

Computational Offloading: Supporting Distributed Team Working Through Visually Augmenting Verbal Communication

Yvonne Rogers (yvonner@cogs.susx.ac.uk)

Harry Brignull (harrybr@cogs.susx.ac.uk)

Interact Lab, School of Cognitive and Computing Sciences,

University of Sussex, Brighton BN1 9QH, UK

Abstract

Distributed team working often involves close-knit groups collaborating over a large geographical space performing time-critical tasks. We present a field study of the way a dispersed team of technicians coordinate their work, highlighting the phenomenon of extraneous ‘detective work’ – where much communication, via walkie-talkies, needs to take place to resolve uncertainty arising in their work. We suggest one way of improving the way team members maintain their awareness of what is going on in different places and times is to offload some of the computation involved, by augmenting the verbal channel with visual information. Using the external cognition framework, we describe how we designed a dynamic visualization that allowed salient verbal information to be re-represented as an external cognitive trace. To test our assumption about externalization and computational offloading, we carried out an experiment, with three different conditions: visualization, pen and paper and no cognitive aid. Our findings showed that allowing users to create and view a dynamic visualization improves awareness of what is going on and the way distributed work is coordinated.

Introduction

Distributed team working often involves close-knit groups having to closely coordinate with each other and be constantly aware of each other’s movements, whereabouts and activities. An example is a security team who has to manage and coordinate a large event (e.g. a presidential visit, a conference, a football game) by roaming a large geographical area (e.g. an airport, a convention center, an arena) while coordinating their whereabouts and movements with each other and a central control base. To keep in touch with each other in this kind of setting public broadcast systems are used (i.e. walkie-talkies) backed up by cell and landline phones.

While broadcast systems are generally robust and effective for supporting the communication and coordination needs of such nomadic teams, they do have their problems. Sometimes members miss or mishear messages and consequently do not move to where they are needed or do what is required at a given time. Other times, especially when a lot is happening,

discrepancies can arise between what the different members understand to be the current situation. Establishing common ground can require much confirming and updating by the various team members (Rogers, 1994).

In this paper, we consider how to support the coordination and communication between members of a distributed mobile team. A field study is presented, that reports on how a team of technicians manage the A/V requirements of a large conference. One of our main observations is the huge amount of extraneous ‘detective’ work that takes place among team members; redundant and time-consuming communication is common, serving the purpose of letting each other find out and keep updated of what is happening. To improve upon this situation, we propose providing an external aid that can reduce the cognitive load involved in keeping updated. To this end, we developed a dynamic interactive visualization that was intended to augment existing forms of verbal communication. Specifically, the visualization was designed to ‘offload’ some of the cognition involved in keeping updated, by enabling users to create, annotate and change a visual representation of what is going on and to use this to refer to, when needing to find information.

To test our assumption we carried out a Wizard-of-Oz experiment, where a distributed problem-solving task was set-up for three different conditions: (i) interacting with the visualization, (ii) using pen and paper, and (iii) a control set-up, with no external aid. Our hypothesis is that visually augmenting verbal communication, through providing team members with the ability to track, represent and update salient events can reduce cognitive load and, in so doing, make the coordination task easier to accomplish. We further propose that interacting with a visualization is a more effective form of externalization than pen and paper, because the interactivity and form of graphical representation provide more powerful ways of aiding cognition.

Theoretical Background

There has been a growing interest in how external representations aid cognition, and, in particular, how

they can reduce the cognitive effort required to solve problems (e.g. Hutchins, 1995; Norman, 1993). A main concern has been to examine the interplay between 'knowledge in the world' and 'knowledge in the head' (Norman, 1988; Vera and Simon, 1993). Analytic frameworks that account for the functional role external representations play in relation to internal representations include Zhang and Norman's (1994) distributed representations model and Wright et al's (2000) resource model. Collectively, this body of work has been called the *external cognition* approach (Card et al, 1999; Rogers, in press), and is central to the *distributed cognition* approach (Hollan et al, 2002) – although the latter is, generally, considered to be a more encompassing cognitive science theory (Hutchins, 1995).

Our approach to external cognition (Scaife and Rogers, 1996) is based on conceptualizing how different kinds of graphical representations (e.g. diagrams, animation, multimedia) are used during cognitive activities. Our framework presents a set of core properties, of which the central one is computational offloading; this refers to the extent to which different external representations vary the amount of cognitive effort required to carry out cognitive activities. It is operationalized in terms of specific forms of offloading, including rerepresentation, graphical constraining and temporal constraining. The concepts provide a way of analyzing the cognitive benefits of using different external representations. They have been further operationalized in terms of a set of *design dimensions*, intended to be used to inform the design and selection of representations when developing cognitive aids or interactive systems.

The particular kind of external representation we are interested in here is *visualization*; this is typically a dynamic form of graphical representation, intended to help individuals carry out complex cognitive tasks, such as forecasting trends and spotting patterns in masses of data (e.g. Card et al, 1999; Hegarty, 2002; Johnson and Shneiderman, 1991). One of its assumed benefits is its ability to amplify cognition (Card et al, 1999). This is achieved through visualizing the relationships between various features of the data being depicted – something that is near impossible to work out, unaided, with raw numerical data – through exploiting human perceptual mechanisms.

Our interest in visualization here is concerned with how it is combined with interactivity (cf Kirsh, 1997). We claim that by adding interactivity to a visualization, more cognitive benefits can potentially be gained; by this we mean actively involving the users in creating, annotating and changing the visualizations (as opposed to only viewing those created by a computer). We argue, that in doing so the users are engaged more in the process of externalizing; keeping them more in the

loop and allowing them to remember more and do more problem-solving. It is well known that externalizing one's thoughts through annotation, tracing and other methods can greatly assist in problem-solving, be it individual or collaborative (e.g. Cox and Brna, 1993; Lee et al, 2001). In addition, allowing users to construct their own representations enables them to lay out information in ways that can help them derive a solution and know what to do next (Reisberg, 1997). User involvement has also been found to radically improve the readability of the visualizations (e.g. Dix and Ellis, 1998; Tweedle, 1997).

The goal of our research is to investigate whether computational offloading – in the form of externalizing through interacting with a dynamic visualization – improves problem-solving.

Field Study of the Coordination Among a Team of Technicians

In order to understand how distributed teams keep in touch and the communication problems that can arise during their work, we initially carried out an ethnographic study of nomadic team working. Two of us shadowed, interviewed and joined an audio/visual (A/V) team of technicians who were responsible for organizing a large international conference, prior and during the event, that was held at the Seattle Convention center. We took notes on the fly and having got permission, also recorded conversations that took place face-to-face and via walkie-talkies, mobile and landline phones, using tape recorders and a scanner. We also interviewed a number of conference organizers and observed technicians at work at other conferences.

Firefighting

The various events that take place at a conference, such as tutorials, workshops, paper sessions and demos, have all been meticulously planned in advance, as to how they are to be run, what A/V support is required and how to set up the rooms. Typically, a data projector, microphone, VCR, overhead projector and screen are set up, although items like wireless mikes, flip charts, cameras, tripods may also be requested. Large conferences can have anything from up to 50 events (rooms) running in parallel each day, each needing setting-up and monitoring. A key objective for the team is to get as much of the equipment set up in advance. Once installed, a main part of their work is to deal with the unplanned events, known as 'firefighting'. Examples range from a speaker demanding different equipment, to equipment malfunctions, to setting up unscheduled events. The number and timing of these 'fires' is unpredictable and many are last minute.

Hence, a key aspect of the AV team's work involves trouble-shooting; dealing with unexpected events in unpredictable places and times. The team needs to be

highly flexible; being able to rapidly group and disband, in different locations depending on the exigencies of the moment. To achieve this, up-to-date information of who is where, who is doing what, what needs to be done, etc., needs to be continuously and effectively relayed between the distributed team members.

Communication Streams and Misalignments

There are many different routes by which problems are detected, reported and subsequently dealt with by the A/V team. Speakers, student volunteers, conference organizers, attendees and the technicians themselves may all be the first to detect a problem. How it is communicated and dealt with can then follow a number of pathways. For example different speakers have different ways of making their demands and complaints known. Some may just tell an A/V technician or a conference organizer and leave it at that, whilst others will tell everyone in sight. Thus, excessive communication can take place, even for dealing with fairly minor problems.

A problem that can arise is for certain pieces of information not to be communicated at the appropriate time or, alternatively, faulty pieces passed on instead. When discovered, uncertainty can kick in requiring the technicians to engage in detective work. Typically, they will try to repair the situation, by backtracking and repeating steps, for example as illustrated in the following excerpt:

S (conference organizer in conference center over the walkie-talkie {WT}): Yeah, R, we're having some audio/visual problems in room 608, they said that there is a buzzing...

A (technician in control center speaks to B): Is that the room we already fixed?

B (technician in control center replies to A): Yeah

R (A/V manager roaming, replies to S over WT): How old is that report?

S <reply is inaudible to all>

R (over WT): yeah, I believe 608 is already fixed.

S (over WT): Umm, someone's been and said that the problem was worked on but the problem was still occurring...

R(over WT): Umm, A and B head up to the 6th floor.

A (over WT): That's the room we already fixed but we'll go check

Many of these communication problems arise because of the uncertainty of what is the current state of an activity or event. The technicians are constantly dealing with multiple jobs and tasks, while speakers and attendees may not know something has been sorted and continue to relay their concern to the conference center. Often the technicians will spend much time working out what is currently true from the disparate representations and messages they receive. When dealing with such uncertainty they will tend to err on the cautious side, choosing to do additional monitoring and checking, that can end up being very time-consuming, frustrating and unnecessary.

Offloader: An Updating and Tracking Visualization

Our analysis highlights the many ways verbal updating takes place. A high degree of mental tracking is needed to work out what needs to be done, who is doing what and who knows what. Transient verbal information is used to do this mental juggling; making it difficult to get and maintain a full picture of what is going on. We argue that such *nomadic awareness* can be improved if salient aspects of the verbal information are transformed into a more permanent visible trace. In so doing, the updating and tracking can be supported both by talking and an external display, updated in an immediate way.

To this end, we designed a visualization called Offloader. It was based around the activity of managing a set of problems, i.e. what is happening where, who is dealing with what and how urgent it is. Our aim was to design a visualization that supported computational offloading through active user externalization, and in so doing, to allow users to readily integrate various pieces of information needed to make rapid decisions.

Graphical Representation

Much human factors research and guidelines exist for informing the layout of displays for time-critical tasks, where rapid decision-making is involved. We took into account the proximity-compatibility principle (PCP) that promotes the physical co-location and organization of information that needs to be mentally integrated (Wickens et al, 1995; Wong et al, 1998). Specifically, we designed a problem status display, partially shown in figure 1, intended to convey how different dimensions co-vary in relation to each other through integrating and overlaying interdependent representations. The vertical columns represent locations in a building, while the horizontal axis represents time. A dynamic timeline moves down the display to show where the problems – labeled as jobs – are in relative time: job strips stretch down along with the timeline, as it moves down the display. This shows at a glance how many jobs are outstanding and how long they have been running for. Each job has a default identifier, which is the room it is located in. The color of a job strip also increases in intensity the longer it remains. An alert sign flashes when a job reaches a critical state. Once a job is completed its state changes; indicated by its color fading and the strip no longer moving with the timeline.

To support user externalization we provided interactive building blocks to represent actions and objects that were considered central to planning and coordination (i.e. problems, people). A new job strip (i.e. problem) is created by the user selecting a color coded icon from a palette and dragging it over to the appropriate column in the job status display area. The

user can add further details, by typing comments into the job strip. Team members are represented by highly distinguishable cartoon icons, with names underneath; again, these are dragged off a palette and overlaid on a job strip or any of the other columns, providing a way of representing that person's current or planned location. This way of superimposing representations of different dimensions also allows users to see at a glance who is working on what and where all the team members are – something which is very hard to do just 'in the head'.

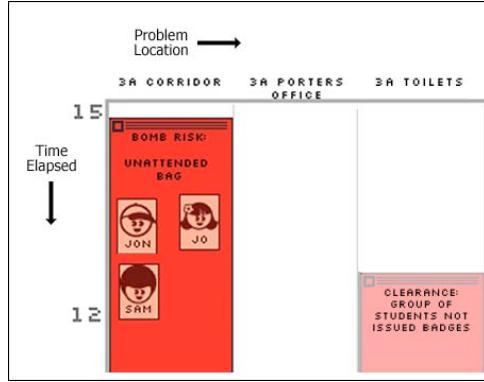


Figure 1: Part of the dynamic visualization, showing technicians superimposed on job strips as part of a chart, with location and time elapsed as its axes

Evaluating Offloader

To test our assumptions about computational offloading and externalizations we carried out an experiment that investigated whether being involved in the creation and manipulation of an interactive visualization improved distributed problem-solving. Three conditions were set up: (i) externalization through interacting with the Offloader visualization, (ii) externalization through the use of pen and paper and (iii) no externalization. A between subjects design was used: 8 participants per condition.

Method. To simulate a nomadic work setting we used the Wizard of Oz method. A scenario was devised where a security team was responsible for checking the security of a university building prior to a visit by a VIP. The participants (including real supervisory security guards/ porters for the building) were asked to imagine they were in charge of running this event. Specifically, they were asked to:

- allocate their team to roam certain parts of the building to deal with any reported incidents in a manner that maximizes their productivity
- keep track of all reported incidents to check they are being dealt with
- ensure that all incidents are dealt with before the VIP arrives

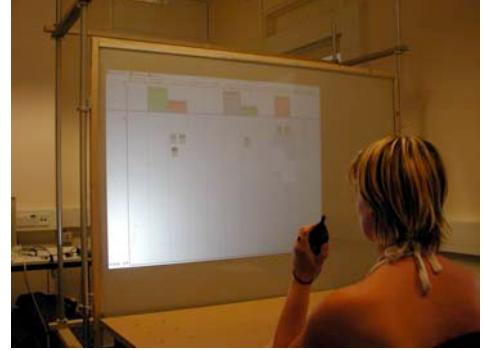


Figure 2: Photo of the Offloader visualization in use

Six other stooges were asked to pretend to be the security guards roaming the building. They sat together in a room away from the participant, communicating with him or her via a walkie-talkie. They were asked to follow a script, detailing a sequence of incidents to be reported at specific times, which they were supposed to have discovered while roaming the building. These ranged from minor events (e.g. coffee spilt in corridor) to severe events (e.g. suspicious package in the area). The script was written in increasing complexity, with more incidents happening in parallel towards the end of the study. The stooges were also required to report, via the walkie-talkie, what they were doing (e.g. switching jobs, completing jobs). The participant could also communicate, using the walkie-talkie, with any of the stooges at any time to allocate jobs and find out what they were doing.

The Offloader visualization was back projected onto a large vertical display in front of the controlling participant (see Figure 2). In the other two conditions the participants sat at the same desk, with either pen and paper or nothing. The stooges and the controlling participant each had a walkie-talkie. All sessions were videoed and the participants interviewed afterwards.

Results A main finding of the study was that the participants who used the Offloader visualization found it easier to plan, manage and coordinate the problem-solving task compared with the participants in the other two conditions. This was most marked when several critical events had to be attended to in parallel, especially towards the end of the session, as shown by the way the participants planned while communicating. In the Offloader condition, the participants interacted 88% of the time with the visualization while talking to the stooges over the walkie-talkie. Creating new jobs was effortless to do while talking as was coordinating team members and deciding where to send them. The video data showed the participants often moving the people icons around when talking to the stooges, moving them around on the job strips (cf. to the way players move their Scrabble pieces around when trying

to work out the best word). In contrast, in the pen and paper condition parallel talking and writing was less frequent (54%). A common strategy was for participants to jot down a few notes while talking over the walkie-talkie, and then once they had finished talking, to re-work them into a representational format (e.g. table, list, diagram). Hence, the planning was done sequentially. When making changes, the participants in this condition typically crossed out what they had previously written and wrote the new information elsewhere. The longer the study went on, however, the more difficult it became for them to make out what of their notes was relevant. Having so many crossing-outs also made it difficult for them to organize the latest information and give it a readable structure. When interviewed, several commented on this problem, e.g., "When it got busy my notes did get confusing... there's only so much you can write on a page" and "It started out well but my table [...] ended in a bit of a mess".

To examine in more detail what kind of communication went on across the conditions, we looked at the kinds of utterances spoken by the participants across the conditions. Utterances were classified in terms of four main types:

- (i) *instructing* others to do something e.g., "go to room 4 and help Joe"
- (ii) *monitoring* ongoing progress, e.g., "How's the clean-up going?"
- (iii) *committing errors* e.g., where stooges were sent to the wrong location or where they were assigned jobs when they were already busy
- (iv) *requests and clarification* where the participant wasn't sure what was going on, e.g., "what's happening over there?" and "haven't you done that already?"

The first two categories give an indication of how well the participants are managing their task. The last two give an idea of when the participants are not coping well, indicating the need for more communication to build up a picture of the current state of affairs. Specifically category 3 shows the amount of incorrect decisions made by the participants about what to do, while category 4 indicates the repair work the participants carried out; working out who is where, why something has not happened, when it should have and so on. (Other kinds of talk, e.g. banter are not included in this analysis).

Figure 3 shows the means and standard deviations of the four communication types across the three conditions. A main difference is the pattern of the communication types. Far more monitoring utterances were made when using the Offloader condition (mean = 9.25) than for the other two conditions (means = 7.75 and 4.37, respectively). A similar trend was found for instructing; far more commands were made in the Offloader and the pen and paper conditions (means = 11.83 and 11.13, respectively) compared with the no

externalization condition (mean = 4.34). This suggests that the participants in the two externalization conditions were much more in control of their task.

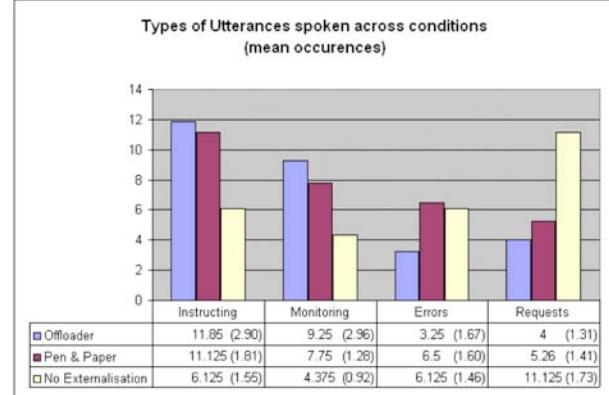


Figure 3: Breakdown of utterance types (mean and S.D) across the three conditions: Offloader, pen and paper and no externalization.

Although there was no difference between the two externalization conditions for number of instructing ($t=0.62$, ns) and monitoring utterances ($t= 1.31$, ns) there was a difference between the number of errors made: the Offloader condition exhibited significantly less errors than the pen and paper condition ($t=3.97$ $p<0.001$). This suggests that the participants in the pen and paper condition were more likely to misread their representations and make inappropriate plans than in the Offloader condition. As might be expected, the participants in the no externalization condition made a higher number of errors than the Offloader condition ($t=3.67$ $p<0.001$). They also spent considerably more time asking the stooges questions about what was happening compared with either the Offloader ($t=9.30$ $p<0.000$) or pen and paper conditions ($t=6.98$ $p<0.000$). This contrast in coping strategy seemed to be most acute when the participants were required to respond to an immediate request for help from one of the stooges. To do this they have to work out whom they can send straight away from their team. In the Offloader condition, the participants 'read off' from the visualization where all their team members currently were and then worked out who was free or least busy to send. In the no externalization condition the participants had no such cognitive aid, and so resorted to calling up on the walkie-talkie each of the team members, to determine whom to assign. This took much longer; sometimes the participants ended up calling the same team member again because they'd forgotten where they were. Likewise, in the pen and paper condition towards the end of the session, the participants found it increasingly more difficult to work out from their representations where everyone was and so phoned up team members to check what they were doing before making a decision as to whom to send.

Conclusions

As our field study showed, distributed team working involves considerable coordination and ad hoc decision making. Much mental juggling and tracking is demanded of the team members. A compensatory strategy is to engage in extraneous detective work, backtracking and confirming with each other as to what is what. To improve nomadic awareness, we suggested providing a cognitive aid. While visualizations and other displays are commonly used in time-critical work in fixed 'office' settings (e.g. emergency services, dispatch management) they have not been used in 'nomadic' settings, like conferences.

The kind of visual display we proposed for supporting distributed team working was an interactive dynamic visualization, where users co-create an external trace of the ongoing events in conjunction with the system, by rerepresenting salient aspects of the stream of verbal communication. In so doing, information, that otherwise would have had to be rediscovered through verbal 'detective work', is made available through the dynamic visualization. Involving the users in building up the visualization also enables them to keep track of the events while using the external representations to plan. In particular, moving the person icons around while discussing the problem with a stooge, enables the participants to both plan which people to assign to the problem and visualize where they will be next.

We argue that it is the externalization that the participants actively engage in which provides the cognitive benefit, especially when workload is at its greatest (i.e. when multiple problems need to be tracked). Memory load is reduced, allowing for more monitoring, updating and planning to take place at the same time – something that is very hard to achieve without having the opportunity to externalize what has just been heard.

Acknowledgements

This research was carried out as part of the Dynamo project, grant no. GR/N01125 awarded by the EPSRC, UK. The authors gratefully thank Rene Audette and the rest of the AVHQ team for their cooperation, Mia Underwood and Greta Corke for their graphic designs and to our partners on the Dynamo project, especially Tom Rodden.

References

Card, S., Mackinlay, J. & Shneiderman, B. (1999). *Readings in information visualization*: San Francisco, CA: Morgan Kaufmann.

Cox, R. & Brna, P. (1993). Analytical reasoning with external representations. *Proceedings of the AI-ED93 Workshop on Graphical Representations, Reasoning and Communication*. (pp. 33-36).

Dix, A. & Ellis, G. (1998). Starting Simple - adding value to static visualization through simple interaction. *Proceedings of Advanced Visual Interfaces - AVI98*, (pp. 124-134) New York: ACM.

Hegarty, M. (2002). Mental visualizations and external visualizations. *Proceedings of the Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hollan, J., Hutchins, E. & Kirsh, D. (2000). Distributed Cognition: toward a new foundation for human computer interaction research. *Transactions on Human-Computer Interaction*, 7(2), 174-196.

Hutchins, E. (1995). *Cognition in the wild*, Cambridge, MA: MIT.

Johnson, B and Shneiderman, B. (1991). Tree-Maps: A space-filling approach to the visualization of information structures. *Proceedings of Information Visualization '91*, (pp. 275-282), IEEE.

Kirsh, D. (1997). Interactivity and Multimedia Interfaces. *Instructional Science*, 25, 79-96.

Norman (1988). *The psychology of everyday things*. NY: Basic Books.

Norman, D. (1993). *Things that make us smart*. Addison-Wesley.

Reisberg, D. (1997). *Cognition: Exploring the science of the mind*. New York: Norton.

Rogers, Y. (1994). Exploring obstacles: Integrating CSCW in evolving organizations. *Proceedings of CSCW'94*, (pp. 67-78). NY: ACM.

Rogers, Y. (2004). New theories for HCI. To appear in *Annual Review of Science and Technology*, 38.

Scaife, M. & Rogers, Y. (1996). External Cognition: how do graphical representations work? *International Journal of Human-Computer Studies*, 45, 185-213.

Tweedle, L. (1997). Characterizing Interactive externalizations *Proceedings of CHI'97*, (pp.375-382). NY: ACM.

Vera, A.H. & Simon, H.A. (1993). Situated Action: A Symbolic Interpretation. *Cognitive Science*, 17, 7-48.

Wickens, C.D. & Carswell, C.M. (1995). The proximity compatibility principle: it's psychological foundation and relevance to display design. *Human Factors*, 37(3), 473-479.

Wong, W., O'Hare, D. & Sallis, P.J. (1998). The effect of layout on dispatch planning and decision making. In *People and Computers XIII HCI 98 Conference*, Sheffield, UK; Springer.

Wright, P., Fields, R. & Harrison, M. (2000). Analyzing Human-Computer Interaction As Distributed Cognition: The Resources Model. *Human Computer Interaction*, 51(1), 1-41.

Zhang, J. & Norman, D.A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, 18, 87-122.