THE PAST, PRESENT, AND FUTURE OF WORKGROUPS IN A THEATRE OF WORK

Tom Gross[1]

The members of workgroups need information about each other in order to cooperate efficiently and effectively. In face-to-face situation this information is available and can be captured naturally. In other situations the distributed users need technological support for this awareness information. Such systems should provide information about things that are going on currently, but also about the history of the actions and interactions in the workgroup. Ideally such systems should also provide predictions about possible future events. In this paper we introduce the Theatre of Work Enabling Relationships (TOWER) environment, which provides an infrastructure for awareness information support. We then discuss how it can be used to present information on the workgroup's past, presence, and future.

1 Introduction

In the CSCW literature it has been emphasised for years that efficient and effective cooperation requires that the cooperating individuals are well informed about their partners activities [7]. They require information about the other persons they are cooperating with, about their actions, about shared artefacts, and so forth. This information is often referred to as awareness (sometimes with prepositions such as *group* awareness [3, 8] or *workspace* awareness [15].

In situations where the cooperating individuals are at the same place this information is often perceived automatically [18]. In other situations where individuals, who are at different places, have to cooperate as a group, technological support for the cooperation process as well as the perception of cooperative activities is essential.

The types of awareness that are supported by technology today range from informal awareness about other people (or presence and availability awareness; or shared awareness [6]; or general awareness [11]) to workspace awareness about shared artefacts [16].

---

[1] Faculty of Media, Bauhaus-University Weimar, Germany
tom.gross(at)medien.uni-weimar.de

In the earliest attempts the information was captured and presented within one single application (e.g., [2, 19]). This first generation can be called *proprietary awareness systems*. In the second generation toolkits were developed that contained components for presenting awareness information. These *awareness widget toolkits* made the development of applications easier, because the software developers did not have to implement their awareness widgets from scratch (e.g., [23]). In the third generation, *awareness information environments* allowed capturing information from various applications and other sources and presented the information in a generic representation such as with tickertapes of pop-up windows on the computer desktop (e.g., [9, 21]). Recently many systems aim to address the users' peripheral perception and use *ambient interfaces* for the presentation of the information [12, 20].

Besides these important challenges concerning technical and psychological aspects of capturing and presenting the information, some fundamental challenges also lie in the structuring and modelling the information. Adequate structuring and modelling of the information has many advantages for the users: it allows to provide the information when it is of most value for the respective user and task; it allows to reduce the amount of less important information; and it allows to reduce disturbance of users.

It is also a prerequisite for the provision of past, present as well as future information. Especially the latter has received particular attention in the HCI and CSCW communities recently [4, 17]. Two important questions have to be answered in this respect: first, what are the parameters that influence the information need, and secondly, how is the information need influenced. Important influential parameters are the current situation of the user in terms of technical and social environment and the current task. These parameters influence the quality, quantity and granularity of the information needed, and particularly the timing of the presentation of the information.

In this paper we briefly describe the Theatre of Work Enabling Relationships (TOWER) environment, which provides an infrastructure for awareness support. We show how it captures and presents information on past and present presence of users, on their activities, as well as on the evolution of shared workspaces and on the artefacts contained in these shared workspaces. We then discuss how it can be used to present information on the workgroup's and their workspaces' future. Finally, we discuss some recent related work of other authors. These related approaches reach from systems that can analyse users' activities and infer their interruptibility, to systems that analyse and present patterns of online behaviour of uses, to systems that automatically adapt the users' reachability, and finally to systems that predict the physical navigation of users.

2 Theatre of Work

The Theatre of Work Enabling Relationships (TOWER) environment aims to support distributed work groups or virtual communities with group awareness in their current work context. It provides an infrastructure for facilitating chance encounters and spontaneous conversations among remote users.

For this purpose, the infrastructure has various sensors capturing information about users and their activities and a range of indicators notifying users about the presence, availability, and current activities and tasks of the other users. *Sensors* capture user activities within the TOWER environment (e.g., logins, logouts), user activities on Win* platforms (e.g., changes to files, sharing of folders and files, starting of applications, opening of documents), user activities in shared workspaces (e.g., a sensor for the Basic Support for Cooperative Work (BSCW) system [5] records all activities in the shared workspaces such as user logins and logouts, folder creation, invitations users to shared folders, document uploads), and access to Web servers. A broad variety of *indicators* present the awareness information. Examples are lightweight indicators such as pop-up windows with pure text or tickertapes displaying messages about the other users and shared artefacts; AwarenessMaps, which provide awareness information in the context of shared workspaces; the TowerWorld, which presents shared artefacts and users in a 3D multi-user environment; ambient interfaces, which present the information in the physical environment of the users; and mobile client presenting light-weight information for users on the road. In this paper we will only have space for briefly describing the TowerWorld. Information on TOWER as a whole can be found in [22]; information on the ambient interfaces can be found in [12] and other indicators can be found in [14].

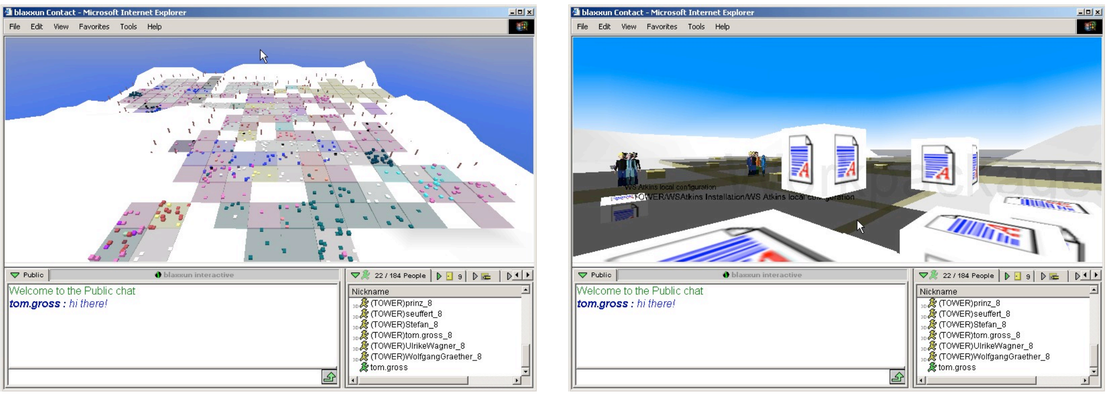

Figure 1. TowerWorld: (a) overview from a distance; (b) details in a close-up.

The *TowerWorld* is a 3D multi-user world; it consists of a stage that is dynamically created based on shared information space such as BSCW workspaces [5] or the content of other document management systems and avatars navigating on this stage and performing symbolic actions. Figure 1 shows screenshots of a TowerWorld: the first picture presents a view from the distance, where users can get a good overview of the whole stage with all its cubicles representing the documents; and the second picture shows a close-up of the same TowerWorld with more details (e.g., icons representing the file types, avatars positioned according to the current activities of the respective users).

The stage evolves in response to the patterns of use in the shared information space. The stage is generated and adapted according to rules and semantic mappings that can be specified by the users. Various attributes of the documents in the shared information space can be visualised such as the type of a document, the size of a document, the frequency of manipulations to a document, the creator of a document, the similarity among documents, and so forth. These attributes can be mapped to the size of the cubicles in the TowerWorld, their shape, their colour, their position, their clustering, and so forth. Another criterion for the stage creation is the granularity of the mapping of document sets into the stage. User workshops yielded different opinions whether a more detailed view or a more abstract overview provides better context for the visualisation of user activities. In the current implementation users can select between different worlds, each created by different selection criteria and rules for generation and mapping. In an overview world for example objects in the TowerWorld represent only folders of the shared information space, while the detailed TowerWorld provides a representation for each document. In the overview world activity spots are easier to recognise while in the detailed world clusters of objects with a similar semantic are easier to identified.

The stage of the TowerWorld is populated with avatars representing users and representing their current activities as symbolic actions such as automatic navigation through the TowerWorld and gestures. The emphasis in symbolic acting is to show the contextual information telling users about where the other users are, who they are and what they are doing right now. With symbolic acting the context dependent actions of all users are shown at all times so the world can seem like a more active place as well as more closely reflecting the activities of a user group. We let the system do the walking–and the acting. This is a very powerful and engaging way of solving problems in mediated communication.

3 Discussion

The TOWER environment it its current state provides a broad range of sensors capturing information, a variety of indicators presenting the information, and convenient means for specifying preferences.

Its particular strength lies in the modelling of awareness contexts and the entailed possibility for flexible adaptation of the information to the respective user, situation, and task. The TOWER environment allows users to specify their personal preferences with respect to the information they want to receive, with respect to the indicators used for the presentation of this information, and with respect to the timing of the presentation. A light-weight, but powerful, context model allows to structure the captured events into semantically coherent aggregations that make more sense to the users and abstract from, sometimes unwanted, details [13]. At the origin of events they are analysed and mapped to awareness contexts. Likewise, the events a user produces are analysed and mapped to a context. Now, the user can specify for each context which kind of information they want to receive, and how it should be presented. Furthermore, they can specify their preferences concerning the timing: they can opt for immediate presentation, for presentation in rhythms (e.g., once per hour, once per day), or particular moments (e.g., upon login, before logout). In order to facilitate the specification of the preferences and in order to allow for better co-orientation among users, the TOWER environment provides functionality for sharing preferences. That is, users can publish their preferences in shared workspaces, and other users can then subscribe to them.

So far, the TowerWorld provides *up-to-the-moment* information about the events and state of the users and their shared environment. A mechanisms, that is called DocuDrama allows to replay *past* states and evolutions of the TowerWorld [24]. An activity report provides a daily overview of the changes of the last 24 hours.

As far as the prediction of *future* events and states of the users and their shared environment are concerned, we have several ideas and plans for future work. One next step is to use the data of the BSCW shared workspace system to analyse patterns of behaviour of its users. The public BSCW server has already more than 100.000 registered users, who produce a vast amount of events. We have already extended BSCW to produce events that can be sent to and that are readable by the TOWER server. The following actions of BSCW users can be analysed: login patterns and patterns of cooperation on shared folders and documents (e.g., are there typical sequences of creating, reading, updating documents?). Furthermore, we can analyse events in the shared BSCW calendars, which are increasingly used since the introduction of the standardised iCalendar format allowing for the flexible exchange of calendar data with other calendaring systems.

As a result the users can then get predictions and warnings in this 2D indicators. For instance, once the system has identified the typical login and logout behaviour of a user, it can use the tickertape to warn other users who work in the same context that they should contact the colleague now (since she is leaving any minute). In an analogy to the fast replay of past states of the TowerWorld, we can animate the TowerWorld to play likely future events and evolutions of the TowerWorld. Finally, in many situations users expect something to happen, and can be notified when this is not the case. For instance, if the users are expected to put their deliverables into a certain workspace and nothing is happening in this workspace, then the corresponding region in the TowerWorld can be highlighted (with bright colours, or clouds over the respective landscape).

4 Related Work

There have been some recent approaches for analysing the behaviour of users of systems from computer-supported cooperative work and computer-mediated communication. This research has been stimulated by results from ubiquitous computing—in particular, the broad availability of sensing technology makes it possible to capture the behaviour of users in both the electronic world as well as in the real world.

Fogarty et al. [10] developed sensors, which run as a background process on users' primary computers and can capture users' activities on their computers (i.e., keyboard and mouse activity). Furthermore, audio sensors were used to detect the conversations and noise in offices; other sensors were attached to the doors and could measure the angle of the door (i.e., if the door was open, cracked, or closed). At the same time users had to do frequent self-reports—describing their interruptibility from highly interruptible (1) to highly non-interruptible (5). The automatically captured information was then compared with the results from the self-reports. The subjects were managers, programmers, and students on internships. On a whole 100 interruptibility self-reports were collected from each subject. On an average the interruptibility was more or less evenly distributed (about 16 to 19 percent per degree; with the exception of about 29 percent for the highly non-interruptible degree). The overall finding was that the model was reasonably adequate in inferring the interruptibility from the sensor data. There were 640 cases where the model correctly predicted values from 1 to 4. In 135 cases the system correctly predicted 5 (or highly non-interruptible). In only 43 it was the other way round—that is, the system predicted 1 to 4, but the users said 5. In 157 cases the model made the wrong judgement that the subject was interruptible (i.e., a value of 1 to 4), whereas the self-report showed non-interruptibility (i.e., 5). On a whole only these last 157 cases are really critical—since the users did absolutely not want to be disturbed, but the system judged that the users could be interrupted. On a whole the model showed an accuracy of 79,50 percent. With future versions of the inference can be improved by adding a component that allows the system to learn from explicit judgements of the users.

Begole et al. [4] also did some thorough analysis of users' activities on computers. This information was combined with information on the location of the activity, scheduling information from electronic calendars, and email activity. Here the goal was to derive typical patterns of user behaviour, which can then in turn be used for predictions of user behaviour. They developed a tool that can generate various visualisations of use patterns. The so-called actogramm shows a matrix with calendar days on the vertical axis and time of the day on the horizontal axis. For each interval, in which the user is active at the computer, the actogram shows a black bar. So, the individual rhythms per day can easily be identified. In another diagram, with the same coordinates both the actions and the meetings taken from the calendar are shown in the matrix. Furthermore, in an aggregated diagram with the aggregated online and meetings data on the vertical axis and the time of the day on the horizontal axis daily patterns can be identified (i.e., the typical time a user starts working, is out for lunch, or leaves in the evening). On a whole these visualizations provide great insight, but have two limitations—pointed out by the authors: the level of detail is one minutes and there is no more detailed information, and the pure information that a user is actively working on the computer does not necessarily mean that the user is interruptible.

The SenSay prototype by Siewiorek et al. [25] goes one step further: it does not only analyse the state of the user, but also performs some automated actions based on the inferences. Basically the Sensing & Saying (SenSay) prototype is a mobile phone, which is context-aware and modifies its behaviour based on its user's state and surroundings. It constantly adapts to the dynamically changing environmental and physiological states of its user. Furthermore, it provides the remote caller with information on the current context of the phone user, and lets the caller decide if the reason for the call is important enough to disturb a user in a given situation. It uses light, and motion sensors, as well as a microphone, which all are placed on various parts of the body of the phone user and connected to a central hub, which the user carries on his/her belt. The prototype then uses rules to infer the state of the user and then consequently adapts the profile of the phone to one of the four basic states: uninterruptible, idle, active, or normal (default state). Each state has influence on a number of phone actions (e.g., if the state is uninterruptible, the phone ringer and vibration are turned off). Although the system certainly is not always 100 percent correct with the inferences, it provides an interesting approach to facilitating the management of various settings of the increasing number of gadgets users are carrying around.

Ashbrook and Starner [1] provide a slightly different approach, but which has some similarity with our TOWER approach. They do thorough capturing of location data of users of GPS and WLAN technology, and then identify patterns of user movement in the real world, and use these patterns as a basis for predictions of the physical locations of users. This approach is analogous to our approach, since in TOWER the same predictions are made for the electronic world—that is, in TOWER we analyse the behaviour of users, try to identify patterns, and make predictions of future behaviour of users. These predictions are then mapped to movements of the avatars in the TowerWorld.

5 Conclusions

In this paper we have motivated the need for thorough analysis of the behaviour of users, which can facilitate the communication, coordination, and cooperation in distributed workgroups. We have described how the past and present behaviour of users are captured, analysed, and presented in TOWER. We have made suggestions as to how this approach can be used as a basis for predictions. Finally we have presented some recent approaches with similar ideas.

It is clear that already the analysis of the past and present behaviour of users has enormous implications on people's privacy. Making detailed analysis and presenting typical behaviour and rhythms of users as well as making predictions on future reachability of users can add tremendously to these privacy challenges. In this paper we did not have the space to deal with these issues. However, there are some ideas on how improve the situation (e.g., reciprocity in the sense that all users are mutually informed about each other, and no users can only be observer of others without being visible to others).

Acknowledgments

The research presented here was carried out by the IST-10846 project TOWER, partly funded by the EC. I would like to thank all colleagues from the TOWER team.

References

1. Ashbrook, D. and Starner, T.E. Using GPS to Learn Significant Locations and Predictd Movement Across Multiple Usres. Personal and Ubiquitous Computing 7, 5 (Nov. 2003). pp. 275-286.

2. Beaudouin-Lafon, M. and Karsenty, A. Transparency and Awareness in a Real-time Groupware System. In Proceedings of the ACM Symposium on User Interface Software and Technology - UIST'92 (Nov. 15-18, Monterey, CA). ACM, N.Y., 1992. pp. 171-180.

3. Begole, J., Rosson, M.B. and Shaffer, C.A. Flexible Collaboration Transparency: Supporting Worker Independence in Replicated Application-Sharing Systems. ACM Transactions on Computer-Human Interaction 6, 6 (June 1999). pp. 95-132.

4. Begole, J.B., Tang, J.C., Smith, R.B. and Yankelovich, N. Work Rhythms Analysing Visualisations of Awareness Histories of Distributed Groups. In Proceedings of the ACM 2002 Conference on Computer-Supported Cooperative Work - CSCW 2002 (Nov. 16-20, New Orleans, LO). ACM, N.Y., 2002. pp. 334-343.

5. Bentley, R., Appelt, W., Busbach, U., Hinrichs, E., Kerr, D., Sikkel, K., Trevor, J. and Woetzel, G. Basic Support for Cooperative Work on the World-Wide Web. International Journal of Human Computer Studies: Special Issue on Novel Applications of the WWW (Spring 1997).

6. Borning, A. and Travers, M. Two Approaches to Casual Interaction Over Computer and Video Networks. In Proceedings of the Conference on Human Factors in Computing Systems - CHI'91 (Apr. 27-May 2, New Orleans, LO). ACM, N.Y., 1991. pp. 13-20.

7. Dourish, P. and Belotti, V. Awareness and Coordination in Shared Workspaces. In Proceedings of the Conference on Computer-Supported Cooperative Work - CSCW'92 (Oct. 31-Nov. 4, Toronto, Canada). ACM, N.Y., 1992. pp. 107-114.

8. Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M. and Richards, J.T. Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of Babble. In Proceedings of the Conference on Human Factors in Computing Systems - CHI'99 (May 15-20, Philadelphia, PE). ACM, 1999. pp. 72 -79.

9. Fitzpatrick, G., Mansfield, T., Kaplan, S., Arnold, D., Phelps, T. and Segall, B. Augmenting the Workaday World with Elvin. In Proceedings of the Sixth European Conference on Computer-Supported Cooperative Work - ECSCW'99 (Sept. 12-16, Copenhagen, Denmark). Kluwer Academic Publishers, Dortrecht, NL, 1999. pp. 431-450.

10. Fogarty, J., Hudson, S.E. and Lai, J. Examining the Robustness of Sensor-Based Statistical Models of Human Interruptibility. In Proceedings of the Conference on Human Factors in Computing Systems - CHI 2004 (Apr. 24-29, Vienna, Austria). ACM, 2004. pp. 207-214.

11. Gaver, W.W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K.A. and Buxton, W. Realising a Video Environment: EUROPARC's RAVE System. In Proceedings of the Conference on Human Factors in Computing Systems - CHI'92 (May 3-7, Monterey, CA). ACM, N.Y., 1992. pp. 27-35.

12. Gross, T. Ambient Interfaces in a Web-Based Theatre of Work. In Proceedings of the Tenth Euromicro Workshop on Parallel, Distributed, and Network-Based Processing - PDP 2002 (Jan. 9-11, Gran Canaria, Spain). IEEE Computer Society Press, Los Alamitos, CA, 2002. pp. 55-62.

13. Gross, T. and Prinz, W. Awareness in Context: A Light-Weight Approach. In Proceedings of the Eights European Conference on Computer-Supported Cooperative Work - ECSCW 2003 (Sept. 14-18, Helsinki, Finland). Kluwer Academic Publishers, Dortrecht, NL, 2003. pp. 295-314.

14. Gross, T., Wirsam, W. and Graether, W. AwarenessMaps: Visualising Awareness in Shared Workspaces. In Extended Abstracts of the Conference on Human Factors in Computing Systems - CHI 2003 (Apr. 5-10, Fort Lauderdale, Florida). ACM, N.Y., 2003. pp. 784-785.

15. Gutwin, C. and Greenberg, S. Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness. In Proceedings of the ACM 1998 Conference on Computer-Supported Cooperative Work - CSCW'98 (Nov. 14-18, Seattle, WA). ACM, N.Y., 1998. pp. 207-216.

16. Gutwin, C. and Greenberg, S. A Descriptive Framwork of Workspace Awareness for Real-Time Groupware. Computer Supported Cooperative Work: The Journal of Collaborative Computing 11, 3-4 (2002). pp. 411-446.

17. Gutwin, C. and Greenberg, S. A Desriptive Framework of Workspace Awareness for Real-Time Groupware. Computer Supported Cooperative Work: The Journal of Collaborative Computing 11, 3-4 (2002). pp. 411-446.

18. Heath, C. and Luff, P. Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. In Third European Conference on Computer Supported Cooperative Work (Amsterdam). Kluwer, 1991. pp. 65-80.

19. Ishii, H. TeamWorkStation: Towards a Seamless Shared Workspace. In Proceedings of the Conference on Computer-Supported Cooperative Work - CSCW'90 (Oct. 7-10, Los Angeles, CA). ACM, N.Y., 1990. pp. 13-26.

20. Ishii, H. and Ullmer, B. Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In Proceedings of the Conference on Human Factors in Computing Systems - CHI'97 (Mar. 22-27, Atlanta, GA). ACM, N.Y., 1997. pp. 234-241.

21. Patterson, J.F., Day, M. and Kucan, J. Notification Servers for Synchronous Groupware. In Proceedings of the ACM 1996 Conference on Computer-Supported Cooperative Work - CSCW'96 (Nov. 16-20, Boston, MA). ACM, N.Y., 1996. pp. 122-129.

22. Prinz, W., Graether, W., Gross, T., Kolvenbach, S., Klein, K.-H., Pankoke-Babatz, U. and Schaefer, L. TOWER: Presenting Activity Information in a Theatre of Work. In Supplement Proceedings of the ACM 2002 Conference on Computer-Supported Cooperative Work - CSCW 2002 (Nov. 16-20, New Orleans, LA). ACM, N.Y., 2002. pp. 91-94.

23. Roseman, M. and Greenberg, S. Building Real-Time Groupware with GroupKit, A Groupware Toolkit. ACM Transactions on Computer-Human Interaction 3, 1 (Mar. 1996). pp. 66-106.

24. Schaefer, L., Pankoke-Babatz, U., Prinz, W., Fatah gen. Schiek, A. and Oldroyd, A. DocuDrama. Virtual Reality 7, 1 (Dec. 2003). pp. 43-53.

25. Siewiorek, D., Smailagic, A., Furukawa, J., Moreveji, N., Reiger, K. and Shaffer, J. SenSay: A Context-Aware Mobile Phone. http://www-2.cs.cmu.edu/~aura/docdir/sensay_iswc.pdf, 2004. (Accessed 11/5/2004).