Putting it all together: information visualizations, display arrangements, and sales transactions

John Halloran
Interact Lab, School of Cognitive and Computing Sciences, University of Sussex, Falmer, Brighton BN1 9QH, UK
e-mail: johnhall@cogs.susx.ac.uk

ABSTRACT
Computer networks offer powerful ways of integrating and coordinating information so that it can be more effectively accessed and used. Often, the focus is on the back end: the systems infrastructure, applications and standards which will enable this. However, with diversification in the range of activities networked information now supports, growth in numbers and types of user, and the appearance of new technologies, new challenges arise. As well as the issue of back end information design, creating networked information systems increasingly involves the issue of how to coordinate this, across a range of activities both single-user and collaborative, with a host of front-end design problems concerning information representation and visualization, display arrangements, interactivity, and human factors. We tackle these issues through a case study of the analysis and redesign of user interfaces used by the travel industry during sales transactions. This involved the development and evaluation of an interactive trip planner, featuring linked information visualizations and multiple displays, which represents networked information in innovative ways, and allows new kinds of access, configuration and discussion.

INTRODUCTION
A major motivation for information networking is that it can radically improve information organisation and retrieval. An area which has been recently described as 'exploding' (1), it involves several concerns. A central issue is the design, implementation, and management of information resources, for example library catalogues (2), or the databases that back-end online stores, local government services, and so on. Related to this is the linkage of information within and across devices – and this involves developing appropriate networked application protocols and standards (3). In addition, because the success of networked information systems is linked to how effectively information is delivered to users, user interface design is also a crucial issue. The question of how we should design interfaces for networked information systems is, however, becoming increasingly complex. Different kinds of people - shoppers, students, business people and educators - all make use of networked information, and interfaces need to be developed appropriate to a wide range of needs (4). Access to networked information services can be single-user or collaborative (for example, kitchen planning (5) or virtual estate agency (6)), and collaboration can be co-present or distant (collaborative browsing (7)). Information can vary from relatively simple, for example regional train timetables, to very complex, as in financial projection (8). These issues take on added urgency as the range of technologies available for retrieving and displaying information expands and diversifies: information networks can now be accessed not only by PCs and terminals, but also a new generation of mobile devices, including WAP-enabled phones and PDAs. There have also been recent breakthroughs in fixed display technologies including largescreens and flatscreens which now offer increased low-cost high resolution screen estate, creating opportunities for massive and/or multi-display interfaces.

A key driver in the development of networked information systems is sales transactions, an area increasingly supported by networked information and where there are interesting challenges for the creation of user interfaces. While e-commerce tends to be restricted to 'objective' products (for example books, clothes, and so on), other kinds of product (for example hi-fi systems, kitchens, and round-the-world tours) consist of many possible components and options, and need to be configured. The data is often so complex that it is hard to deal with in the absence of a knowledgeable expert who can offer guidance and advice. Creating configurable products, then, appears to require a synthesis of digital and personal services. Focussing on a case study in the creation of user interfaces for this purpose, this paper explores some central issues and questions concerning the design of information for users in new technological contexts: How should information be designed for activities whose form is changing and evolving? How should user interfaces be designed for different types of device? What kinds of display should we use for different kinds of information, including complex data? What sorts of interactivity are appropriate? And what kind of human factors do we need to consider, especially where use of networked information might be collaborative, as well as single-user?

USER INTERFACES: DESIGN ISSUES

One reason it is hard to create configurable products online is that it is difficult to find representations that effectively integrate and present the complex data involved. Configuring a product like a kitchen, for example, involves matching constraints like budget, space and deadline against a host of parameters including cost, size, appearance, facilities and availability in a context where the product is constantly changing as components are added or removed. Representing constraints and parameters dynamically involves delivering a lot of information in many different forms, including text, visualizations, and diagrams. This presents not only practical problems (e.g. screen overcrowding); some problematic theoretical issues are also involved.

Work in external cognition (9) reveals the importance of ‘external representations’ (ERs) when interacting with information. There are many different types of ER including text, visualizations, and diagrams. These may have different properties, for example being static or dynamic (e.g. animated or otherwise changing), and/or fixed or interactive (e.g. user-manipulable). Whatever the type, an effective ER brings about ‘cognitive offloading’: it reduces the amount of mental work people need to do to make use of information. Examples of effective ERs include bar charts, which, in making it visually clear how one item relates to another (for example, class size), makes comparison easier than where the same information is presented numerically; or (to take a specific example) the London Underground map, which by rendering the system rectilinearly constrains visual search making it easy to find one’s way from A to B.

Information visualization

These examples are both of one class of ER: information visualizations. Information visualizations are potentially highly effective ways of presenting information to users. However, designing effective information visualizations is complex, as there are diverse forms, functions and uses. One major issue is how much information can realistically be presented by a visualization. This bears on how much information can be effectively delivered to users through interfaces.

Much of the work on external cognition involves comparing information visualizations to see how far they can produce cognitive offloading (e.g. 10). This research often seems to suggest that for cognitive offloading to be achieved, *single* visualizations are required: although a scatterplot, for example, may involve more than one variable, it is a single representation. Some tasks, however, may require more than one visualization, for example topographical maps with separate keys. Sweller (11) argues that where multiple representations have to be coordinated, this can cause high cognitive load. In solving geometry problems, for example, students have to cognitively integrate graphical and textual material in different spatial locations. This involves remembering information (e.g. the apex of a triangle) while other information (e.g. text referring to the apex) in the same modality - here, visual - is searched for. The implication is that representations should be visually integrated to save the work of cognitive integration: graphical and textual items that are related should appear in the same place.

This example involves cognitive integration of visually interpreted material in two different media, graphical and textual. Only one of these could be called an information visualization, but it raises the question of whether separate information visualizations like a map and its key should be integrated in the same representation for the same reason: to remove the need for cognitive integration through visual search. This issue is addressed by MacKinlay (12), who discusses alternative information visualizations for selecting a car on the basis of best mileage and price. Here, cognitive integration requires visual search if the two variables are represented as different graphics (for example, two separate barcharts). Mackinlay uses a system known as APT to generate various solutions. Some are easy to understand, for example a scatterplot where one axis is cost and the other mileage - we need only look for the bottom right item. This, however, is problematic because labels have to be left off for reasons of occlusion, so we do not know which car we are selecting. Another solution is a single axis (mileage) where each car is represented as a dot, the varying sizes of which indicate relative costs. This allows labelling but is much harder to interpret.

These examples demonstrate issues in designing and comparing complex visualizations – 'complex' because the information presented is multidimensional, involving more than one variable. There has been much debate on their efficacy. Cox and Brna (13) argue that where a complex information visualization is novel, possible cognitive benefits may be outweighed by the learning overhead. This is an argument for conventionality (i.e. sticking to what people are generally used to or familiar with) in design of visualizations. However, if visual integration of multidimensional information should not be done because it compromises conventionality, this implies multiple separate visualizations, which leaves us with search issues. Others argue that if they are well designed, complex information visualizations can be unconventional but still lead to understanding – and even increased insight - without major difficulties (14).

Display technologies and arrangements
The literature on information visualization is inconsistent on the issue of whether complex information, of the type needed to put together configurable products, should be visually integrated. Many examples appear to assume static presentation, without going into the possible roles and properties of different display technologies. However, this question is becoming increasingly important. There are now many different types of display, size and configuration.

Novel displays may create new insights and opportunities for information visualization design. Grudin (15) reports on an increasingly common practice - that of single users using multiple monitors (two or three instead of just one). With the increase in high-resolution screen estate, users partition different tasks across different monitor screens, so that, for example, a wordprocessor may be up in one screen while e-mail is up in a second. Grudin reports that this has a number of effects. Users perceive bezels as separators and different screens are regarded as different workspaces. The arrangement of these spaces around the user makes use of focal and peripheral awareness, so that while one screen is the focus, the importance of the others is demoted while they still remain in the field of vision ready to be attended to if needed. An important finding is that with a multi-screen desktop, coordination across different tasks in different locations is achieved by switching gaze, rather than by manipulating windows (through moving or bringing to front, for example). Grudin claims that this leads to more seamless task switching and integration.

Grudin's work suggests that the ability to place visual material in space and switch between items by means of gaze may have cognitive benefits. Cognitive benefits of various kinds may also accrue from using largescreen technologies. Single Display Groupware (SDG) uses largescreens to support co-located groups (16) where size and resolution enable co-present collaborative use because of increased salience of shared material. InteracTable (17) is a shared digital table which places the screen horizontally enabling shoulder-to-shoulder collaboration with similar effects. DiamondTouch (18) capitalises on this arrangement by adding interactivity which allows users to share control: being able to see what another is doing with material can help one understand it. Largescreen resources also allow not just proximal collaborative use, but also distal viewing by multiple users, enabling involvement in collaborative work without direct control. Recent research looks at some of the issues in largescreen arrangements, showing how different arrangements and interaction styles give rise to different cognitive effects when collaborators access information to solve tasks (19, 20). There are now many new analytical axes for considering displays, including small/large, vertical/horizontal, proximal/distal, and multiple/single. One question this research implies but does not explicitly address is how different types and arrangements of display might interact with different information visualization designs to produce different cognitive effects.

Interactivity

A further issue raised by the research into new display technologies is that of the relationship between interactivity and cognitive integration. In multi-display use, for example, the need to promote or demote windows (i.e. bring to front or send to back) can be replaced by glance - a change in physical interactivity - leading to improved ability to switch between and integrate different tasks: a change in cognitive integration. However, if use is collaborative, users' moving windows, rather than glancing, may make their actions more salient to others. A further important issue is dynamic interaction between different representations – or *dynalinking* (21) – where a change in the state of one representation is accompanied by a simultaneous change in another, for example where an animation in one location triggers different textual descriptions in another as it unfolds.

There is another aspect to interactivity, which is how far users are able not just to look at, move or place an information visualization, but how far a user can actually *create* it: whether the visualization is a non-manipulable entity or is constructed by the user. Rogers and Scaife (21, *ibid.*) show how user construction of a representation through direct manipulation helps promote understanding of food chains in children. In addition, Cox (22) shows that self-constructed visualizations can have valuable cognitive benefits over pre-constructed ones.

Recent research into information visualizations, display arrangements and interactivity suggests there is much crossover between these areas. KidStory (23) gives us a flavour of some of the issues here, showing how information visualization, display configurations, and interactivity are becoming interlinked research areas which are responding to issues of how networked information, especially where data is complex, should be represented across different types of interface, how information visualizations should be designed, and how they might be displayed. A critical question, on top of whether different cognitive effects might accrue from interactions between different kinds of display arrangement and information visualizations, is what the effect of different kinds of interactivity - including dynalinking and user-manipulation - might be. In the following sections we discuss how we approached the issues empirically, through the design and evaluation of a user interface for sales transactions.

INFORMATION ARRANGEMENTS IN SALES TRANSACTIONS

Working with a high-end London travel agency specialising in round-the-world tours, we carried out a six-month ethnography of how configurable products are put together and sold (for a fuller report on the ethnography, see (24)). We were interested in the uses of networked information systems, especially how this information is accessed, configured and used during consultations between customers and agents.

Sales transactions as information processes

The travel agency's networked information system consists of 'Traveller', a fares database dedicated to flight packages; and 'Gecko': an in-house booking and ticketing system which includes pricing information on other products like hotels, tours, and hire cars, as well as text versions of the agency's brochure materials. The information is delivered to separate windows on each agent's PC screen.

Traveller and Gecko support a multi-stage information process which typically takes some weeks to conclude. The first stage is typically a face-to-face consultation where a tour plan is worked out between the agent and the customer. During this stage (which we call the approach stage), the networked information system needs to cater for the needs of two users – the agent and the customer – in such a way that satisfactory input for the rest of the process (i.e. the tour plan) can be obtained. During the next stage –the development stage - the agent works alone, translating the plan into a detailed quote form that is then sent out to the customer. The following is the commitment stage: if the customer agrees to the quote, s/he sends in an initial deposit and the agent starts the booking process. If the customer wants to make changes, s/he may send in suggested revisions, or visit the agency again, and the transaction can backtrack to new approach and development stages where the agent will need to produce a new quote form until the customer is happy. Finally, at the closure stage, the balance is paid and the holiday is sold.

Problems at the start of the transaction

Networked information is used at all the stages of the transaction, but we became especially interested in the approach stage. Getting this stage right is critical. It is here that a relationship between customer and agent is established, a plan sketched out, and a product opportunity realized: in other words, the foundation for the rest of the transaction is laid. However, despite agents' skills and experience, the arrangement of networked information sources, and the design of interfaces onto these sources, means the approach stage can often go wrong.

Physical and representational arrangement of the interface

When a customer sits down with an agent, the arrangement of information and technology is typically as appears in Figure 1 (below). The two parties sit on either side of a desk, facing each other. The agent has a PC facing him/herself, which is turned away from the customer. The networked information sources used during this stage – Traveller and Gecko – are displayed on the PC screen, form-based applications which are difficult for the untrained to understand. An effect of this arrangement of information representations, display and physical positioning is that the networked information is very difficult to share. While the agent has an interface onto the networked information used to construct the product, the customer does not: for the customer, the *agent* is the interface.

The planning model

In order to build the quote, the agent implements a product-ordered planning model: the flight package is sorted out first, followed by hotels, and finally 'land sales' like tours, safaris, and hire cars. The planning model takes this form because deposits are taken and products booked in the same order, and because it is easier for agents to create constraints at general levels before detail is added to the tour. The planning model dictates an ordered query structure. The first information the agent will attempt to get from the customer is where s/he wants to go, so that the first item, the flight package, can be established. Then s/he will go on to the other products.

Verbal representation of the product

During the approach stage, the customer often has only an outline of what s/he want to do and needs the agent's advice. However, s/he will immediately be asked specific questions as per the planning model, typically focussing, to begin with, on where s/he wants to go on what dates. The discussion can quickly become complex and difficult for the customer to remember. For example, the customer may say s/he wants to visit the US, South Africa and Australia in that order, and the agent, entering the information into Traveller, retrieves results enabling them to verbally quote various alternative flight packages with different stop-off options. Trying to weigh up the alternatives in the head can make it difficult for the customer to confirm on the spot, or to be able to deal with further questions. The approach stage frequently comes to a halt at this point, with the customer requesting, and/or the agent offering, a written summary, showing that an ER is needed to enable the customer to overview the material and make decisions.

An important issue at the approach stage, then, is that because the interface is designed for the agent, the representational modality, for the customer, is verbal. This modality, however, is insufficient. One reason is that conversation is evanescent – to keep track of a topic participants need to remember what has been said, and in complex trip-planning this can quickly become impossible. On top of this, conversation analysis (25, 26) is consistent with research into cognitive load in showing how, given the type of problem represented by trip-planning, conversation can follow a recursive structure with pairs of utterances (for example a question and an answer) separated by further pairs. In other words, an initial question cannot be immediately answered but has to be postponed until intervening questions and answers are provided. Such 'insertion sequences' can only be negotiated up to a point before people lose track (27, 28). This can be exacerbated by a 'granularity' issue - where the customer tries to interrupt in order to discuss a detail (for example, a side trip from major city) where the dictates of the agent's query structure oblige him/her to try to steer the conversation back to more general issues (whether this city can be included in the flight package at all): the two parties are on different tracks so that questions may not be perceived as valid and answers are *never* forthcoming, although they may be waited for. The problem of customers wanting to discuss details that diverge from the query structure also reflects that travel agencies, unlike banks, do not have effective ways to 'configure' customers – that is, implicitly coach them, through forms of question and answer, as to what kind of responses they need in order to progress the enquiry (29). During the approach stage customers can often simply be unaware of what the query structure is or why it is important to fit in with it.

Social interactional issues
Problems with information coordination during the approach stage can arise from the design and positioning of the interface, pressure on the representational modality (speech), and the fact that customers are not 'configured', leading to problems in negotiating the query structure. However, on top of this, there are social interactional issues to do with sharing and evoking information interpersonally.

The two parties face each other as strangers, and have no particular method of ice-breaking. As Goffman (30) points out, this kind of mutual orientation is intimate. Goffman draws a distinction between verbal information and other face-to-face phenomena like gaze, orientation, gesture, and proximity. Direct eye contact - of the type that occurs during an approach stage - has a meaning of 'I am paying attention to you'. Together with the need to negotiate emotive questions which may be difficult to ask and/or answer, including, for example, 'How serious are you?', 'How much do you have to spend?', or 'Have I exceeded my budget yet?', the situation can create social awkwardness. However, during the approach stage, direct gaze alternates with aversion of gaze, because the agent needs to attend to his/her PC to input and retrieve information. According to Goffman, aversion of gaze is interpreted as meaning 'I am now doing something else'. This can be confusing for the customer.

Heath (31) shows that aversion of gaze is 'licensed' in some activities – in other words its social meaning is understood and accepted. In medical consultations, the aversion of gaze by a doctor means not only that the doctor is doing something else, but is also a cue for the patient not to interrupt – patients therefore go silent (32). Because it is licensed, disattention is understood not as something anomalous (or possibly even rude), but as a standard part of what doctors do. However, disattention during the consultations we observed does not appear to be licensed in the same way. Instead of going silent, customers may appear uncomfortable, interrupt, or try to regain eye contact, as if to say, 'Concentrate on me, please; not that machine'. An example of this can be seen in Figure

1, where the customer does not understand the information the agent is dealing with (1a, 1b) and uses direct gaze to try to get this clarified (1c). Thus we see how the physical arrangement of people, as well as technology, and the problem of esoteric information representations, can create counter-productive social-interactional effects which can create social discomfort and confusion. These effects demonstrate how the lack of an effective customer interface makes it hard for the agency to gain the input (the trip plan) which is needed for the information process the agency's networked information system is designed to support (i.e. the transaction). This reflects the critical importance of effective user interfaces for networked information systems.

FIG. 1a [left]; b [middle]; c [right]: *Social interaction and communication problems*

Another means of communicating face-to-face is gesture. Gesture can signal emphasis, hesitation, anxiety, satisfaction, and so on (33). It can also be used to point things out, or to 'draw' something 'in the air'. Again, in the interactions we saw, this 'channel' tends to become ambiguous. Agents may point and 'draw' on the basis of information they have, but this may have no clear meaning for the customer, since s/he does not have equivalent access to this information. Equally, customers may produce meaningful gestures which are not being attended to because the agent is looking at the PC screen (one such is the customer's placement of the hand to the chin in Figure 3b, conveying concentration but also frustration).

Social interactional issues like these can create difficulties in communicating at a time when there is potential cognitive overload. In other words, the social situation can put a brake on what/how much information the customer volunteers, its quality, and its usefulness as input for the transaction. Also, such issues are not likely to be conducive to an enjoyable customer experience, one where s/he feels relaxed and involved; and can compromise the quality of the personal and working relationships the two parties need to establish during the approach stage.

Customer pre-planning as a potential solution

One way to avoid the problems that arise at the approach stage is to leave this stage out. This can potentially be achieved through customer pre-planning. Here the customer works up an itinerary alone and can send it to the agency, where an agent then works directly on development without a face-to-face consultation ever having taken place. The agency we worked with sends out a range of detailed, comprehensive brochures on request and our research shows that having received a set of these, customers are prepared spend a good deal of time and effort working out an itinerary.

Although the information is attractively presented and persuasive, it is diverse, may exist in different brochures, and can be arranged esoterically within individual brochures. This creates issues of retrieval and coordination, as well as the search and memory issues familiar from the literature. These are exacerbated when we consider the complex nature of the task: specifying this type of product involves setting constraints (including remaining budget and available dates) and coordinating numerous parameters (including cost, time, sequence, and destinations) simultaneously – constraints and parameters which change dynamically as the product is developed. In other words, computational work is required. Because this work is hard (if not impossible) to do in the head, people often use pen and paper, post-its, and calculators. They may go through several iterations to work up their itinerary, which can start as a set of scribbled notes and end up as a neat, word-processed 'final product' ready to be sent in or taken to the agency.

Brochures: a 'proxy interface'

Unsurprisingly, because of all the work customers put in, they tend to develop fairly firm ideas of what they want during this process. However, while customer pre-planning using brochures would appear to be ideal for agents' purposes, it can create various problems. Customer itineraries are typically arranged chronologically - not according to product order - so that the agent has to perform a time-consuming translation. In addition, an itinerary may not be actionable, so that it has to be redesigned or even scrapped. The reason for this is that a set of brochures when used by customers working alone acts as proxy interface onto the same information that the agency uses – but one which does not carry all the information. For example, customers do not have access to the often bewildering range of flight package options, especially late-breaking deals. This means they may come up with flight packages which turn out to be unrealistically expensive or which cannot be put together at all, because of timing issues not revealed by brochures. Another problem is where customers take an itinerary to the agency in person, and look for confirmation of the entire plan on the spot. This is not realistic until development (typically involving several hours' work) has been done. It makes for consultations which can be unsatisfying for the customer after the amount of work put in.

While customer pre-planning using brochures may be an effective way of getting past the approach stage of the transaction, then, problems can arise relating to divergent planning models and lack of information, as well as customer over-commitment to their plan. This way of doing things is also labour-intensive for both parties and misses out a crucial service many customers still want and the agency feels it is invaluable to deliver: face-to-face consultations where customers can get the benefit of expert agent attention and advice during the planning process.

Knock-on effects

Frequently, the problems that occur during the approach stage mean the sale is lost, the customer never returning. However in many cases the approach stage is negotiated satisfactorily enough for the development stage to be entered. What happens during the approach stage has a big influence on the rest of the transaction. If an actionable, customer-agreed plan is forthcoming, the rest of the transaction is often likely to proceed smoothly. But if a plan cannot be actioned (because, for example, the limited information available from brochures can lead to errors), there may need to be repair work including, in some cases, starting from scratch. If substitutions are made by the agent, customers frequently want to query points or make changes before sending in deposits, rewriting the itinerary and sending it in in text form. This creates an update issue - agents have to change the electronic quote form accordingly, involving translation; as well as a recap issue: where changes are proposed agents and customers frequently have to remind each other of the entire itinerary as neither can remember it.

DESIGNING A SOLUTION

The problems that arise at the approach stage are, we would argue, the result of inadequate user interfaces onto the networked information systems used to support sales transactions. The lack of an effective customer interface, especially, creates significant problems in eliciting the main input for this process – the itinerary. Our research shows that putting an itinerary together (and by implication product specifications in other domains) involves considering a whole range of interacting issues: the availability of information, its design and representation, its display, its physical context of use, and its social context of use. In designing a solution we wanted to see how we might improve the approach stage by creating an interface that would address these issues and could deliver information to both users in appropriate ways, enabling tour plans to be more effectively created, increasing sales potential, and reducing problematic knock-on effects.

Information visualizations

One reason it is hard to keep track of the product as it develops during the initial consultation is that there is no shared ER. Rather, the product is heavily represented by the verbal medium. As we saw, this creates cognitive load problems attendant on tracking the product by means of one modality – talk – not least because of the need to deal with multiple constraints and parameters at the same time. What we wished to do was to be able to explicitly represent constraints and parameters as external representations in such a way that their relations could be clearly perceived. How might we approach representing this complex data in such a way as to reduce cognitive load? Information visualizations seemed appropriate because (a) they can provide a persistent trace of what is being discussed to remove cognitive load from working memory; (b) there is potential for integrating many variables in one visualization; (c) numerical data (including timespans and costs) can work better visually than textually; (d) canonical visual representations like maps are highly effective ways of representing geographical data, central to this type of product; and (e) visualizations can remove representational pressure from the verbal modality, which may still need to be used.

Multiple or complex visualizations?

Our initial ideas concerned creating a map-based tour visualization on which the different components could be shown, together with timeline and budget visualizations. Thus all of the relevant parameters might be simultaneously displayed. Figure 2 shows some early paper prototypes.

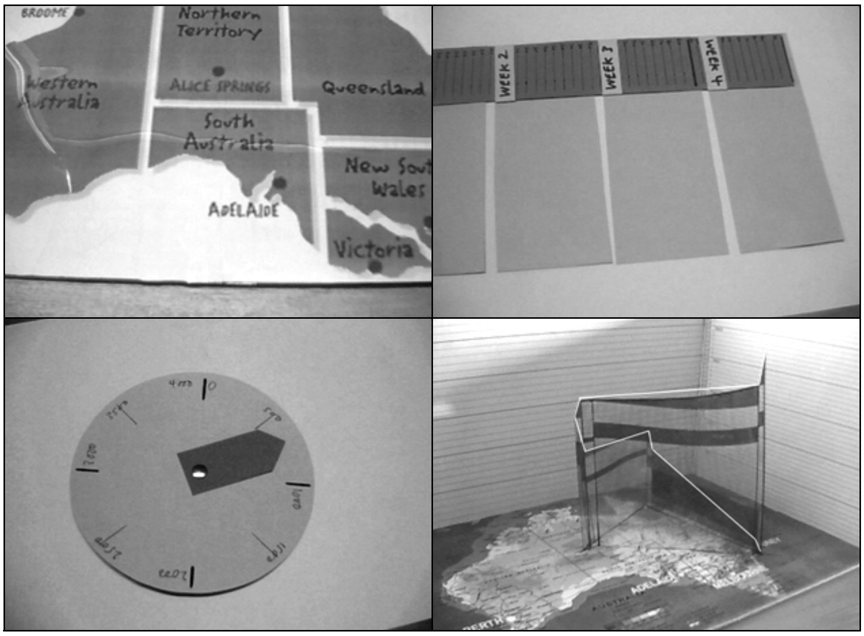

FIG. 2: *Paper prototypes.* (a [top left]) *map;* (b [top right]) *timeline;* (c [bottom left]) *budget;* (d [bottom right]) *complex visualization*

The map-based visualization (Figure 2a, which, for prototyping purposes, we limited to Australia) is covered with transparent film on which components of the trip could be drawn. The timeline (Figure 2b) represents weeks divided into days with space for text descriptions of the different components of the trip. The budget (Figure 2c) is a pie chart where slices represent the price of each component as a proportion of total available budget. The concept was that these would be implemented electronically with automatic update of timeline and budget according to what events took place on the map, the only visualization intended to be user-manipulable.

The three different visualizations (2a, 2b and 2c), if implemented separately as shown, would require cognitive integration. A change in the state of the map (for example, a hotel is added to a destination) would be associated with an automated change to the timeline (we add text to show what the component is and when it occurs) and the budget (we add a slice). The relations between the visualizations would have to be realized internally through cognitive integration. The relationship of cost to time, or cost to location, for example, could only be worked out by examining different visualizations at different points in space, which, as we have seen, can create memory and search issues. Because of the issues that can arise from using multiple visualizations, we produced an alternative prototype visualization which integrates all the different variables. This is shown as Figure 2d.

In this visualization, which is complex rather than multiple, the map is augmented by a three-dimensional ‘sail’. The rising height of the sail represents time elapsed (highlighted in white); the wedges are time sections, and the rectilinear sections are amount spent. As we have seen, complex visualizations can to be hard to interpret without some training if they are unconventional. This was borne out empirically in our user testing scenarios. Apart from the visualization being hard to ‘read’ there were occlusion problems including difficulty in reading destinations, and big issues with accurately registering the different components with the numerical scales - allowing read-off of cost and time - that appear on the back and right walls (which involved actions including crouching to defeat perspective).

These issues raised an important problem: which, and how many, parameters can or should be integrated in a single information visualization? According to Peuquet and Kraak (14, *op. cit.*), maps are a representation which can be easily understood by the general population, down to children as young as four years. The fact that a map is an analogue of its referent may be an explanation, but on top of this, people are readily able to accommodate related concepts, for example relative location, direction, distance, and sequence (e.g. of destinations along a route). This offers clues as to what parameters might be successfully integrated in one complex information visualization.

Interactivity
One way our map prototype differs from many map visualizations is that from the start it was designed to be user-manipulable: we intended the trip visualization to be built dynamically by its user, which potentially adds another layer of cognitive complexity – interactive interface objects. Following research into intuitive widgets (34, 35), and consistent with the notion that conventional representations are likely to work best, we wanted to use simple drag-and-drop functionality as well as flyout menus, both highly familiar to windows interface users. This would enable user construction of visualizations, which has been shown to be an effective way of creating meaning. Our final map prototype appears as Figure 3.

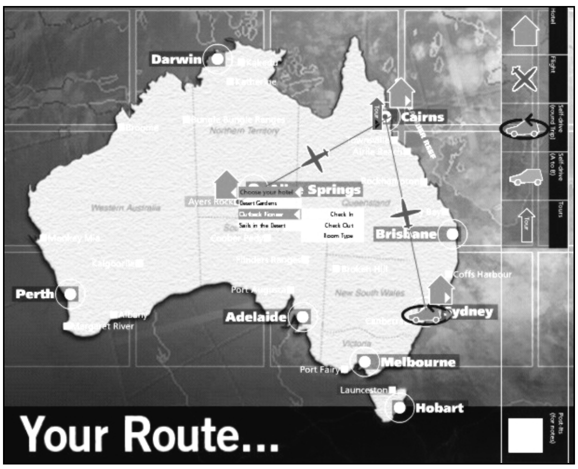

FIG. 3: *Electronic map prototype.* Objects can be dragged onto and between locations, and options can be selected through dynamic menus

Of interest is the fact that the resulting representation loses two parameters, time and cost, from our 'sail' visualization, with a large gain in comprehensibility. This shows that complex visualizations can work where they avoid the unconventional, for example the attempt to integrate geographical with temporal and budgetary information. This suggested leaving visualizations of time and budget (which seem conceptually distinct from each other as well as from geography) separate, and as conventional as possible. Therefore we represented time like a year-planner, and budget as a pie-chart, both conventional for these domains. We also wanted to display brochure information dynamically. The resulting electronic prototypes appear as Figure 4.

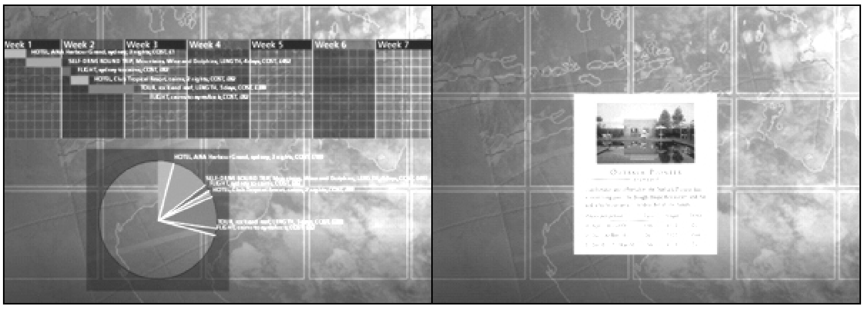

FIG. 4: *Electronic timeline, budget and brochure visualizations.*

These visualizations were not designed to be alterable by direct user manipulation but would update automatically when the map was used. So for example, when a hotel is selected at a location, relevant brochure information is displayed (right). When the hotel is configured (i.e. time and cost established through the menu), it is simultaneously inserted into the timeline and the budget (left). Thus the four visualizations (map, budget, timeline, brochure pages) are dynalinked, a form of interactivity that can enable cognitive integration of multiple visualizations by representing relationships.

Display arrangements

Our initial display design consisted of a large ($1m^2$) horizontal back-projected workspace with the three visualizations – map, budget/timeline, brochure display - in different windows. These windows could be promoted, demoted, minimized, maximized, resized and moved. This arrangement, although it promoted some collaborative success (36), did create some difficulties. The size of the display meant its top was hard to see by people working next to each other. In addition, screen estate was not realized satisfactorily as the display is a blowup of a single PC screen with loss of resolution. Perhaps most important, though, is that people's ability to move the visualizations about in order to free space to look at others continuously re-created a search problem: users had to remember where, for example, they put the budget. In other words, we created exactly the single-modality search issue which has been shown to create cognitive load.

One way around both the screen estate and resolution problems was to use multiple screens, which can be oriented in different ways so that those further away from users could be tilted toward them. At the same time, while Grudin's work on multiple monitors (15, *op. cit.*) shows benefits of partitioning different tasks across multiple displays, little work has been done on partitioning components of the *same* task across different screens. We were interested in whether and how this might be possible or useful as a solution to our own issues.

In designing a user interface we wanted to explore whether there might be ways of linking multiple information visualizations in ways which would not increase cognitive load, but *decrease* it – against many of the findings in the literature. One way of achieving this, we felt, would be to 'lock' representations in specific spatial positions, so that their locations relative to each other would be persistent. If the user knew where to look, this could cut out search. At the same time, both parties could easily refer to the same thing, rather than one collaborator having to work out what the other has done with the display. This also capitalises on Grudin's insight that coordination across different tasks in different locations is much more easily accomplished through switching gaze than by manipulating screens (e.g. promote/demote). Spatial locking could be achieved by placing the visualizations on screens whose position is fixed, by embedding them in a console.

The locking of representations to locations is potentially a way to cut down search and so promote cognitive integration of multiple information visualizations. As we have seen, dynalinking can also support this. We felt that a combination of spatial locking and dynalinking could reduce, if not obviate, problems of search attendant on using multiple representations.

Human factors

As we saw earlier, accessing and developing information in a sales transaction is not just to do with availability of resources but also the quality of the social interaction that goes on. One striking finding was how, during the approach stage, the physical orientation of the collaborators relative to each other (opposite/facing) creates social awkwardness – even inhibition. At the same time it means that the visual material being worked with (displayed on the PC screen) can only be seen by one party. We wanted to create a shared interface, and the corollary was that both parties should sit shoulder-to-shoulder in front of it. This, we felt, might help remove social awkwardness, as well as create persistent content to be discussed, potentially reducing (a) convoluted conversational structures (with many insertion pairs), and (b) the need for one modality – speech – to represent the product.

The final prototype

Our final prototype consists of dynalinked information visualizations displayed across three screens. Spatial locking is achieved by embedding the screens in a custom-built table. The result appears as Figure 5.

FIG. 5: *The final prototype*

The interface is linked to the agency's networked information sources, Traveller and Gecko. To create a plan objects are dragged and dropped from the 'console' to the right of the interactive map (front left screen, also shown as Figure 3 above), onto destinations. When this is done, the different product types (hotels, tours, hire cars) trigger flyout menus dynamically built from information delivered by the databases, allowing selection of different product possibilities at that destination. When particular products are selected, the relevant brochure information appears (back screen), helping the user to make a decision on whether they will accept this product, which can then be configured (e.g. for dates and number of people) and the results added to the timeline and budget visualizations (front right screen). Flights are added in a similar way, except that these are dragged between, as well as onto, destinations. Products can be removed by dragging them off the map, or inserted at particular points. All the visualizations are appropriately updated simultaneously.

TESTING THE SOLUTION

In partnership with the travel agency we worked with, we were able to try out our solution at a major UK travel trade show. Here there were many opportunities to see whether our design works as a shared user interface. Thousands of people passed through the show and many different kinds of people and groupings tried out the planner together with an agent.

This paper has shown how sales transactions involve a complex interaction between information representation, display, interactivity, and human factors. Our solution aimed to reduce the problems we identified by designing an interface which would simultaneously address all these issues, that is, effectively redesign the interaction that occurs during the approach stage to enable tour plans to be more effectively produced and used during the remaining stages of the transaction. While further research needs to be carried out (especially in-store), our findings from the trade show already suggest some success.

Information visualizations, multiple displays, and interactivity

Early on in the paper we looked at issues concerning the efficacy of multiple information visualizations – visualizations which are separately presented and at different locations. In our system, many kinds of information are represented simultaneously but separately. Might this not cause cognitive overload? In fact, we saw very little evidence of this. Attention was generally focussed on the main map, with the other visualizations – budget, timeline, brochure information – as useful peripherals. This is consistent with findings which suggest that the allocation of representations to focal and peripheral attention assists with organisation, and that spatial location of representations enables people to cognitively promote and demote information through glance. This means that, if not wanted, parameters are defeasible through removal of attention: ongoing cognitive integration of the different visualizations is not necessary. In fact, the budget/timeline and brochure visualizations seemed to be integrated at intervals through glance.

The glance behaviour suggests that our locking of the visualizations to locations was effective in our aim of facilitating cognitive integration. At the same time, the ability to place the visualizations not only in space but also on different planes (vertical, horizontal) is a way of organising material so it can be easily seen. People clearly got used to looking at different planes, and the fact that the planes were different may have acted as a cue in remembering the spatial location of the different visualizations. These findings go against those familiar from the literature, that cognitive load is incurred where multiple representations are used. We were able to remove this effect through use of multiple dynalinked screens where visualizations are spatially locked. This is an example of interaction between information visualization design and display arrangements which has gone largely unaddressed, but which appears to be important.

We have claimed that delivery of networked information to users depends critically on interfaces. One of the problems we saw was that at the start of the customer-agent transaction there was, effectively, no customer interface, making it difficult for the two parties to remember and discuss the product during the approach stage. The challenge facing us was how to create shared representations and promote shared understanding. The decision to create an interface that is identical for both parties appeared to be a successful way of doing this, and several positive effects resulted.

It was much easier for the two parties, especially the customer, to keep track of what they were talking about, since there is a persistent representation in front of both. By enabling users to see information, our system releases networked information in a representational format which is salient not only to the agent but also to the customer, and we observed this leading to much fuller, more detailed discussion at an early stage. Because customers can access information previously unavailable to them, they are in a much better position to make decisions: every time a product is configured, its time and cost implications become clear. While the interface is visually identical to both parties, it is designed to cater for different cognitive needs: customers can consider the information chronologically, and agents in terms of product order. The planner reduces the need for a rigid query structure since product order can be perceived regardless of what planning model is used to construct the itinerary. At the same time the reduced need for a rigid query structure means that customer configuration is no longer an issue. For related reasons, the system helps reduce the 'granularity' effect, where the two parties are working at different levels of detail. It is much easier to switch between levels of granularity. We saw a customer and an agent spending a good deal of time talking about sights within a city before they established the flight plan for the continent. This can be done because the planner allows constraints to be generated at any level of detail: constraints do not need to follow product order. Rather, the interface allows planning of products in any order.

The interface does not just display information: it allows configuration of information through interactivity. Our drag-and-drop widgets proved to be readily learned by users after some 'playing', to the extent this functionality was hardly commented on. Nested flyout menus did present some difficulties but both customers and agents were able to 'drive' the planner, sharing initiative. Customers' being able to control the planner themselves appears to have created increased involvement and understanding, consistent with the literature on self-constructed representations. In addition different users, especially agents, repeatedly commented on how much they liked the fact that menus triggered relevant brochure material without any need for search. These observations show not only that sticking to conventional widget functionality appears effective, but also how different forms of interactivity like manipulation and dynalinking can empower users to interact with information.

Human factors

Other effects were observed, which show how an interface can promote better collaboration, communication and social interaction. Linking to Traveller and Gecko and re-representing the information visually implies that agents do not have to translate queries to these systems into verbal form so that the customer can understand it – customers can perceive the meaning of information immediately and agents are freed up to do what they are good at: describe and recommend products. The creation of a shared visual representation means that the product can be represented in two modalities – visual and speech. This has effects on conversational content and structure, and also on social ease and ice-breaking – people relax far faster and may even joke together (Figure 6c). There is much less nesting of insertion pairs, and the socially odd mix of attention and disattention disappears. However despite the reduction of insertion pairs - a way of dealing with complexity in planning - the planning process can, paradoxically, deal with much greater complexity. People are much better able to accommodate detail, switch subjects, and jump from one temporal point in the itinerary to another. This makes for an easier but also more satisfying and substantial interaction which is less socially confusing.

FIG. 6a [left]; b [middle]; c [right]: *Social interaction and communication improvements*

The approach stage as currently configured tends to defeat or confuse the meaning of gaze and gesture as communicative resources. We saw synchronization of gaze, and shared gaze shifts (for example that which occurs between Figures 6a and 6b) which were seamless – both parties appeared to know where to look. These appear to be an indicator of shared understanding of the trajectory of the discussion. We also observed meaningful use of gesture. Parties point things out to draw attention to what they are talking about (Figure 6a), but also to hypothesize. In Figure 6b gesture is used as a way of drawing imaginary possibilities on the basis of what is already there. This turns out to be surprisingly powerful. We observed several sessions in which parties undertook parallel planning imaginatively, that is, they had a plan configured on the map visualization, and were able to imagine alternatives and discuss these coherently despite only being represented internally. This is consistent with Peuqeut and Kraak's (14, *op. cit.*) claim that conventionalized visualizations based on maps can give rise to many kinds of imaginative and creative effect.

Knock-on effects

As well as trying to alleviate some of the problems that occur during the approach stage, we wanted to create a stronger foundation for the rest of the transaction. We were especially interested in the possibility that reconfiguring information resources during one consultation – the approach stage - could have implications for the entire transaction. By looking at customer information access, communication and social interaction problems, we feel we have gone some way to removing a major stumbling block: the loss of a potential sale. Using the system in tandem with an agent, customers can work up an itinerary quickly, understand and remember it, and form a positive relationship with the agent from the off. There is no need to do pre-planning, removing issues to do with lone use of brochures; and no need for agents to have to ask for definite ideas from the customer when s/he may not have these. This means that the basic difficulty at the approach stage – producing input for the remainder of the transaction – is ameliorated. At the trade show we saw episodes where people would take the plan derived from our system to other agents in order to progress it there and then – evidence that the development stage of the transaction can be reached more effectively and in less time.

One of the problems we saw during transactions is the frequent need to do repair work. There are various reasons for this: customers have not chosen an actionable itinerary; and where they have, there may still be issues of update and recap. We have tried to address these issues in the following ways. The system prevents customers making decisions on itineraries that cannot be actioned, since configuration is based on networked information resources which are more accurate and better cross-referenced than brochure materials. In addition, it has been designed to facilitate and encourage parallel planning. Agents and customers can work together on different plans at the same time, and save these. This means that if one itinerary turns out not to be actionable, another, already discussed, can be used. On top of this, the system generates abstract state representations of itineraries, enabling them not only to be saved, but also reloaded, edited and represented in different media. What this means is that a plan can be re-invoked at different times and revised without the need to start from scratch and that the result can be re-represented in whatever medium is appropriate, for example, text. This eases recap and update problems as well as the need to manually translate the itinerary into different media. These innovations provide some examples of how user interfaces can quite radically change the nature of the information processing required by sales transactions.

CONCLUSION

Within the domain of information networking, user interfaces are becoming an increasingly important research area. Back end integration and coordination of information is an essential concern but this effort can be wasted if user interfaces are inadequate. Our research into sales transactions shows that networked information systems are frequently designed to support information processes. The requirements of these processes may not be adequately satisfied where the needs of the different users, who need to work together to provide input, are not satisfactorily met. Here, designing effective user interfaces also becomes an essential issue.

How should effective user interfaces for networked information systems be designed? A basic question is how information can be delivered in a useful, understandable way. Central issues here are information representation, display arrangements and interactivity. Our research tends to go against the view that multiple information visualizations are difficult to cognitively integrate because of search and memory issues. We have been able to reduce these problems through a novel approach: dynalinked, spatially locked information visualizations. Our results suggest that there are interactions between information visualizations, forms of display, and interactivity that can produce important cognitive effects. Further research is needed to understand this, especially in a technological context which supports greater data complexity at the interface (multiple displays and largescreens being two examples), and commercial contexts that would benefit from solutions to problems of complex information representation. However, an immediate implication of our findings is that complexity of information as well as its cognitive integration may be constrained or promoted in different ways given different form factors. This means that user interfaces may not be equally effective where display types differ.

Another central issue in designing user interfaces for networked information systems is human factors, especially where interfaces need to be used collaboratively. We have shown how interfaces that preserve the salience of important social-interactional cues like attention, gaze and gesture can lead to less confusing collaboration. On top of this, our findings have implications for spreading information coordination work at the interface across modalities. Removing representational pressure from the verbal medium and offloading it to the visual allows much richer and more complex planning to go on. This is important in a context where social interaction needs to go smoothly for approach stages in sales transactions to be effective. It suggests that multimodal (as well as multimedia) designs for collaborative user interfaces could be a fruitful direction.

Our research has implications for other domains, wherever collaborative work involving retrieval, discussion and configuration of complex networked information needs to take place around and through the same display. Such domains may include not only sales transactions but also collaborative learning and computer-aided design, to take just two examples. The research presented here shows that in such domains 'networking' information becomes as much a representational issue at the interface as an organisational issue restricted to the back end.

ACKNOWLEDGEMENTS

This paper is dedicated to the late Mike Scaife, who initiated the research. Without his influence and inspiration it would not exist. The author gratefully acknowledges the support of a grant from the UK ESRC/EPSRC PACCIT programme. Thanks to Bridge the World London, especially Jerry Bridge, James Bell, Gillian Hughes and Jo Williams, for providing fieldwork and system evaluation opportunities; and to my research partners Yvonne Rogers (Sussex University), Tom Rodden and Ian Taylor (Nottingham University). Additional thanks go to Yvonne for comments on drafts.

REFERENCES

1. MCCLURE, C. R., and BERTOT, J. C. *Evaluating Networked Information Services: Techniques, Policy and Issues.* Information Today, 2001.
2. BAWDEN, D., and ROWLANDS, I. Digital libraries: developing a conceptual framework. *New Review of Information Networking,* 5, 1999, 71-89.
3. PITTI, D. V. Encoded archival description: an introduction and overview. *New Review of Information Networking* 5, 1999, 61-69.

4. BAKRY, S. H., and AL-DHELAAN, A. H. Planning information networks: the scope and the customer issues. *International Journal of Network Management*, 9, 1999, 28-37.

5. http://www.2020design.com

6. http://www.ehome.ltd.uk

7. TWIDALE, M. B., and NICHOLS, D. M. Collaborative browsing and visualization of the search process. *Proc. ASLIB,* 48, 1996, 177-182.

8. STRAUSFELD, S. Financial viewpoints: using point-of-view to enable understanding of information. *Proc. CHI'95,* 1995, 208-209.

9. SCAIFE, M., and ROGERS, Y. External cognition: how do graphical representations work? *International Journal of Human-Computer Studies,* 45, 1996, 185-213.

10. ZHANG, J., and NORMAN, D. A. Representations in distributed cognitive tasks. *Cognitive Science* 18, 1994, 87-122.

11. SWELLER, J. Cognitive load during problem solving: effects on learning. *Cognitive Science* 12 (2), 1998, 257-285.

12. MACKINLAY, J. D. Automating the design of graphical representations of relational information. In S. K. Card, J. D. MacKinlay, and B. Schneiderman (Eds) *Information Visualization: Using Vision to Think*. Morgan Kaufmann, 1999, 66-82.

13. COX, R., and BRNA, P. Reasoning with external representations: supporting the stages of selection, construction and use. In P. Brna, S. Ohlsson and H. Pain (Eds) *Proceedings of the World Conference on Artificial Intelligence in Education*, Charlottesville, VA: Association for the Advancement of Computing in Education, 1993, 185-192.

14. PEUQUET, D. J., and KRAAK, M-J. Geobrowsing: creative thinking and knowledge discovery using geographic visualization. *Information Visualization,* 1, 2002, 80–91.

15. GRUDIN, J. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. *Proc. CHI'01,* 2001, 458-465.

16. STEWART, J., BEDERSON, B. B., and DRUIN, A. Single display groupware: a model for co-present collaboration. *Proc. CHI'99,* 1999, 286-293.

17. STREITZ, N. A., GEIßLER, J., HOLMER, T., KONOMI, S., MÜLLER-TOMFELDE, C., REISCHL, W., REXROTH, P., SEITZ, P., and STEINMETZ, R. i-LAND: an interactive landscape for creativity and innovation. In *Proc. CHI'99,* 1999, 120-127.

18. DIETZ, P., and LEIGH, D. DiamondTouch: a multi-user touch technology. *Proc. UIST'01,* 2001, 219-226.

19. HASCOET, M., and SACKS, F. Exploring interaction strategies with Wall-screen. *Proc. IV'02,* 2002, IEEE, 719-724.

20. ROGERS, Y., and LINDSLEY, S. Vertical versus horizontal surfaces: linking, integrating and interacting with information across shared displays. Unpublished manuscript, 2003.

21. ROGERS, Y., and SCAIFE, M. How can interactive multimedia facilitate learning? CDROM/WWW *Proc. Multimedia Interfaces,* AAAI. Menlo Park, CA, 1999.

22. COX, R. Analytical reasoning with multiple external representations. PhD thesis, University of Edinburgh, 1996.

23. ABNETT, A., STANTON, D., NEALE, H., and O'MALLEY, C. The effect of multiple input devices on collaboration and gender issues. *Proc. Euro-CSCL'01*, 2001, 29-36.

24. SCAIFE, M., HALLORAN, J., and ROGERS, Y. Lets work together: supporting two-party collaborations with new forms of shared interactive representations. *Proc. COOP'02, International Conference on Cooperative Systems Design,* 2002, 123-138.

25. BOWERS, J., and CHURCHER, J. Local and global structuring of computer mediated communication: developing linguistic perspectives on CSCW in Cosmos. *Proc CSCW'98,* 1998, 125-139.

26. COULTHARD, M. *Discourse Analysis,* Longman, 1997.

27. GOFFMAN, E. Replies and responses. *Language in Society* 5 (3), 1976, 257-313.

28. SACKS, H., SCHEGLOFF, E., and JEFFERSON, G. A simplest systematics for the organization of turn-taking for conversation. *Language* 50 (4), 1974, 696-735.

29. HUGHES, J., O'BRIEN, J., RANDALL, D., RODDEN, T., ROUNCEFIELD, M., SOMMERVILLE, I., and TOLMIE, P. Getting to know the 'customer in the machine'. *Proc. Group'99,* 1999, 30-39.

30. GOFFMAN, E. *The Presentation of Self in Everyday Life.* Allen Lane, 1959.

31. HEATH, C. *Body Movement and Speech in Medical Interaction.* Cambridge University Press, 1986.

32. GREATBATCH, D., LUFF, P., HEATH, C., and CAMPION, P. Interpersonal communication and human-computer interaction: an examination of the use of computers in medical consultations. *Interacting with Computers* 5 (2), 1993, 193-216.

33. MYERS, G. Analysis of conversation and talk. In M. W. Bauer and G. Gaskell (Eds) *Qualitative Researching With Image Text and Sound: A Practical Handbook*. Sage, 2000, 191-206.

34. PREECE, J., ROGERS, Y., and SHARP, H. *Interaction Design.* Wiley, 2002.

35. HALLORAN, J. Can we afford it? Issues in designing transparent technologies. *Proc. CT'01, 4th International Conference on Cognitive Technologies,* 2001, 141-149.

36. RODDEN, T., ROGERS, Y., HALLORAN, J., and TAYLOR, I. Designing novel interactional workspaces to support face to face consultations. To appear in *Proc. CHI '03,* 2003.