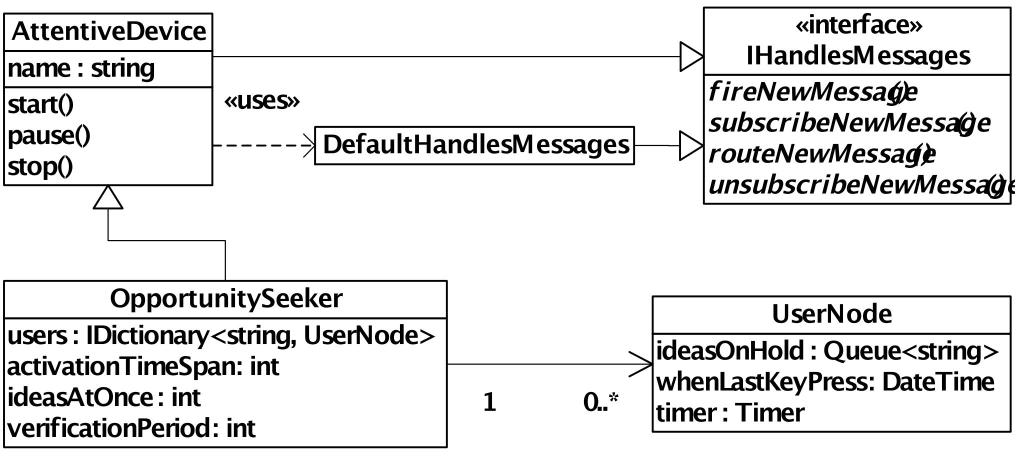

Figure 6: Class diagram of the opportunity seeker implemented in ABTool

Figure 6 shows that the opportunity seeker derives from the AttentiveDevice generalization, which actually implements immediate feedthrough of ideas. The OpportunitySeeker class alters this default behavior by maintaining separate buffers, one per user, containing ideas that have been put forward by the other users. The buffer is stored in the UserNode class, which also keeps a Timer object that every verificationPeriod milliseconds checks the time of the most recent key press by the user, and if it was more than activationTimeSpan milliseconds ago, then it delivers up to ideasAtOnce ideas to the user.

The AttentiveDevice and OpportunitySeeker classes implement three methods: start() is executed when a session starts or resumes; pause() is executed when, for some reason, the session needs to be paused; and stop() is executed at the end of a session. Other methods handle the reception and forwarding of messages, but we omitted those for brevity. We show in Figure 7 two screenshots of the client application with the opportunity seeker running.

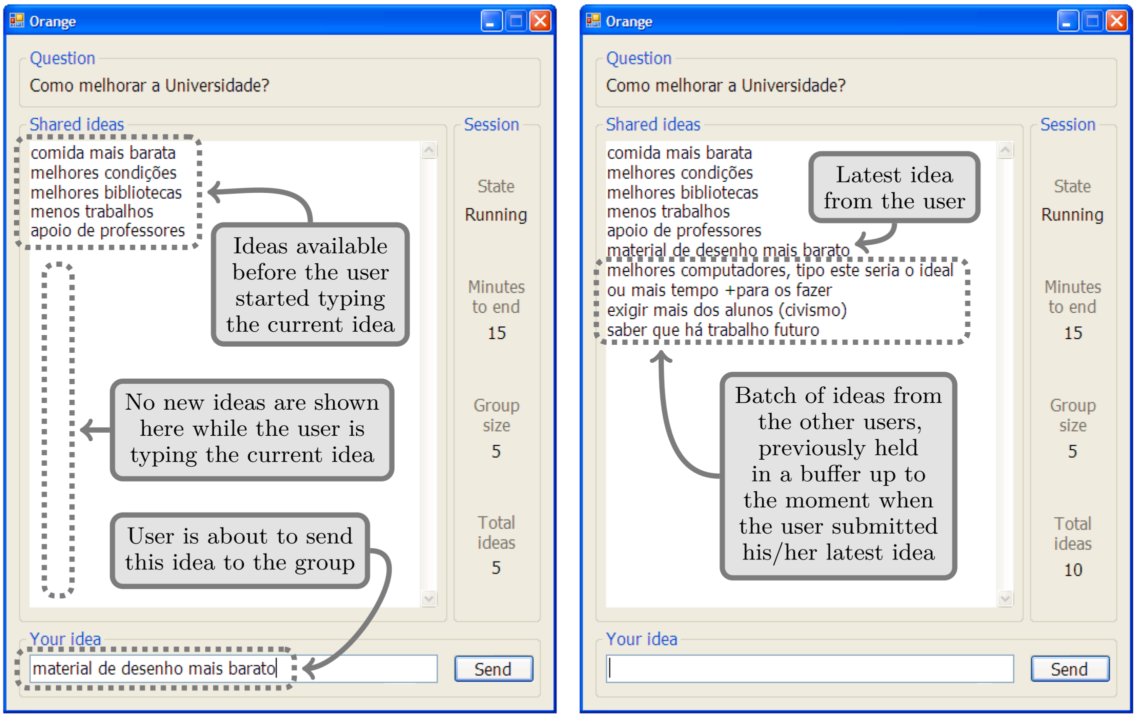

Figure 7: ABTool. Left: while typing an idea, the user receives no new ideas from the group. Right: when the user submits an idea, the others' ideas are displayed

5. Laboratory experiment

We now describe a laboratory experiment that we set up using ABTool to test the theory that the capacity to attend the group and the group performance improve when groups are exposed to the opportunity seeker.

5.1. Participants

A total of 11 groups of 5 people, for a total of 55 volunteers (44 men and 11 women) participated in the experiment. The median age was 23 years (min. 20 and max. 29). 51 participants were students (40 undergraduate, 10 M.Sc., 1 Ph.D.), and the remaining 4 were professionals. A convenience sampling was used to select the participants, who were recruited from social contacts and posters on corridors at the University of Lisbon. No monetary reward was offered and the only information available was that the experiment would concern brainstorming.

5.2. Apparatus

The experiment was conducted with 5 laptops having identical hardware (Intel Pentium M at 1.2 GHz, 1 GByte of RAM) and software specifications (Microsoft Windows XP SP2, .NET Framework 2.0), interconnected by a dedicated 100 Mbit/s Ethernet network. Keyboard sensitivity, desktop time display resolution, and brightness were controlled. Each computer had screen-recording software (ZD Soft Screen Recorder 1.4.3), and a web-camera (Creative WebCam Live!) affixed to the top of the screen. The client application of ABTool was installed on all laptops.

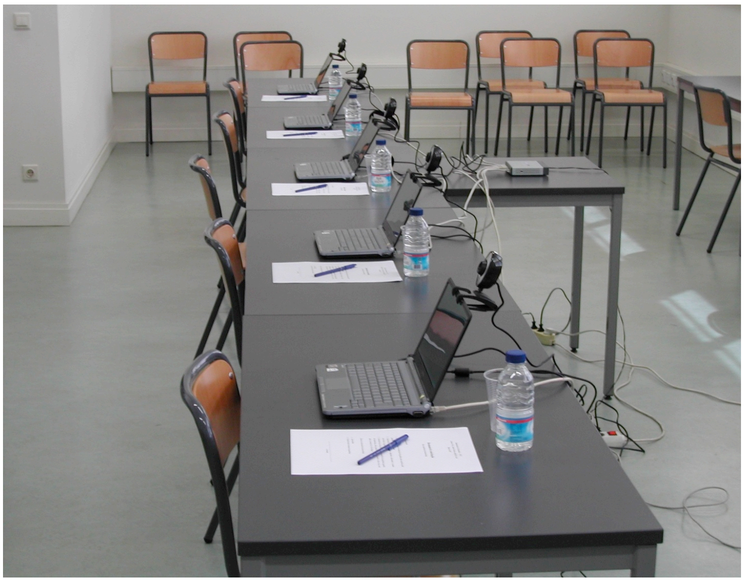

Figure 8: Laboratory room

Figure 8 shows that the layout of the laboratory room where the experiment was conducted comprised a single row of five work desks. We chose this configuration to minimize eye contact between the participants (to reduce distractions) while keeping staging costs manageable. After some training, a team of two was able to make all necessary preparations in about half an hour.

5.3. Tasks

The participants were subject to practice and test tasks, both related to brainstorming. The *practice* task allowed them to get familiar with ABTool. In the *test* task, the participants were given a question and then asked to generate as many ideas as possible by typing on the keyboard and looking at the computer display. Speech and other forms of communication were disallowed to mitigate extraneous influences.

Both practice and test brainstorming tasks were anonymous and allowed parallel input from all participants. In this way, we would be using technology to alleviate two well-known motivation losses, namely evaluation apprehension and production blocking, and creating the conditions for participants to generate more ideas (Hymes and Olson 1992; Connolly, Jessup et al. 1990). This, in turn, demands greater attentional resources from the group members, which is fundamental to this research.

5.4. Design

A repeated measures design was chosen for the experiment. The independent variable was *attention management* and every group of participants was under the influence of a control treatment, with no group attention management, and an experimental treatment, using the opportunity seeker to manage group attention.

Figure 9 illustrates the differences between the two treatments in a simulated brainstorming session, from which we highlight that under the control treatment the user immediately received all nine ideas generated by the group in contrast with three batch deliveries of ideas under the experimental treatment. Note that for illustration purposes we do not show the propagation of ideas 3, 11, and 12 to the group, and we also limit the number of ideas delivered at once to five instead of ten, which is the actual value in ABTool.

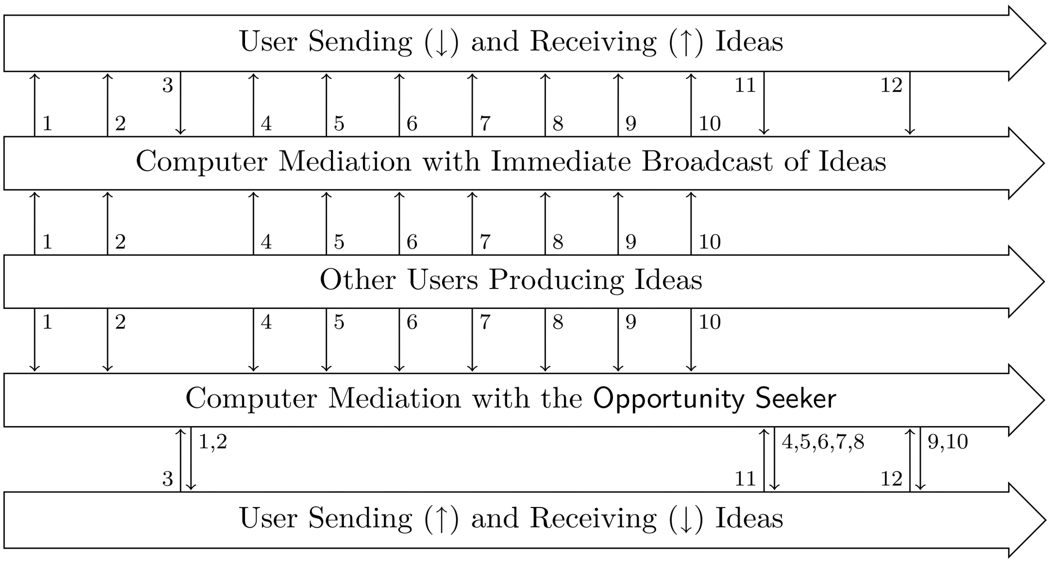

Figure 9: Simulation of group and user activity under the control (upper region) and experimental treatments (lower region, with the opportunity seeker). In both cases the user produces three ideas (numbered 3, 11, and 12) but the exposure to the nine ideas s/he received from the other users is different

The dependent variable, *group performance*, was calculated from the sum of the number of ideas produced by each user on the group per brainstorming session.

	Groups										
Treatment	1	2	3	4	5	6	7	8	9	10	11
Control	1/C	2/D	4/C	3/B	1/B	1/A	2/C	3/B	2/B	3/C	1/A
Experimental	3/B	1/A	2/B	4/C	3/C	2/B	3/A	1/C	1/C	2/A	3/B

Table 1: Session order/brainstorming question per group and treatment. The questions were: A - how to preserve the environment; B - how to promote tourism; C - how to improve the university; and D - how to stimulate sports practice

The order of exposure to the treatments and the brainstorming questions are depicted in Table 1, which shows, for example, that group 1 was under the influence of the control treatment in the first session, in which ideas for question C were put forward (this corresponds to the top-left cell marked 1/C). We note that sometimes the session order is greater than two and that four questions were used, because we are reporting here a part of a larger experiment with two additional treatments, involving similar brainstorming tasks.

5.5. Procedure

A trial started when a group of participants arrived at the laboratory room. An introduction to this research was given and the participants were informed on their privacy rights and asked to sign a consent form. Next, the participants completed an entrance questionnaire about gender, age, and occupation. Written instructions with the rules of brainstorming and ABTool functionality were then handed to all participants and read out loud by the experimenter.

The participants were asked to carry out the practice task for 5 minutes, after which an inquiry about the ABTool was answered. The group then performed the test tasks in succession, each lasting for 15 minutes, with a brief rest period in between. This session length was chosen to stress the importance of generating ideas at a fast pace for a relatively small duration, in which participants would need to remain attentive (interestingly, Dennis, Valacich et al (1996) suggest that time constraints increase the rate of idea generation), and also because we wanted to avoid fatigue induced by the repetitive measures design.

At the end of the trial, answers were given to any questions the participants had about this research, comments were annotated, and the experimenter gave thanks in acknowledgement of their participation in the experiment.

6. Experimental results

The results are organized in three parts: firstly, an analysis of the overall group performance, which is central to our theory; secondly, a decomposition of group performance into consecutive periods over a brainstorming session; and finally, results from a post-hoc analysis of group attention based upon more fine-grained data collected at the user level.

6.1. Group performance

As mentioned above, we used the total number of ideas produced by the group to measure group performance. In the following, the brainstorming sessions using the opportunity seeker will be designated either as experimental or OSB, while the control brainstorming sessions will also be designated CB.

	Groups													
Treatment	1	2	3	4	5	6	7	8	9	10	11	**Sum**	***M***	***SD***
Control	152	83	133	91	264	77	48	53	66	104	70	1141	103.7	62.0
Experimental	192	108	113	117	258	77	68	61	76	116	65	1251	113.7	60.8
Difference	40	25	−20	26	−6	0	20	8	10	12	−5	110	10.0	17.2

Table 2: Number of ideas produced by groups under the influence of the control (CB) and experimental (OSB) treatments

The results from our experiment indicate OSB produced an average of 9.6% more ideas than CB (OSB: *mean* = 113.7, *stdev* = 60.8; OSB: *mean* = 103.7, *stdev* = 62.0), as shown in Table 2. The Shapiro-Wilk normality test indicated that the normality assumption could not be accepted for both the control and experimental data distributions (W = 0.795, p = 0.008; and W = 0.797, p = 0.009, respectively). Therefore, we applied the non-parametric Wilcoxon signed-ranks test, which revealed a 3.7% probability of chance explaining the differences ($W_+ = 45.5$, $W_- = 9.5$).

We also analyzed possible confounding influences from the questions or session order on group performance to see if there was a bias introduced by popular questions or a learning effect due to the nature of the repeated measures design. We applied the Wilcoxon signed-ranks test to both scenarios, which found no significant influences ($p > 0.205$ and $p > 0.343$, respectively). Given this evidence, we can accept the hypothesis that group performance improved when groups were exposed to the opportunity seeker.

6.2. Group performance over time

Concerning the analysis of group performance through the duration of the brainstorming sessions, we broke down the 900 seconds that each session lasted into consecutive periods of 300, 150, and 30 seconds and counted the number of ideas put forward during each period. By following this approach we tried to identify any specific periods where the opportunity seeker could be more or less effective. For example, a brainstorming session may be decomposed in the beginning, when users usually have plenty of ideas, the middle, and the end, when users are typically more passive. So it would be possible the opportunity seeker be less useful in the beginning of the brainstorming session.

This decomposition is depicted in Figure 10. It shows that in all three periods of 300 seconds OSB produced more ideas than CB. We obtained similar results at the 150 seconds level of aggregation. Finally, if we consider periods of 30 seconds, OSB performed better in 21 out of 30 cases. From the evidence collected, there seems to be no particular phase when the opportunity seeker could be considered worse than the control case.

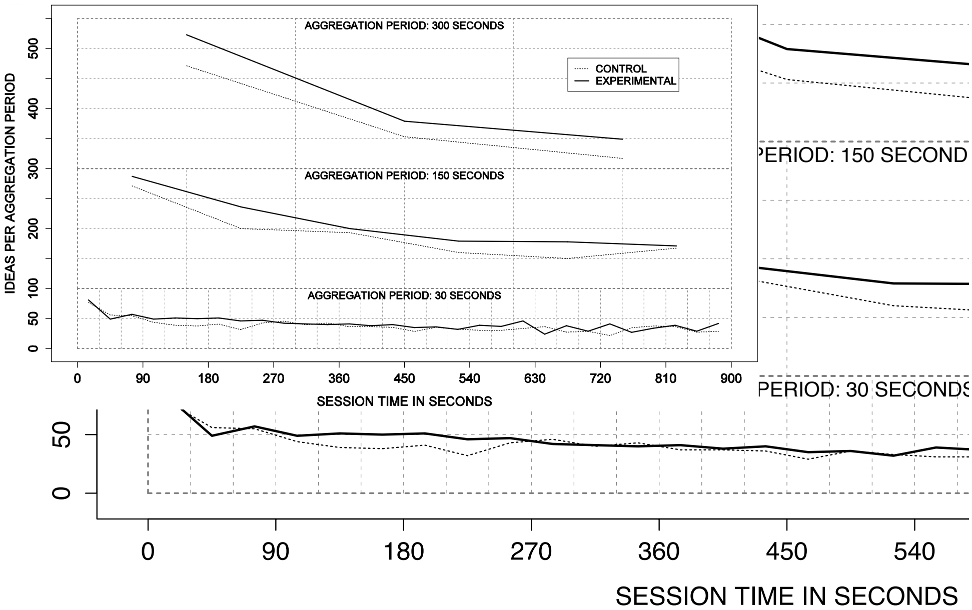

Figure 10: Group performance through the duration of the brainstorming sessions under the control and experimental treatments. Top: number of ideas per period of 300 seconds. Middle and bottom: same, considering periods of 150 and 30 seconds, respectively

6.3. Post-hoc analysis of group attention

We also performed a post-hoc analysis of group attention using fine-grained data collected with ABTool. Our main purpose was to analyze the users' capacity to attend to the others' ideas during the brainstorming sessions, with and without the group attention management provided by the opportunity seeker.

We considered the following variables: TBDL, seconds between consecutive deliveries; TIDEA, seconds to write an idea; PAUSE, seconds between a user submitting an idea to the group and restart typing; CIDEA, characters per idea; CHARS, total number of characters typed per user in a session; and DCHARS, total characters deleted per user per session.

	Control		**Experimental**		**Difference**		**Wilcoxon Test**		
Variable	*M*	*SD*	*M*	*SD*	*M*	*SD*	W_+	W_-	*p*
DLVR	82.7	48.1	46.2	14.6	−36.5	37.4	0	1540	0.000
TBDL	13.7	5.9	21.2	6.1	7.5	3.2	1540	0	0.000
TIDEA	25.7	17.3	21.5	11.8	−4.2	12.9	422	1118	0.004
PAUSE	34.1	34.3	27.7	19.2	−6.4	21.7	469	1071	0.012
CHARS	1044.8	511.2	1110.4	529.8	65.6	321.4	936.5	603.5	0.164
CIDEA	45.6	12.7	43.9	12.9	−1.7	9.5	613	872	0.266
DCHARS	206.7	163.0	199.3	133.3	−7.4	121.9	724.0	816.0	0.703

Table 3: Results of post-hoc analysis at the user level, ordered by relevance using *p*-value

Table 3 shows a summary of the results we obtained at the user level, including separate descriptive statistics for the two cases in which the users were under the exposure of the control and experimental treatments (creating two large conceptual groups, without and with the opportunity seeker), as well as the output of the Wilcoxon signed-ranks test, which we use here to prioritize the data presentation rather than to do null hypotheses significance testing. Thus, no family-wise corrections were made.

Starting with the DLVR variable, the opportunity seeker reduced feedthrough, that is, the number of deliveries of group ideas that reached a user per session, by 44.1%. The difference from a mean value of 82.7 deliveries per session to 46.2 was due to each delivery having comprised a batch of 1.9 ideas on average (*stdev* = 1.2), with up to 5 ideas in 99% of the cases and a maximum batch size of 9 ideas (happening only once), unlike the control treatment, in which new ideas were immediately broadcasted, one by one, to the group.

Another consequence of the opportunity seeker captured in variable TBDL is that users had 54.7% more time, on average, to think about and type ideas without being interrupted. This corresponds to uninterrupted periods with a mean duration of 21.2 seconds, instead of 13.7 seconds in the control case.

The opportunity seeker trades immediate attention to one idea for less frequent attention to a batch of ideas. This could have aggravated the alternation between production and stimulation if, for instance, the users had slowed down because of lack of cognitive stimulation or information overload dealing with the batches. In fact, TIDEA reveals that users on average needed −16.3% time to write an idea in OSB sessions, corresponding to an average cut down of 4.2 seconds per idea when users typed their ideas without being interrupted with feedthrough. We also found through variable PAUSE that on average the users switched 18.8% more rapidly (or 6.4 seconds faster) from production to stimulation (presuming the pauses typing correspond to reading the others' ideas).

For the remaining variables in Table 3, the results revealed small differences between OSB and CB sessions, thus likely explained by chance. CHARS, the number of characters typed per user in a session, was 6.3% higher in OSB sessions, influenced by the higher number of ideas produced, but balanced by slightly fewer characters per idea (CIDEA had a mean variation of −3.7%). Finally, the number of deleted characters, DCHARS, was 3.6% lower in OSB sessions.

7. Discussion

In this section, we elaborate on how users act when they receive feedthrough, then we analyze the potential problem of some of the ideas not being delivered because of the buffering technique employed by the opportunity seeker, and, finally, we discuss the limitations of this study, in particular concerning the lack of a qualitative evaluation.

7.1. Validation of patterns of user activity

Earlier, we discussed three patterns of user activity in brainstorming sessions with immediate broadcast of ideas. These patterns are important because they supply the basic assumptions adopted by the opportunity seeker.

We now provide evidence for the first two patterns (that users typically do not stop typing when they receive feedthrough and that they usually pause after putting forward an idea) based upon fine-grained data collected during the laboratory experiment.

On the one hand, in the first 5 seconds after the reception of feedthrough, users continued typing their idea at a mean rate between 1.4 and 1.6 key presses per second (*stdev* between 0.7 and 0.8). On the other hand, after submitting an idea to the group, the users almost stopped typing for at least 5 seconds, with a mean rate between 0.1 and 0.2 key presses per second (*stdev* between 0.2 and 0.3). This provides evidence to validate the two patterns mentioned above.

7.2. Batch size and inactivity timeout

We adopted a batch size for feedthrough delivery of 10 ideas. Of course we do not have strong evidence regarding this empirical decision. However, we note the average delivery of ideas, shown in Table 3, is 1.9. Furthermore, 95% of deliveries had 5 ideas or less. These results indicate the specified batch size actually had no impact on the experiment, and indeed may be of secondary importance.

However, care should be taken when considering larger groups. For instance, if the group had 20 or more users, and assuming an average rate of 1.5 ideas per minute per user and an average time to write an idea of 21.5 seconds (data from Tables 2 and 3, for the experimental condition), then the Poisson probability of more than 10 ideas arriving during the time to generate an idea is greater than 0.5. In other words, the number of ideas arriving would likely be greater than the batch size, in which case there would be increasing delays in the delivery of ideas, ultimately leading to several not being shown to the users. This might hamper group creativity and productivity although the benefits of delivering those extra ideas in larger batches (at a rate of 30 per minute) would certainly be outweighed by the cost of information overload.

We also adopted an inactivity timeout of 10 seconds. Again, we do not have strong evidence regarding this empirical decision. But analyzing the numbers given in Table 3 we observe the participants in control groups had pauses averaging 34 seconds. This might indicate the established timeout should have been longer. Interestingly, we also observe the participants using the opportunity seeker had pauses averaging 28 seconds. This might indicate the timeout effectively served to balance the user's alternations between production and stimulation, which supports our departing assumption.

7.3. Undelivered ideas

One of the concerns of buffering ideas during brainstorming sessions, instead of immediately broadcasting them, is that the ideas submitted near the end of the session may not be delivered to some of the users. This may happen when a user is less productive than the others, either because s/he types very slowly or does not type at all due to lack of inspiration. As explained earlier, in these circumstances the opportunity seeker delays feedthrough until the user finally submits the idea to the group or until a timeout occurs.

Undelivered ideas	0	1	2	3	4	5	6	7
Number of sessions	40	7	4	1	1	1	0	1

Table 4: Sessions with undelivered ideas

Since it is undesirable to have undelivered ideas, we measured group production in each OSB session and subtracted from it the number of ideas from others actually received by each user. Table 4 shows that in 72.7% of the cases (or 40 sessions out of a total of 55) all ideas were delivered to the users and that in 20.0% of the times one or two ideas were not delivered; the remaining 7.3% were for cases with between 3 and 7 undelivered ideas, each occurring only once. In other words, these data reveals that the users' natural work rhythm was rapid enough so that less than one idea (*mean* = 0.6, *stdev* = 1.4) was not delivered at the end of a session with the opportunity seeker, which seems reasonable.

7.4. Limitations

We had to accept several compromises for this study, most of them related to the absence of a qualitative analysis of both the users' ideas and the videos that were captured during the brainstorming sessions.

Firstly, we did not evaluate nor compare the quality of the ideas due to the subjective nature of this task and also because it would have required several evaluators, which have not been available so far. Quantity is one of the goals of brainstorming (Osborn 1963) and there is evidence that quality is positively linked to quantity (Briggs, Reinig et al. 1997). On the other hand, we shall also account for research pointing out that brainstorming should be evaluated by measuring good ideas (Reinig and Briggs 2006).

Secondly, we did not investigate duplicate ideas, something that could be explicitly addressed in a qualitative analysis. The interest here would be to know if the opportunity seeker artificially inflated the number of generated ideas by causing users to unknowingly submit ideas equivalent to those stored in the buffer but not yet displayed. However, with immediate feedthrough users may not be able to keep up with the others, which might also lead to duplicate ideas. Thus, a comparison between the two conditions on this topic is appealing and its results could eventually help fine-tune the opportunity seeker.

Thirdly, we always used the same values for the two parameters in the opportunity seeker: no more than 10 ideas are delivered at once and the timeout period for delivering feedthrough is 10 seconds. We could have considered other values but that would have increased the complexity of the experimental design beyond our current logistic capacity.

Fourthly, we faced many problems while examining the video feeds of the computer screen and the user's faces. The purpose was to make observations related to the three patterns of user activity identified earlier: (a) if users are able to attend feedthrough and write an idea simultaneously; (b) if the pause in typing activity after the submission of an idea coincides with the user attending feedthrough; and (c) if periods of inactivity correspond to lack of imagination, distraction, or to engaged reading. However, the videos showed users who appear to be indiscriminately focused on the task and the computer screen most of the time.

Very occasionally, there was an outward reaction to reading an idea such as a frown or a smile. It was also infrequent to observe users acting distracted, for instance, staring somewhere else than the computer screen. Given this data, it was impossible to accurately distinguish when a user was reading ideas, pausing, or distracted, so we had to discard these data.

Finally, we did not directly assess the degree to which users actually experienced information overload. There exist several techniques that could provide insight into this, such as physiological measures and self-assessments of mental workload (Wickens and McCarley 2008), which could be applied in future experiments.

8. Conclusions and future work

We highlighted the need to develop active compensation mechanisms to maintain the capacity to attend the group in group support systems. We made contributions to conceptualize the group attention problem and studied the problem in the specific domain of synchronous brainstorming.

We developed the opportunity seeker, a component responsible for managing feedthrough information. The opportunity seeker considers the natural rhythms of group work to identify the most adequate time to raise the users' attention to the group. We showed how the opportunity seeker may be implemented in synchronous brainstorming tools, and especially how the most adequate time to raise attention may be detected by analyzing the users' keyboard activity.

We provide evidence that the opportunity seeker may increase the work done by groups. The improvement amounts to 9.6% more ideas produced in synchronous brainstorming.

In addition, results from a post-hoc analysis show that the opportunity seeker reduced feedthrough by 44.1% by combining ideas in batches. This resulted in 54.7% more time for users to think about and type ideas without being interrupted. In these conditions, the users were 18.8% faster in alternating between generating an idea, which they did in 16.3% less time, and being stimulated by the other users' ideas.

We believe the opportunity seeker addresses the today's and tomorrow's information overload problems: on the one hand, even if the users in our experiment were not overloaded with information, the number of ideas produced was, nonetheless, higher; on the other hand, the opportunity seeker affords synchronous brainstorming sessions with larger group sizes, because it ensures that each user will be exposed to new ideas from others at his/hers own natural rhythm, thus automatically mitigating information overload.

As for future work, we are considering several research paths: one is to address the limitations presented earlier; another is to experiment with the opportunity seeker in other types of group work; finally, we have plans to introduce an eye-tracker in future experiments.

Acknowledgements

This work was supported by the Portuguese Foundation for Science and Technology, through project PTDC/EIA/67589/2006 and the Multiannual Funding Programme, and by the Departamento de Postgrado y Postítulo of the Vicerrectoría de Asuntos Académicos of the Universidad de Chile.

References

Anderson, J. (2005). Cognitive psychology and its implications. New York, NY, Worth Publishers.

Bailey, B. and J. Konstan (2006). "On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state." Computers in Human Behavior 22(4): 685-708.

Briggs, R., B. Reinig, M. Shepherd, J. Yen and J. Nunamaker (1997). Quality as a function of quantity in electronic brainstorming. Proceedings of the thirtieth Hawaii international conference on System sciences, Washington, DC, IEEE Press.

Brooks, F. (1975). The mythical man month. Reading, Massachussets, Addison-Wesley.

Card, S., T. Moran and A. Newell (1983). The psychology of human-computer interaction. Hillsdale, NJ, Lawrance Elrbaum.

Chen, D. and R. Vertegaal (2004). Using mental load for managing interruptions in physiologically attentive user interfaces. CHI'04: Extended abstracts on Human factors in computing systems, Vienna, Austria, ACM Press.

Conway, A., N. Cowan and M. Bunting (2001). "The cocktail party phenomenon revisited: The importance of working memory capacity." Psychonomic Bulletin & Review 8(2): 331-335.

Connolly, T., L. Jessup, J. Valacich (1990). "Effects of anonymity and evaluative tone on idea generation in computer-mediated groups." Management Science 36(6): 689–703.

Dehaene, S., J. Changeux, L. Naccache, J. Sackur and C. Sergent (2006). "Conscious, preconscious, and subliminal processing: A testable taxonomy." Trends in Cognitive Sciences 10(5): 204-211.

Dekker, S. (2006). The field guide to understanding human error. Hampshire, England, Ashgate.

Dennis, A. and B. Reinicke (2004). "Beta versus vhs and the acceptance of electronic brainstorming technology." Management Information Systems Quarterly 28(1): 1-20.

Dennis, A., J. Valacich, T. Connolly and B. Wynne (1996). "Process structuring in electronic brainstorming." Information Systems Research 7(2): 268-277.

Dennis, A. and J. Valacich (1993). "Computer brainstorms: More heads are better than one." Journal of Applied Psychology 78(4): 531-537.

Dennis, A. and M. Williams (2003). Electronic brainstorming: Theory, research, and future directions. Group creativity: Innovation through collaboration. New York, Oxford University Press.

Douglas, S. and A. Kirkpatrick (1999). "Model and representation: The effect of visual feedback on human performance in a color picker interface." ACM Transactions on Graphics 18(2): 96-127.

Endsley, M. (2003). Designing for situation awareness, Taylor & Francis.

Eppler, M. and J. Mengis (2004). "The concept of information overload: A review of literature from organization science, accounting, marketing, MIS, and related disciplines." Information Society 20(5): 325-344.

Fogarty, J., A. Ko, H. Aung, E. Golden, K. Tang and S. Hudson (2005). Examining task engagement in sensor-based statistical models of human interruptibility. Proceedings of the SIGCHI conference on Human factors in computing systems, Portland, OR, ACM Press.

Gallupe, R., L. Bastianutti and W. Cooper (1991). "Unblocking brainstorms." Journal of Applied Psychology 7(1): 137-142.

Gallupe, R., W. Cooper, M. Grise and L. Bastianutti (1994). "Blocking electronic brainstorms." Journal of applied psychology 79(1): 77-86.

Gallupe, R., A. Dennis, W. Cooper, J. Valacich, L. Bastianutti and J. Nunamaker (1992). "Electronic brainstorming and group size." Academy of Management Journal 35(2): 350-369.

Henninger, W., A. Dennis and K. Hilmer (2006). "Individual Cognition and Dual-Task Interference in Group Support Systems," Information Systems Research 17: 415-424.

Hill, J. and C. Gutwin (2003). Awareness support in a groupware widget toolkit. Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work, Sanibel Island, Florida, ACM Press.

Hollnagel, E. and D. Woods (2005). Joint cognitive systems: Foundations of cognitive systems engineering, CRC Press.

Horvitz, E., C. Kadie, T. Paek and D. Hovel (2003). "Models of attention in computing and communication: From principle to applications." Communications of the ACM 46(3): 52-59.

Hymes, C., G. Olson (1992). Unblocking brainstorming through the use of a simple group editor. Proceedings of the 1992 ACM conference on Computer-supported cooperative work, Toronto, Canada, ACM Press.

Iqbal, S., P. Adamczyk, X. Zheng and B. Bailey (2005). Towards an index of opportunity: Understanding changes in mental workload during task execution. Proceedings of the SIGCHI conference on Human factors in computing systems, Portland, OR, ACM Press.

Jenkin, T., J. McGeachie, D. Fono and R. Vertegaal (2005). Eyeview: Focus+context views for large group video conferences. CHI'05: Extended abstracts on Human factors in computing systems, Portland, OR, ACM Press.

LaBerge, D. (1999). Attention. Cognitive science. B. Bly and D. Rumelhart. San Diego, CA, Academic Press.

Lyman, P. and H. Varian (2003). How much information?, School of Information Management and Systems, University of California at Berkeley.

McDaniel, S., G. Olson and J. Magee (1996). Identifying and analyzing multiple threads in computer-mediated and face-to-face conversations Computer Supported Cooperative Work '96, Cambridge MA, ACM Press.

Osborn, A. (1963). Applied imagination. New York, Charles Scribner's Sons.

Perrow, C. (1999). Normal accidents, living with high-risk technologies. Princeton, New Jersey, Princeton University Press.

Pinsonneault, A., H. Barki, R. Gallupe and N. Hoppen (1999). "Electronic brainstorming: The illusion of productivity." Information Systems Research 10(2): 110-133.

Reason, J. (1990). Human error. Cambridge, UK, Cambridge University Press.

Reason, J. (1997). Managing the risks of organizational accidents. England, Ashgate.

Reason, J. (2008). The human contribution: Unsafe acts, accidents and heroic recoveries. Surrey, England, Ashgate.

Reinig, B. and R. Briggs (2006). Measuring the quality of ideation technology and techniques. HICSS '06. Proceedings of the 39th Annual Hawaii International Conference on System Sciences, Hawaii.

Roda, C. and J. Thomas (2006). "Attention aware systems: Introduction to special issue." Computers in Human Behavior 22(4): 555-556.

Roda, C. and J. Thomas (2006). "Attention aware systems: Theories, applications, and research agenda." Computers in Human Behavior 22(4): 557-587.

Simon, H. (1971). Designing organizations for an information-rich world. Computers, communication, and the public interest. M. Greeberger. Baltimore, MD, Johns Hopkins University Press: 37-72.

Simons, D. and R. Rensink (2005). "Change blindness: Past, present, and future." Trends in Cognitive Sciences 9(1): 16-20.

Sternberg, R. (2003). Cognitive psychology. Belmont, CA, Wadsworth.

Tung, L. and E. Turban (1998). "A proposed research framework for distributed group support systems." Decision Support Systems 23: 175-188.

Van de Ven, A. and A. Delbecq (1976). "Determinants of coordination modes within organizations." American Sociological Review(41): 322-338.

Vertegaal, R. (1999). The gaze groupware system: Mediating joint attention in multiparty communication and collaboration. Proceedings of the SIGCHI conference on Human factors in computing systems, Pittsburgh, PA, ACM Press.

Vertegaal, R. (2003). "Attentive user interfaces: Introduction." Communications of the ACM 46(3): 30-33.

Vertegaal, R., I. Weevers, C. Sohn and C. Cheung (2003). GAZE-2: Conveying eye contact in group video conferencing using eye-controlled camera direction. Proceedings of the SIGCHI conference on Human factors in computing systems, Ft. Lauderdale, FL, ACM Press.

Vertegaal, R., J. Shell, D. Chen and A. Mamuji (2006). "Designing for augmented attention: Towards a framework for attentive user interfaces." Computers in Human Behavior 22(4): 771-789.

Weick, K. (1995). Sensemaking in organizations. Thousand Oaks, CA, Sage.

Weick, K. (2001). Making sense of the organization. Oxford, UK, Blackwell.

Wensveen, S., J. Djajadiningrat and C. Overbeeke (2004). Interaction frogger: A design framework to couple action and function through feedback and feedforward. Proceedings of the 2004 conference on Designing interactive systems, Cambridge, MA.

Wickens, C. and J. McCarley (2008). Applied attention theory. Boca Raton, FL, CRC Press.