

Designing Dynamic Interactive Visualisations to Support Collaboration and Cognition

Yvonne Rogers, Harry Brignull and Mike Scaife*
Interact Lab,
School of Cognitive and Computing Sciences,
University of Sussex,
Brighton BN1 9QH
UK
yvonner, harrybr@cogs.susx.ac.uk

Abstract

Dynamic interactive visualisations (DIVs) are intended to help coordination and collaboration, through augmenting existing forms of synchronous communication (i.e. phones, face to face, walkie-talkie). A central feature of a DIV is active user involvement: users are required to create, annotate and change the information visualisation to represent the changes in the activity space they are concerned with. A main benefit of doing so is to enable users to externalise and offload some of the cognitive effort involved in problem-solving, by laying out information in ways that can help them derive a solution and know what to do next. In this paper we describe how we went about designing a DIV to support nomadic team working. We begin by describing our experimentation in designing a DIV. We then show how our computer-based DIV substantially improved performance for a complex collaborative task, which involved much communication and cognition.

Keywords

Dynamic information visualisations, external cognition, graphical representations, cognitive amplification, broadcast communication, collaboration, team working

1: Introduction

Traditionally, information visualisations (IVs) have been designed to help individual users carry out complex

cognitive tasks, such as forecasting trends and spotting patterns in masses of data [1, 11]. The graphical representations are intended to ‘amplify cognition’ ([1], p.6) through showing relationships between various features – something that is near impossible to do with the human eye using only raw numerical data. They do so largely through exploiting the benefits of differing graphical formats based on human perceptual mechanisms [8] and cognitively-based principles (e.g. [13]).

In contrast, the relatively new area of collaborative information visualisation environments (CIVEs) has a different focus: supporting groups of people, who are distributed over time and space, to interact remotely. In particular, a main objective is to enable such people to collaboratively access information and be able to communicate with each other about this. To this end, CIVEs are intended to be more extensive than IVs: providing visualisations and other sources of information to enable users to carry out their tasks while also supporting communication with each other *through* them. An example of a CIVE is Babble [6] which was designed to provide a dynamic visualisation of the participants in an ongoing chat-room conversation alongside a text stream of the conversation itself.

The potential benefits of these kinds of hybrid IVs is that collaboration, information access and decision-making can all be facilitated. To enable this to happen, however, requires that the CIVE be designed to *support* both cognitive and collaborative processes. In other words, they need to be both cognitive amplifiers and

* Mike Scaife died suddenly and unexpectedly while we were revising the paper. His contribution to the research was instrumental.

communicative facilitators. Most CIVEs, however, have so far been concerned with how to support synchronous communication between remote participants through using various chat facilities, with visualisations being provided more as an adjunct. Our approach to designing CIVEs is different: we are concerned with enabling nomadic people, who are remotely located, to work together by providing *dynamic interactive visualisations (DIVs)* that can be used to coordinate and manage their collaborative work needs, through augmenting their existing ways of communicating (i.e. via face-to-face, broadcast systems and phones). In other words, our focus is on supporting existing forms of communication through providing new forms of cognitive amplification. A main reason for this difference in focus is that we are interested in supporting people ‘on the move’ who have to closely collaborate with each other ‘there and then’ rather than ‘stationary’ people who collaborate over time from separate offices. Hence, the needs for information access and communication are quite different.

Central to our approach is user interactivity. In particular, we believe it is important to let the users be very much involved in creating, annotating and changing the information visualisations. This is contrary to the way most information visualisations have been designed, where the system creates and changes the representations using various algorithms to represent the underlying data. The rationale for our approach is based on the premise that actively involving users enables them to externalise and offload some of the cognitive effort required to do their work [21]. It is well known that externalising one’s thoughts through annotation, tracing and other methods can greatly assist in problem-solving, be it individual or collaborative (e.g. [4, 10, 18]). 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. It also has been found to radically improve the readability of the visualisations (e.g. [5, 24]).

However, there are obviously advantages of having the system create the visualisations automatically and then update them dynamically to reflect changes in the underlying information. A key question is what kind of user interactivity is optimal? In some ways, this question can be regarded as a form of the well-known Human Factors concern to do with person/machine allocation. We would argue that it is a more complex issue here, in that it is not just about letting users and machines do what they are best at, respectively, but concerns the interplay between a number of interdependent factors. These include deciding what are appropriate kinds of interactions, that will enable users to externalise their thoughts effectively, and how these will interact with system-created visualisations, intended to support

awareness, collaborative planning and decision-making activities.

Thus, a central concern of our research is to determine (i) how to design a CIVE that gives users a set of interactive building blocks to externalise with, and (ii) system-created dynamic visualisations, that when combined with (i) can support users in their collaborative problem solving tasks. Specifically, this paper describes how we designed a DIV to support nomadic team working. By nomadic team working we mean close-knit groups collaborating over a large geographical space performing time-critical tasks. Typically, such groups of people have to closely coordinate with each other and be constantly aware of each other’s movements, whereabouts and activities. An example is a technical support team who have to manage and coordinate a large event (e.g. a conference) by roaming a large geographical area (e.g. a convention center) while coordinating their whereabouts and movements with each other and a central control base. The aim of our CIVE was to improve this kind of nomadic awareness, by allowing dispersed users to re-represent verbal communication as DIVs on interconnected wireless handheld devices and a large fixed display.

The first part of the paper describes how we went about designing our DIV, focusing on the various alternative representations that were experimented with to show how and why we came up with our overall design. The second part of the paper discusses the findings of an empirical study we carried out to assess the extent to which our DIV was able to both amplify cognition and facilitate communication.

2: The Problem Space

The domain space we chose was distributed nomadic team working, where members often have a difficult time keeping track of each other, of what co-workers are up to and what needs to be done. Typically, such teams work in settings where much on-the-fly problem-solving takes place, e.g. emergency control teams dealing with unexpected events like fires and accidents. This kind of work requires a lot of monitoring and updating of other team members, relying on the verbal channel as the primary means of support (e.g. [9]). We were interested in seeing if some of this monitoring and coordination work could be ‘offloaded’ onto a DIV.

To begin, we carried out an in-depth ethnographic study of a technical team who are responsible for installing and maintaining the audio-visual (A/V) equipment at conferences. Detailed findings of the study are reported elsewhere (19). Here we summarize some of the key issues.

A lot of the technicians' work involves troubleshooting and dealing with unexpected events in unpredictable places and times. The team need 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 takes place through a combination of walkie-talkies, mobile phone and opportunistic face-to-face meetings.

One of the main difficulties confronting nomadic workers trying to keep in touch of the latest happenings via these communication channels, is that the amount and pace of the changing information can sometimes be overwhelming. Moreover, different messages about the same event can be passed on by different people, some of which have been superseded by newer events. This requires the workers to engage in much mental juggling of information, trying to maintain a coherent picture of the latest, i.e. what is happening, where everyone is, where the origin of certain messages came from and so on. Inevitably, there are times when old or faulty information is accepted as representing the current state of affairs, resulting in wrong decisions being made about what to do next. Much redundant verbal repair work then has to take place to enable the members to realign their understanding of what is going on [20] – all of which takes up valuable time. Hence, much time is spent by the technicians working out what is currently true from the disparate representations and messages they receive of what is happening and what is meant to be happening. How might we design a CIVE to reduce this?

2.1: Transforming our analysis into design needs.

One of our main proposals, based on the findings of the study, was that nomadic teams involved in on-the-fly problem solving tasks could benefit if provided with the opportunity to create an external memory or 'cognitive trace' [21] of all the on-going problems, both past and present. In particular, we assumed that it could unite the team members' disparate pieces of knowledge into an integrated and up-to-date shared DIV of the current state of affairs.

To achieve this, we needed to find a format that could represent the important information that the technicians constantly needed to know, share and update, in a visually effective and appealing way. We began by conceptualising the kind of information needed to manage and monitor unexpected problems in this kind of setting. This included the following set of inter-related parameters:

- Problem type and details
- Urgency of problem
- Spatial location of team members at time of problem reported
- Status of problem (e.g. being solved / on hold)
- Team member attending problem
- Time problem reported
- Time elapsed since problem started
- Time until deadline

A key design question was how to provide a set of graphical components that the users, in conjunction with the system, could build up into a visualisation that represented the parameters in a way that was meaningful and easy to understand. To begin we looked at what would be the most effective kinds of information visualisations that the system could depict.

2.2: Designing an information visualisation based on a concrete composite representation

We explored the possibility of using Chernoff faces [2], as an information visualisation to represent the variables. These were originally developed in the field of statistical psychology for conveying multidimensional data. Chernoff faces are iconic objects which allow the features to represent changes in the values of different parameters (see Figure 1). Changes in facial expression are then read as changes in data points. For example, the raising of the eyebrows and the extent to which a mouth smiles could indicate a change in the trend of the data (rising, decreasing, etc.). The main appeal of Chernoff faces as a candidate visualisation format here is that humans have a strong ability for recognizing facial expressions and thus we might detect multiple data values just as we do multiple facial expressions [16, 17, 25].




Figure 1: A Chernoff face with annotations (from [16])

We began by mapping the core facial features onto a sub-set of our variables. We tried to see if any of the mappings were more ‘natural’ than others. For example, nose size seemed appropriate for representing time elapsed since problem reported as it could be expanded to reflect change in time, e.g. the bigger the nose got, the longer the problem had been running for. Similarly, we used eye size to represent urgency of problem. Hence, at a glance, it could be possible to see how urgent and how long it had been running by checking nose and eye size. In this way, each reported problem could be represented by a different face displayed adjacent to each other. This made it possible to compare the various states of them, allowing users to know which needed most attention and which could be kept on hold. To test our assumptions we carried out a ‘quick and dirty’ user test with four participants. They were instructed on a coding scheme for the faces (eye size represented complexity; nose size represented urgency; mouth size represented time elapsed), and were then given a series of paper sheets containing a selection of the Chernoff faces. They were then asked to select the faces that represented different problem states, starting with individual variables (e.g. ‘Which is the newest/ oldest problem?’), through to interacting variables (e.g. ‘Which problem is the most urgent and most complex?’).

The main finding was that the participants took considerable time to read off the overall state of a problem from a complete face, even for relatively simple combinations. The most likely reason for this is that there are inherent problems with the visualisation format itself. The very fact that faces are holistically perceived may mean that they are not well-suited for parsing into separate variables and then subsequently recombining them to make complex decisions [14]. This is particularly marked when there are a number of simultaneous faces present, increasing target confusion.

2.3: Designing an information visualisation based on an abstract composite representation

The next strategy we employed was to use less socially-potent representations than the face, trading on the human visual system’s capacity for pre-attentive processing – the rapid, automatic detection of basic perceptual features such as size, hue and orientation [8]. Such a method has been used successfully for building effective static and dynamic visualisations (e.g. [3], [8]) To do this we worked in collaboration with a professional graphic designer. We came up with a number of sketches for visualisations, in which we experimented with brightness, colour, size and outline shape. Some dimensions, like size and brightness, are ‘separable’, i.e. they can be attended to selectively,

while others, like the saturation and brightness, are ‘integral’ (together they define colour) and are hard to separate [7].

Using pre-attentive visual features may enhance detection but still leaves the issue of ensuring accurate interpretation of what they might signify within a representation. In particular, we wanted to enable each end of a feature dimension (e.g. very bright and very dull) to be readily associated with different underlying states. We were concerned to find ways of combining these parameters into a holistic representation that would enable identification of the underlying state of the problem. After several attempts, we came up with the metaphor of ‘ripening fruit’, whereby transitions through shape (complexity), colour (urgency) and brightness (time elapsed) resembled the ripening of a piece of fruit. Hence, an unripe piece of fruit was likened to a simple, not urgent and recently identified problem being depicted as an amorphous, pale and light pink object. A ripe piece of fruit was likened to a complex, very urgent and long-time elapsed problem and looked like a spiky, very bright and vivid red object (see Figure 2).




Figure 2: Comparison of the two modes of representation used: concrete (Chernoff face) and abstract (ripening fruit metaphor)

We used the same user testing method as with the Chernoff faces study. Again our findings were disappointing. Even though the participants found it easier to understand the coding scheme and were faster at deriving the underlying problem state using the ‘ripening fruit’ representations, they still reported that they had difficulty making comparisons between them, particularly when comparing multiple parameters (e.g. ‘Which is the oldest and most complex problem?’). Again, it seems that considerable cognitive effort was required to be able to make these complex judgments – something that we were actually trying to reduce. The lesson seemed to be a quite general one: that the use of these kinds of glyph/iconic

formats, while successfully employed in psychophysical research, were unsuitable for supporting the kind of collaborative decision-making process we were designing for, requiring considerable cognitive effort to make comparisons between the representations to understand them.




Figure 3: The spatial encoding of temporal information, superimposed with type of problem and team member location

2.4: Designing an information visualisation based on spatial information

At this stage, we decided to examine the possibility of using multiple representational formats to reflect the relationships among the variables. In particular we took inspiration from Tufte's idea of 'multiple parallelism' whereby *spatial* information is used to enable visual comparisons between objects: "*Parallelism grows from a common viewpoint that relates like to like... [It] is not simply a matter of design arrangements, for the perceiving mind itself actively works to detect and indeed to generate links, clusters, and matches among assorted visual elements.*" ([23]; p.82). We were also guided by the proximity-compatibility principle (PCP) that promotes the physical co-location and organization of information that needs to be mentally integrated [26, 27].

After sketching out several spatial representations, we came up with a simple chart structure using vertical spatial location to encode temporal information and horizontal spatial location to represent the different locations of the problems. We then decided to *superimpose* the other representations onto this spatial structure. These included (i) using different coloured strips filling the bounded space to represent different kinds of problems, and (ii), placing icons over these to show where people were and what they

were working on. Thus we combined different representational formats by layering them on top of each other rather than try to aggregate or combine them as a composite whole (see figure 3). In so doing, it should enable properties to be read off rather than have to be computed [12].

User testing, following the same 'quick and dirty' procedure as before, revealed that participants were able to read off much easier the various parameters and understand the relationships between them. The chart metaphor, being a familiar structure, was also easy to understand and participants knew how to read the information that was overlaid on it. Thus, this kind of 'spatial and superimposition' representational format proved to be the most effective way of representing multiple variables that co-vary and which we decided to use as the basis for our system-constructed information visualisation.

2.5: The design of our DIV

To decide on which part of our information visualisation to make as the interactive building blocks, we needed next to develop a conceptual model of the CIVE we were designing for (see [19] for more details). Briefly, we designed a system called Offloader, which comprises two interlinked components: a mobile 'pocket-loader' part, intended to be used by technicians on the move and a fixed 'wall-loader' part, intended to be interacted with in a control center, by a manager or controller. An underlying objective was that the interlinked components would be used by the different technicians to augment their existing way of working and communicating: the wall-loader helping controllers manage the nomadic team workers and the pocket-loaders helping the nomadic team workers key in latest information to be relayed to the control center, while also being able to find out latest updates, from information keyed in by other technicians.

2.6: The system-created components of Offloader

For the wall-loader, three interlinked views of information about the setting were designed; the current plan views of the day's events in the upper left corner, a topographical floor plan of the building in the lower left side and a 'job status' display, showing problems as moving strips on the right side. Hence, different kinds of information were provided on the same fixed display, enabling the users to compare existing plans of the day's events with overlaid representations of the unexpected events as they unfolded. For the pocketloader, information about problems was designed as minimalist representations (taking into

account the small size of handheld displays), in the form of text-based to-do-lists and problem reports. Up-to-date information about the availability and progress of jobs was also provided in this format.

2.7: The interactive building blocks of Offloader

To support user ‘externalisation’ we decided to use a familiar GUI interaction mode of ‘drag and drop’. Actions that were considered to be central to the planning and coordination of the user’s task were chosen as building blocks. For example, a new job strip is created by the user selecting a color-coded icon from a palette on the left side of the screen and then dragging it over to the appropriate column in the job status display area. The user can then add further details, by typing in comments into the job strip (see figure 3). Team members are also represented by easy to identify cartoon icons of people’s heads, with their names underneath (see figure 4); again these are dragged off a palette and overlaid on a job strip or any of the other columns, to represent that person’s current or planned location. Moving the icons around can also help users optimize their plans of which people to designate to problems – in the same way that physically moving Scrabble [22] pieces around helps when trying to construct the best word.




Figure 4: Examples of the labeled cartoon icons used to represent team members

The constraints on the pocket-loader design were far more restrictive, especially taking into account the problem of interacting with a handheld device while talking on the walkie-talkie and moving around. Hence, we designed the mode of interaction to be minimalist, requiring only a small number of taps on the screen to enter data, via the use of predefined menus and pre-ordered forms. Our aim was to provide a quick and easy way of entering and looking up information.

2.8: Creating the Dynamic Interactive Visualisation

Having decided which parts of the visualisation that the user needs to build to help create the information visualisation, we then needed to determine how it should

evolve to reflect ongoing changes. This required deciding which components of the IV to then make dynamic. We decide that certain aspects should always remain fixed and constrained, such as the location of events in relation to each other, to reflect their spatial location in the real world while others should be transformed into ‘live’ objects to show they were changing. The changes were to be shown through using various forms of dynamic animations. Examples of parameters which were animated included time, urgency and termination. For example, a dynamic time-line moved down the display to show change in time. Job-strips were also stretched down by the system to move with the descending time-line. They were also designed to change in intensity of colour to reflect the change in urgency of a job. The more intense the colour, the more urgent the job had become. Once a job had been completed, the job-strip could then be deactivated by the user by clicking a button inside it. This had the effect of returning it to a static object, no longer stretching with the time-line and also fading in colour. Redundant coding was also used to emphasize critical events: accompanying sound and flashing were used alongside a job-strip that needed immediate attention.

3: How effective were our Dynamic Interactive Visualisations?

To assess the potential benefits of our DIV to amplify cognition and facilitate communication we carried out a more extensive user study evaluating the core features of Offloader. In particular, we wanted to determine whether:

- The various interactive building blocks would be used by the users to externalise their planning and in so doing reduce cognitive effort
- Our conceptual model of having static and live objects was effective at enabling users to readily understand multi-dimensional information and rapidly act upon this
- Our DIV was successful at augmenting broadcast-mediated communication
- Our DIV helped with collaborative decision-making

We decided to focus our evaluation of the DIV designed for the wall-loader component, as this involved a more complex mode of interaction between the system and the users, enabling us to examine in more detail the above assumptions. We devised a study where participants had to manage an event, via a broadcast system in conjunction with using wall-loader. Specifically, they had to imagine they were in charge of a security team responsible for checking the security of a building prior to a visit by a VIP. They were required:

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

Six other stooges ('roamers'), were asked to pretend to be the security guards roaming the building. They were required to sit together in a room away from the participant, communicating with him/her via a walkie-talkie. They were asked to follow a script, detailing a sequence of incidents to be reported by the various stooges 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 session. 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 wall-loader simulation was set up in front of the controlling participant and projected onto a large vertical display. Six sessions were run with different participants. A control condition was also carried out with six further participants doing the same task, but without access to the Offloader system. Instead they had the choice of creating their own representations, being provided with pens, paper, printed floor plans of the building and stickies. All sessions were videoed and the participants interviewed afterwards.

3.1: Main findings:

A main finding was that the participants who used the wall-loader prototype found it much easier to plan, manage and coordinate their work compared with those who had to create their own representations by hand. This was most marked when several critical events had to be attended to in parallel. Having co-created the DIV with the system and subsequently having it update dynamically changes in the activity space, substantially reduced the need for the participants to call up the various roamers on the walkie-talkie to check what they were up to. In contrast, the group who had to create their own externalisations spent more time on the walkie-talkie trying to find out who was where and who was doing what.

Participants in the Offloader condition were also able to quickly develop an efficient and systematic procedure

of externalising what was being relayed to them verbally into an external representation of the state of affairs. Hence, there did not seem to be any problems of interference between using the two modes or any translation overheads, even when the workload was high. In contrast, the own externalisation participants often hurriedly scribbled down notes, as they attempted to re-represent what was being relayed to them over the walkie-talkie, while at the same time deciding what to do. When trying to make changes to existing activities, they would typically cross out what they had previously written and write the new information elsewhere. The longer the study went on, however, the more difficult it was for the participants to make out what they had written was relevant. Several participants could be seen to be 'hovering' their pen over their externalisation, trying to work out what connected with what. Having so many crossing-outs also made it difficult for them to organize the latest information and give it any structure. This made it much harder for them to see connections between their different externalisations. In their post task interviews several commented on this problem, e.g., one participant said "When it got busy my notes did get confusing... there's only so much you can write on a page." As a compensatory strategy, they would often spent more time asking the stooges questions about what they were doing.

A main way of updating information using Offloader was through 'dynamic replacement'; changes were made to the display, by new information being externalised by the participant (e.g. the icon of Joe was placed on a new jobstrip, showing he has moved to room 206), with the effect of removing 'old' information (the icon of Joe was no longer on jobstrip 305, indicating he had moved from there). In addition, the dynamic features of the Offloader visualisation were also useful at guiding the participant's attention to the necessary information for them to make decisions. For example, from the video data, it can be seen that the participants often moved the people icons around when talking to the roamers over the broadcast system, sometimes placing them back and moving different ones onto a jobstrip. The participants also continuously scanned the wall-loader visualisation throughout the sessions, and on noticing a job getting 'critical' called up the roamer responsible for it, to check up on progress. This did not happen in the own externalisation condition (since they were not reminded), making them much more prone to committing errors. These included sending their team members to the wrong location or trying to assign them jobs when they were already busy. They also often forgot to update their own representations of ongoing tasks especially during busy periods.

The period when there was most marked differences between the two conditions, in terms of how well the

participants coped with the demands of the task, was towards the end of the session. During this busy time, the participants were having to monitor several simultaneous jobs, each with their own unpredictable demands. At one point they were confronted with a distressed garbled message from one of the stooges, requesting immediate back-up. This required the participant working out whom they could send along to help out from their team of roamers. As can be seen from the excerpt in figure 5 the participants in the Offloader condition were able to work out straight away from the DIV who to ask to help out. In contrast, the participants in the own externalisation condition spent much more time communicating with the other team members on the walkie-talkie to determine who to assign. Simply, they did not have the information available and so needed to re-create a mental model of the current situation by verbally requesting information – all of which could simply be ‘read off’ the Offloader DIV.

4: Conclusions

One of the main arguments we have advanced here is involving users in the creation of CIVEs can enhance collaboration and cognition. A main finding from our empirical study supported this: allowing users to be co-constructors in creating IVs (once an appropriate graphical format had been found) while designing the DIV to animate aspects of the IV to reflect changes was an effective way of supporting them. Above all, our DIV Offloader, enables people managing complex organisational tasks to be much more ‘in the loop’ during problem-solving. The key to this, is the power of the dynamic visualisation to effectively free-up users’ cognitive resources, lifting them from actors to directors.

Another finding from our research is that it is not straightforward as to which kinds of graphical representations to use in CIVEs. Little attention has been paid to this concern, although there have been one or two suggestions that simple abstract forms are effective at representing core features of dynamic processes (e.g. [6]). A key to finding the right kind and level of representation, is to consider carefully what the users’ tasks are, especially what decisions they have to make about the information they are monitoring and using. Our experience with the Chernoff faces and ‘ripening fruit’, metaphor suggest that these kinds of composite and continuously varying visualisations are probably best suited to judgments requiring selective attention, that is decisions that have the observer attend to one stimulus dimension while ignoring the others (e.g. [15]). In our case, however, what we required were visualisations that support *integrations* of judgments about dimensions. To

do this successfully we had in the end to design a format that exploited the capacity of humans to use spatial organisation as a means for relating variables one to the other. We also exploited the ability of humans to be able to understand the semantic meaning of having overlaying representations.

Team member Jon (J) discovers a group of rowdy students when roaming and radios in to the participant (P) requesting immediate back-up. Meanwhile all other team members are busy, some more than others. The participant has to decide who they can ask.

(i) Participant in offloader condition:

J: “Control, this is Jon, I need some backup in the 4a common room, over”

P: looks at Offloader to determine who to send: “OK Jo and Bob can you help John on the 4a corridor, over”

(ii) Participant in ‘own externalisation’ condition:

J: “...Come in control, this is Jon, I’ve got a group of students making a lot of noise in the 4a common room... I need some back-up, over.”

P: “...” [hesitates for a few seconds, looks through notes] “Sam, can you help Jon, there’s some protesters... Jon, can you confirm where that was?”

J: “Yea it’s 4a, the 4a common room”

P: “So Sam, go to the 4a common room”

S: “Hi Control, this is Sam, I’m busy dealing with this locked door, can you confirm, you want me to stop this and go down to 4a?”

[participant hesitates for a few seconds and curses to herself]

P: “...Ummm, Nicky what are you doing? Jo? Bob? Who isn’t busy? I need you to join Jon on the 4a common room to deal with the protesting students...”

Figure 5. Excerpts from the broadcast-mediated conversations that took place between team members in the two conditions during a very busy period.

In sum, our research into designing DIVs visualisations suggests that providing users with the ability to use various interactive building blocks helps them externalise their planning and in so doing reduces cognitive effort. We also found that our CIVE was successful at augmenting transient verbal communication, by allowing more permanent traces of what was being said to be displayed. And finally, our study showed that DIVs can significantly help collaborative decision-making, alongside information access and communication.

Acknowledgements

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

References

1. Card, S., Mackinlay, J. & Shneiderman, B. (1999) *Readings in Information Visualisation: Using Vision to Think*. San Francisco, CA: Morgan Kaufmann.
2. Chernoff, H. (1973) The use of Faces to Represent Points in k-Dimensional Space Graphically, *Journal of the American Statistical Association*, 68, 361-368.
3. Christ, R.E. (1975). Review and analysis of color coding research in visual displays. *Human Factors*, 17, 542-570
4. Cox, R. & Brna, P. (1993) Analytical reasoning with external representations. In: R. Cox, M. Petre, P. Brna & J. Lee, Eds. *Proceedings of the AI-ED93 Workshop on Graphical Representations, Reasoning and Communication*. August, Edinburgh. pp. 33-36.
5. Dix, A. and Ellis, G. (1998). Starting Simple - adding value to static visualisation through simple interaction. In *Proceedings of Advanced Visual Interfaces - AVI98*, Eds. T. Catarci, M. F. Costabile, G. Santucci and L. Tarantino. L'Aquila, Italy, ACM Press. pp. 124-134.
6. Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M., Richards, J.T. and Bradner, E. (1999) Socially translucent systems: social proxies, persistent conversation and the design of "Babble". In *Proc of CHI'99*, Pittsburgh, PA. ACM.72-79.
7. Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Lawrence Erlbaum
8. Healey, C.G., Booth, K.S. & Enns, J.T. (1995) Visualizing real-time multivariate data using preattentive processing. *ACM Transactions on Modelling and Computer Simulation*, 5, 3, 190-221. ACM Press, NY.
9. Heath, C. and Luff, G. (1992) Collaboration and control: Crisis management and multimedia technology in London Underground line control rooms. *Computer Supported Collaborative Work*, 1(1-2):69-94
10. Hutchins, E. (1995) *Cognition in the Wild*, MIT.
11. Johnson, B and Shneiderman, B. (1991) Tree-Maps: A space-filling approach to the visualisation of information structures. *Proc of IEEE Information Visualisation '91*, 275-282, IEEE
12. Larkin, J. and Simon, H.A. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11 91987) 65-99
13. Lee, M.D., & Vickers, D. (1998). Psychological approaches to data visualisation. *Defence Science and Technology Organisation Research Report* <http://203.36.224.190/cgi-bin/dsto/extract.pl?DSTO-RR-0135>
14. Lee, M.D., Butavicius, M.A., & Reilly, R.E. (2001). Visualisations of binary data: A comparative evaluation <http://www.psychology.adelaide.edu.au/members/staff/michaellee/homepage/viseval.pdf>
15. Maddox, W. T. & Bogdanov, S. (2000) On the Relation Between Decision Rules and Perceptual Representation in Multidimensional Perceptual Categorization, *Perception & Psychophysics*, 62, 984-997
16. Mohr, B (2001) 'Faces 2.0' Available at <http://www.bradandkathy.com/software/faces.html>
17. Morris, C.J., Ebert, D.S. & Rheingans, P. (2000) An experimental analysis of the effectiveness of features in Chernoff faces. In: W.R. Oliver (ed). *Proc 28th AIPR Workshop: 3D Visualisation for Data exploration and Decision Making*, *Proc. SPIE*, Vol 3905, 12-17.
18. Norman, D. (1993) *Things that make us smart*. Addison-Wesley
19. Rogers, Y, Scaife, M., and Brignull, H. (2002) Supporting "nomadic awareness" and team working through visually augmenting broadcast communication. Submitted to *CSCW'02*, New Orleans, November 2002.
20. Rogers, Y. (1994) Exploring obstacles: Integrating CSCW in evolving organizations. In *Proc. of CSCW'94*, 67-78. ACM.
21. Scaife, M. and Rogers, Y. (1996) External Cognition: how do graphical representations work? *International Journal of Human Computer Studies*, 45, 185-213.
22. 'Scrabble' [Board Game] (1997) J.W. Spear & Sons PLC
23. Tufte, E.R. (1997) *Visual Explanations*. Graphics Press, Cheshire, Conn.
24. Tweedle, L. Characterizing Interactive externalisations *Proc of CHI'97*, 375-382.
25. Ware, C. (2000) *Information Visualisation: Perception for Design*. San Mateo, CA. Morgan Kauffman.
26. Wickens, C.D. and Carswell, C.M. (1995) The proximity compatibility principle: it's psychological foundation and relevance to display design. *Human Factors*, 37(3),473-479.
27. Wong, W., O'Hare, D. and Sallis, P.J. The effect of layout on dispatch planning and decision making. In *People and Computers XIII HCI 98 Conference*, Sheffield, UK; Springer. Sheffield, UK.