

Attentive Groupware

António Ferreira
Pedro Antunes

DI-FCUL

TR-06-12

July 2006

Departamento de Informática
Faculdade de Ciências da Universidade de Lisboa
Campo Grande, 1749-016 Lisboa
Portugal

Technical reports are available at <http://www.di.fc.ul.pt/tech-reports>. The files are stored in PDF, with the report number as filename. Alternatively, reports are available by post from the above address.

Attentive Groupware

António Ferreira Pedro Antunes

Department of Informatics, University of Lisbon, Portugal

July 2006

Abstract

In this paper we present a combined perspective over groupware and attentive user interfaces. The problem that we address is how to attenuate the distracting effects of the multiple interruptions that characterize computer-mediated collaboration, while maintaining the necessary notion of group awareness. We introduce a framework of group attention consisting of collaboration facets, information flows, and attentive devices. Based on known psychological factors we propose four such devices: the time separator, the opportunity seeker, the change emphasize, and the activity anticipator. Finally, we demonstrate the framework role in the design of an electronic brainstorming tool.

1 Motivation

A continuing trend in groupware research aims at improving the sense of proximity among group members, whether by enabling geographically distant people to work together, or by supporting ongoing activity at different times. Researchers often use terms such as “group awareness” to describe techniques that enhance group consciousness by providing *ever greater* information about the actions of its members. A main argument is that communication channels mediated by computers are relatively poor when compared to more natural settings, such as face-to-face meetings [1]. However, a problem with this trend is that it fails to recognize that sometimes *more is less* due to the limitations in the human attentive capacity, especially as we become surrounded with computers and, not necessarily useful, information.

During the late 1990’s several researchers from the Human-Computer Interaction (HCI) field became interested in Attentive User Interfaces, or AUI. 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 [2]. Therefore, instead of assuming the user is always focused on the whole computer screen and willing to take responsibility for optimizing his or hers attention (e.g. by raising or lowering windows), AUI negotiate, rather than impose, the users’ attention by establishing priorities for presenting information.

Most AUI research is directed towards single-user activity, the main assumption being that individual performance degrades with the number of simultaneous requests for attention. Experiments in this area typically study the effects on a primary task caused by interruptions that require processing of a secondary task, possibly unrelated and unexpected, such as coding a computer program

while occasionally having to make a mental multiplication [3]. Based on these postulates, researchers are enhancing both input and output devices so that the user remains focused on the primary task without getting too much distracted by the secondary task. This is presently accomplished by sensors that can detect our state of attention (e.g. eye-gaze, body orientation) [2, 4], by statistical models of human interruptibility that determine the best time to communicate with the user [5, 3], and by displaying information with various levels of detail, depending on the user’s current focus of attention [2].

There exists some research on applying AUI to groupware but it is rare and situated in video conferencing [6]. We argue that groupware design presents a privileged environment for extending the AUI body of knowledge and hypothesize that group performance improves if groupware mediation uses AUI concepts. We put forward the following motivating postulates:

- People involved in group work attend to interruptions much more often than individuals. This is directly linked with the need to maintain group awareness and with the greater number of information flows;
- Collaborative work is more fragmented than individual work. Instead of executing a single, extensive task, group members tend to perform a series of intertwined tasks, one at a time. This work strategy improves group awareness and enables faster recovery from errors;
- In group work the primary and secondary tasks are typically related and can both contribute to the shared goal.

Our objective in this paper is to raise interest in the groupware community by introducing a framework of group attention for aiding groupware design. The framework is targeted at groupware supporting collaborative activities and makes no assumption regarding time and place. It integrates the above postulates and is organized in three parts:

- Collaboration facets that group members attend to while working;
- Information flows for maintaining attention between group members;
- Specialized devices for sensing and enhancing group attention.

We explain the framework of group attention in Sect. 2. In Sect. 3 we describe an application to an electronic brainstorming tool and outline a laboratory experiment for testing the hypothesis about group performance improvements. In Sect. 4 we conclude the paper with a discussion regarding our approach and with prospects for future work.

2 Framework of Group Attention

The starting point for interpreting the framework of group attention is to acknowledge that humans, and therefore *groups*, have limited information processing capacity. This limitation is particularly relevant in scenarios where groups need to reach a common goal, and even more so if computers are used to mediate communication between group members, due to technological and design

constraints. The purpose of this framework is to help designers operationalize the concept of group attention by classifying the aspects of collaboration that group members attend to during group activity (Sect. 2.1), the information flows supported by groupware to maintain group attention (Sect. 2.2), and the specialized devices that can be installed to sense and enhance group attention (Sect. 2.3).

2.1 Collaboration Facets

The first part of the framework identifies facets that characterize the intricacies of collaborative work. These facets are generalizations of our prior study of workspace collaboration [7]. The focus here is directed towards capturing collaboration features that ultimately drive group performance. This is information that group members recurrently must attend to, to keep themselves aligned with the shared goal and ongoing activities. We distinguish the following complementary facets: production, opportunities, and restrictions.

The *production* facet informs about elements of work in production or that have already been produced. It might be organized separately into group and self production especially for comparison purposes. More production means the group is reaching the shared goal. In single-user software design this facet can be used for assessing individual goal achievements. However, with collaborative work, information about group performance is more than a simple combination of individual productions. We try to capture this with the other two facets.

The *opportunities* facet is related to the intertwined nature of collaboration: if one or more group members stop, then the group may also halt, eventually never reaching the shared goal. This suggests that collaboration among group members is bound by opportunity dependencies created by the achievement of individual goals. For instance, a group member may attend to the outcomes of other members in order to choose the next move. In other words, opportunities represent starting points for future production. Having opportunities means the group may progress faster towards the shared goal.

The *restrictions* facet reflects possible negative consequences of coordination and communication during collaborative activity. On the one hand, the prevention of conflicts and duplicate efforts (positive consequences of coordination) may slow down or even impede the work of some group members; on the other hand, communication delays may significantly reduce group performance. Therefore, more restrictions means the group may progress slower towards the shared goal. For example, a group member may be unable to proceed while perceiving that another is producing elements of work. This waiting pattern may be tolerated if the number of targets of attention is reasonable, but collaborative work can rapidly turn it into something overwhelming for the human attentive capacity.

These facets of collaboration contribute to the design process with a modular structure for thinking about attention and group performance. Our vision is that they should be general enough to cover a broad range of collaborative scenarios, while at the same time being flexible enough to be applied individually, if appropriate. For instance, attention in electronic brainstorming sessions may be analyzed using the production and opportunities facets (ideas generate new ideas), but probably not from the restrictions facet if group members do not have to coordinate themselves to contribute their ideas; on the other hand, the

restrictions facet may provide useful information to users sharing a workspace with graphical objects and locking mechanisms.

In more practical terms, designers can attach many types of representations to the information group members attend to: graphical maps, current values, rates of change, activity labels, to name a few. For example, production can be shown in absolute terms or as tendencies over time. These representations have to be conveyed using information flows mediated by computers, alongside with other manifestations of attention (such as requests or guidance). These flows are the subject of the next section.

2.2 Information Flows

The second part of the framework examines information flows that can be used for maintaining group attention. The focus here is on the different forms of interaction between group members and on the mediation provided by the groupware. These information flows should reflect the outcomes of collaborative activity, and support mutual awareness and interdependence. Our strategy is to reuse and generalize flows that we investigated in a previous study about shared workspaces design [8], the result being a characterization in terms of explicit communication, feedthrough, and back-channel feedback.

The *explicit communication* flow addresses information produced by a group member and explicitly intended to be received by other members [1]. For instance, a group member may request permission to use a locked object by calling the attention of the current owner with a reminder note; another example might be an instructor providing online guidance to students during collaborative problem solving. This information flow can be supported by a groupware interface capable of multiplexing information from input devices to several output devices.

The *feedthrough* flow concerns implicit information delivered to several group members reporting actions executed by one member [9]. Feedthrough provides group awareness and helps construct meaningful contexts for collaboration. For instance, a shared workspace may show currently selected menus for each group member that is manipulating objects. However, because such high level of detail may overwhelm the human attentive capacity, feedthrough information can also be less detailed, for example by being linked to concise views of production, opportunities, and restrictions. The groupware interface can generate feedthrough flows by capturing each group member's inputs and then multiplexing feedback information (replies to a single-user) to the other members.

The *back-channel feedback* flow concerns unintentional information initiated by one group member and directed towards another member to facilitate communication [10]. It may be automatically captured and produced by the groupware interface using sensors of eye-gaze, body orientation, and others [2, 4]. Apparently, no significant content is delivered through back-channel feedback since it does not reflect cogitation from the user. However, as we mentioned in the motivation, from the AUI perspective this information flow is very important because it allows the groupware to sense and adapt to human states of attention. For example, in a video conferencing experiment eye tracker sensors were utilized to prolong eye contact perception, therefore minimizing distractions during the group turn taking process [6].

Groupware designers may use information flows to maintain attention between group members, whether by allowing explicit interruptions, by keeping the group informed about ongoing activity, or by disseminating the attentive states of group members. We expect some designs may combine all information flows, especially in groupware for synchronous collaboration, while other designs may require only explicit communication or feedthrough, such as in asynchronous or same-place groupware.

In the next section we approach the problem of enhancing group attention, given that, as evidenced earlier, the amount of information generated by group members and by the groupware itself may overcome the human attentive capacity and, therefore, may decrease group performance.

2.3 Attentional Devices

The third part of the framework introduces devices that manipulate information flows to enhance the attentive capacity of group members. The design problem that we address here is how to attenuate the distracting effects of the multiple interruptions that characterize computer-mediated collaboration, while maintaining the necessary notion of group awareness. Our approach is to start by categorizing existing groupware devices, more specifically those related to collaborative interaction, into awareness and coupling devices, and then describe specialized devices that can be embedded into the groupware interface to dynamically manage attention according to known psychological factors.

We define *awareness devices* as devices dedicated to sense and display information about the collaborative activity within the group. This involves conveying representations of collaboration facets using the groupware information flows so that the group remains aligned with the shared goal. Several examples of awareness devices, such as radar views and video images, can, for instance, be found in a study organized around their capabilities for answering questions about who, what, and where [1]. We also include in this category sensors that capture human attention (e.g. eye-gaze, body orientation) [2, 4] and computer displays that vary the level of detail according to the user's attention focus [2].

Another feature of the awareness devices is that they allow group members to perceive the limitations and mediation role of the communication infrastructure. This is particularly relevant when the Internet is used to support the information flows, where, for instance, feedthrough delays are significantly longer and less predictable than feedback delays [11]. These limitations also affect the explicit communication and back-channel feedback flows, and contribute to design tradeoffs, since, otherwise, the delays could turn into frequent interruptions that distract and reduce group performance.

The second category of groupware devices are the *coupling devices*, for letting group members loose the link between executed actions and group awareness [12]. We consider two types of coupling control: on the one hand, group members may control coupling at the origin to specify what and when information should become public (e.g. committing source code changes in a revision control system); on the other hand, coupling can be controlled at the destination by specifying filters that restrict group awareness to some selected objects and actions (e.g. moving the region of interest in a shared workspace with a viewport).

Coupling devices require manual discrimination and control of awareness information delivery, thus penalizing individual performance. However, this

disadvantage is offset by the capacity to limit the amount of information and number of interruptions, which may improve group attention. This tradeoff sets the stage for introducing specialized devices that dynamically manage group attention.

We propose a set of devices dedicated to collect and combine information received from awareness sensors associated with each group member, and to automatically managing information delivered to their awareness displays, according to known psychological factors. The novel aspect in our approach, in comparison with AUI research, is that processing is done at the group interaction level, meaning that the devices have access, not only to information that reaches a user, but also to information that originates either from the same user or from other group members. The *attentional devices* are:

Time separator Delivers information to the awareness displays after a predefined amount of time has elapsed since the previous delivery. The purpose of this device is to attenuate the effects of the “psychological refractory period” (also called “attentional blink”) during which response time is slower [13].

Opportunity seeker Combines awareness information with the attentive state of a group member to determine the most adequate time to deliver the information to the corresponding awareness displays. The objective of this device is to seek opportunities that minimize distraction due to interruptions.

Change emphasize Highlights changes that occurred since the previous delivery of information to the awareness displays. The purpose of this device is to attenuate “change blindness” by facilitating the detection of changes in order to avoid having to attend to the full group awareness picture [14].

Activity anticipator Senses activity that may affect group performance and delivers preliminary information to the awareness displays. The objective of this device is to prepare group members to be attentive to upcoming collaborative outcomes, thereby enabling faster response times [15].

These attentional devices contribute to the design process with explicit support for mediating group attention in a broad range of collaborative scenarios. For example, in asynchronous groupware, the **change emphasize** may be used to highlight differences between two discrete states of group work. We provide a more detailed example in the next section, where we apply all these attentional devices in an electronic brainstorming tool, in the context of a laboratory experiment for testing the effects on group performance caused by the devices.

3 Application and Experimental Setup

We applied the proposed framework to the development of an electronic brainstorming tool. The following description is organized according to the three parts of the framework.

Regarding the collaboration facets, production corresponds to submitted ideas, opportunities are ideas that generate new ideas, and we assume there are no restrictions, since group members may make suggestions at any time

during the brainstorming session. In essence, from the designer’s point of view, brainstormers may increase their performance by attending to ideas, either their own or from others.

Concerning information flows, the interactions within the group are solely done via feedthrough, that is, the brainstorming tool automatically distributes ideas to all group members as they are being generated, without requiring explicit communication or back-channel feedback between one group member and another.

Finally, regarding the attentional devices, the brainstorming tool implements a **time separator** showing new ideas only at time intervals defined by the designer, an **opportunity seeker** that waits for the group member to stop typing to display new ideas from others, a **change emphasize** highlighting ideas from others that have been recently delivered, and an **activity anticipator** that senses group inactivity and alerts brainstormers that the session will soon end unless new ideas are put forward.

We now briefly describe a laboratory experiment that uses the attentional devices in the brainstorming tool as independent variables that we are setting up in order to test the hypothesized improvement in group performance. The collaborative work scenario is characterized by a group of about 10 participants gathered in the same room and using the brainstorming tool. Each session lasts at most 60 minutes, during which the tool measures the number of ideas, the number of ideas from others that were received while typing ideas, and the time of inactivity immediately after the reception of new ideas from others. At the end, every participant independently fills out a questionnaire regarding the perceived state of attention during the brainstorming session. We are currently finalizing the tool development and expect to run this experiment soon.

4 Discussion and Future Work

The study of the mediating role of computers on group attention is largely an unexplored research area. While current trends keep aiming at conveying ever greater information about the group, we suggest a route that explicitly recognizes the limitations of the human, and therefore the group, information processing capacity. This route is consistent with AUI research, almost entirely directed at single-user activity, and we argue that the existing body of knowledge can be extended into the groupware field. To this end we introduce a framework of group attention for groupware design, and hypothesize that group performance improves if groupware tools make use of specialized attentive devices.

Many questions remain unanswered: does group performance significantly improve? What attentive devices are best suited for different collaborative scenarios? Can groups be made larger while remaining attentive and productive? We are currently addressing some of these questions with an electronic brainstorming tool. The road is open for many more experiments and applications.

References

- [1] Carl Gutwin and Saul Greenberg. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work*, 11(3):411–446, 2002.
- [2] Roel Vertegaal. Attentive user interfaces. *Communications of the ACM*, 46(3):30–33, 2003.
- [3] James Fogarty, Andrew J. Ko, Htet Htet Aung, Elspeth Golden, Karen P. Tang, and Scott E. Hudson. Examining task engagement in sensor-based statistical models of human interruptibility. In *CHI'05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 331–340, Portland, Oregon, USA, April 2005.
- [4] James Fogarty, Scott E. Hudson, Christopher G. Atkeson, Daniel Avrahami, Jodi Forlizzi, Sara Kiesler, Johnny C. Lee, and Jie Yang. Predicting human interruptibility with sensors. *ACM Transactions on Computer-Human Interaction*, 12(1):119–146, 2005.
- [5] Eric Horvitz, Carl Kadie, Tim Paek, and David Hovel. Models of attention in computing and communication: From principles to applications. *Communications of the ACM*, 46(3):52–59, 2003.
- [6] Roel Vertegaal, Ivo Weevers, Changuk Sohn, and Chris Cheung. GAZE-2: Conveying eye contact in group video conferencing using eye-controlled camera direction. In *CHI'03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 521–528, Ft. Lauderdale, Florida, USA, April 2003.
- [7] António Ferreira and Pedro Antunes. Quantitative evaluation of workspace collaboration. In *CSCWID'06: Proceedings of the tenth international conference on Computer supported cooperative work in design*, pages 1065–1070, Nanjing, China, May 2006.
- [8] Pedro Antunes, António Ferreira, and José A. Pino. Analyzing shared workspaces design with human-performance models, 2006. Accepted for publication in the CRIWG'06: Twelfth international workshop on Groupware.
- [9] Jason Hill and Carl Gutwin. Awareness support in a groupware widget toolkit. In *GROUP'03: Proceedings of the 2003 international ACM SIG-GROUP conference on Supporting group work*, pages 258–267, Sanibel Island, Florida, USA, November 2003.
- [10] Sonya Rajan, Scotty D. Craig, Barry Gholson, Natalie K. Person, and Arthur C. Graesser. AutoTutor: Incorporating back-channel feedback and other human-like conversational behaviors into an intelligent tutoring system. *International Journal of Speech Technology*, 4(2):117–126, 2001.
- [11] Carl Gutwin, Steve Benford, Jeff Dyck, Mike Fraser, Ivan Vaghi, and Chris Greenhalgh. Revealing delay in collaborative environments. In *CHI'04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 503–510, Vienna, Austria, April 2004.

- [12] Prasun Dewan and Rajiv Choudhary. Coupling the user interfaces of a multiuser program. *ACM Transactions on Computer-Human Interaction*, 2(1):1–39, 1995.
- [13] Michael W. Eysenck and Mark T. Keane. *Cognitive psychology: A student's handbook*, pages 119–150. Psychology Press, fourth edition, 2000.
- [14] Daniel J. Simons and Ronald A. Rensink. Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9(1):16–20, 2005.
- [15] David LaBerge. Attention. In Benjamin Martin Bly and David E. Rumelhart, editors, *Cognitive Science*, pages 44–98. Academic Press, San Diego, CA, USA, 1999.