Communication of Knowledge from Synthesised Web Sites

1 Previous Track Record

This project combines expertise from the Glasgow Interactive Systems group (GIST) at Glasgow University; the Human Communication Research Centre (HCRC) at Edinburgh; and the Software Systems and Processes group at Edinburgh. We survey below the track record of the project members from each site.

Chris Johnson is a Professor of Computing Science at Glasgow University and is part of the multi-disciplinary, Glasgow Interactive Systems group (GIST). GIST has recognised strengths in the development of animation tools (Patterson) as well as auditory and haptic displays (Brewster). We also have an international reputation for the evaluation of computer assisted learning (Draper), software architectures for interactive systems (Gray) as well as information retrieval and visualisation systems (Chalmers and Dunlop). Prof Johnson's work has focussed on techniques that can be used to represent and reason about the human contribution to major accidents. His work on improved techniques for accident investigation was cited in a Queen's Award for Teaching and Research (1998) and an award from the International Systems Safety Society (1998).

Prof Johnson also has an established record at the intersection of information retrieval and user interface design. He was an invited speaker at the first workshop on Human Factors on the Web in 1995, this talk was subsequently published as [11]. He was co-chair of the first two joint workshops organised by the British Computer Society's Information Retrieval SIG and the British HCI Group. The proceedings of these meetings have been published in four numbers of Interacting with Computers and a special edition of Information Processing and Management [14, 4].

He has been involved in three research council grants that contribute to this proposal. The Temporal Aspects of Usability project involved Steve Draper, Paddy O'Donnell, Phil Gray and David England (Joint Council Initiative in HCI and Cognitive Science, Grant No. 9201233). This focussed on the strategies that users exploit in order to combat retrieval delays over low bandwidth and unreliable networks [8, 10].

Temporal And Graphical Primitives For Declarative Graphics (EPSRC Grant No. GR/K69148) forms a bridge between the Glasgow work and that of the Department of Artificial Intelligence in Ediburgh. This project has used logic-based programming languages as well as parallel functional systems to develop prototypes for the future generation of datalink Air Traffic Control systems, in collaboration with the ESPRIT MEFISTO project and National Air Traffic Services [26].

Linking User and System Models to Analyse the Causes of Major Accidents, UK (EPSRC Grant No. GR/L27800) provides further links with the work of the HCRC in Edinburgh. This project has been using higher-order logics to construct cognitive models that support an analysis of the human contribution to systems failure. The current proposal builds on this research because there must be some means of representing and reasoning about the psychological percursors to accidents, as well as the more observable events that occur during major failures.

John Lee is deputy director of the Human Communications Research Centre at Edinburgh. The HCRC is an Interdisciplinary Research Centre which focuses on cognitive science approaches to human communication. The Centre integrates computational, psychological, linguistic and logical methods in theories of human information processing. Several research themes at HCRC involve the design, analysis and evaluation of human responses to information in multiple media and modalities.

Projects at HCRC that have considered information presentation in multiple media and modalities include: the EU funded project GRACE (Graphical Communication in HCI); the ESRC funded project VICAR (The Vicarious Learner); and the EPSRC funded projects ILEX (The intelligent labelling explorer), Communication in Safety Cases and Understanding Software Architectures.

GRACE involved study of the appropriateness of various interface modalities for presenting any particular piece of information. This study considered the nature of the information, as well as the culture, psychology and physiology of the intended user. The work in HCRC focused on three areas: theoretical groundwork on the concepts involved in multi-modal interfaces and the semantics of modalities, taxonomising graphical representations and building an interface toolkit. An experimental evaluation of how people naturally classify representations was carried out using statistical clustering techniques, and a taxonomy workbench tool to help viewing and classifying images was described and partially developed.

VICAR focused on issues in learning and cognitive modelling. In particular, VICAR carried out a number of empirical studies of the effectiveness of various multimedia presentations of information for distance learning.

ILEX studied the automatic generation of web-based descriptions of objects in a museum gallery. Descriptions were generated dynamically, utilising theories of natural language generation, and tailored to particular users. Detailed evaluation by both casual users and domain experts (museum curators) was an important aspect of ILEX.

Corin Gurr is an RA at HCRC who worked on the Communication in Safety Cases project (EPSRC/DTI grant GR/J/58619) - an interdisciplinary project concentrating on the problems of communication between experts in the domain of safety critical systems. His previous research on this project focused on problems of communication between experts in differing technical disciplines. Important elements of this work included: understanding how various stakeholders in the development and assessment process with varying motivations and domains of expertise are able to reach agreement on safety and economic aspects of systems; a theory explaining the appropriateness of various modalities (such as text, logic and graphics) for representing complex information; and case studies of industrial best practice. These research themes have been carried on to the current EPSRC funded Understanding Software Architectures project.

David Robertson is the leader of the Software Systems and Processes group in Informatics at the University of Edinburgh. His group specialises in the use of formal methods to assist in the modelling of problems and the synthesis of specifications. The work of the group combines ideas from formal methods, human factors and target application domains. We are currently working on topics such as: formal support for business process modelling ([2]); the relationship between logic and diagrams ([21]); and the use of dialogue in describing requirements ([15]). Edinburgh has one of the longest traditions of study of domain-specific synthesis of specifications (see for example [25],[24]) - a subject which is enjoying a renaissance in the formal methods community (see [18], [19]) and in the software engineering community (see the Software Engineering and Programming Languages subgroup report of the ACM Strategic Directions in Computing Research at `http://www.medg.lcs.mit.edu/sdcr/reports`). This is one of the foundations of our proposal.

Robertson has been a grant-holder on EPSRC projects which contribute to the current proposal. The Techniques Editor project (EPSRC/ESRC/MRC grant G9030396) explored the use of structured editing methods based on conceptual rather than purely syntactic components to help novices understand how to build logic programs. The Communication in Safety Cases project (mentioned earlier) studied forms of communication used to connect different technical communities when developing a safety case for a system. One of the results from this project was a system which combined automated synthesis of specifications for safety shutdown systems with design constraints derived from textual guidelines for best practice ([6]). Robertson's work on the construction of specifications in a distributed style ([23]) is currently being used in Spain (supported by the Spanish government) to develop support for modelling of process and inspection in part of their meat industry. He is also involved in a joint project with Brazil (supported by the British Council) on the logical foundations of knowledge sharing. He currently holds an EPSRC Advanced IT Fellowship on the application of techniques from formal specification and knowledge based systems to the problem of requirements capture for logic programs.

The research proposal of Section 2 is predicated on the idea that it is practical to separate, formally, the information content of Web sites from their presentational form and to derive from from content via automated synthesis. Robertson's group have demonstrated that this can work by using a domain-specific synthesiser to produce and maintain their public Web site (synthesised site at `http://www.dai.ed.ac.uk/groups/ssp/index.html`). We claim that the cost of producing the synthesis system was recovered in the first three months of its operation (via savings in routine maintenance) and we believe we are the first to have demonstrated that this idea can be worthwhile. One of the aims of our current proposal is to gain additional benefit from the information produced by this sort of synthesis mechanism in order to support better communication between site designers.

2 Description of the Proposed Research and its Context

2.1 Background

All of the six strategic themes highlighted by the Human Computer Interaction subgroup of the ACM Strategic Directions in Computing Research (`http://www.medg.lcs.mit.edu/sdcr/reports`) relate directly to the Worldwide Web. A single issue underlies three of these themes: how to access distributed sources of knowledge and make use of it in human terms. Much effort is being spent on innovative ways of browsing the Web; visualising information; and standardising languages for knowledge sharing. These tend to be driven by the current language standards and technology (languages like HTML, XML and VRML in the visualisation communitiy; standards like UML, KIF and PIF in the knowledge and software engineering communities). Sharing of the knowledge is done in ignorance of the way in which the knowledge has been engineered. We know from "pre-Internet" experience that this doesn't work well. We need to understand how a repository of knowledge has been built in order to guess what we might find in it (and where it might be found if the repository is large). To use the information we find we must be able to understand it in the context of our own work, which may require re-interpretation of its original form. To re-interpret reliably we need to know how the original interpretation was made. Information at this kind of level can be encoded in various of the languages mentioned above. One of the recently most popular, XML, is sometimes seen as just a rather more general replacement for HTML, to be interpreted as presentation markup in browsers. However, it is actually a quite abstract means of structuring information and attaching further levels of specification to it. Often, XML markup is applied to free text, but it need not be – it can be combined with material already structured in some other formal way. A framework like this can be used to add information to a representation of almost any kind, at any stage in its construction or use, indicating further details about its context, the interpretation applied to it, etc. In this proposal, we do not consider ourselves committed to a particular standard format, but will evaluate XML and a range of other such languages for adequacy in terms of our objectives in knowledge-sharing.

An attractive way of supporting this form of communication is as a by-product of the synthesis of Web sites. This is analogous to the synthesis of other forms of software from domain-specific problem descriptions. The basic architecture for such systems is as shown below

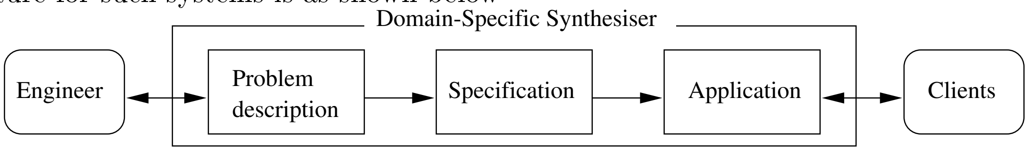

where an engineer interacts with a domain-specific problem description system to build a description of the problem which is used (again with the help of domain-specific synthesis strategies) to build a specification of the solution in a restricted formal language. This specification is then used to generate an application in a target language appropriate to the needs of a target group of clients. There are many applications of this form of domain-specific synthesis (for example, [18], [19]) but the human factors of using such systems are seldom explored beyond the point of getting some sort of interactive system running. This was (arguably) not such a high priority for synthesis systems which work within a closed domain but becomes the highest priority when we use domain-specific generators to assist in sharing information in a highly distributed system like the Worldwide Web. To our knowledge, nobody has investigated this problem in depth. We plan to combine the domain-specific synthesis expertise from the Institute for Representation and Reasoning at Edinburgh with the HCI expertise at Glasgow and the Human Communications Research Centre at Edinburgh in order to do that (see Section 1).

2.2 Programme and Methodology

The hypotheses tested by this project are:

- That for significant domains of application it is practical to build domain-specific languages in which core domain knowledge can be represented in the form of visualisation-neutral problem descriptions.

- That such domain-specific formal languages can be used by semi-automatic generation mechanisms to construct Web sites for visualising problem descriptions from differing points of view.
- That the structure provided by automated generation can provide a skeleton which human Web site designers can embellish with additional material aiding communication in a controlled manner.
- That maintenance of sites can be reduced by pushing as much of the revision process as possible to the level of the problem description language.

If these hypotheses are supported our view of Web site design is fundamentally changed but to investigate them we need an appropriate application domain; a trusted method of problem description and synthesis; and convincing experimental analyses. We describe each of these below.

2.2.1 Application Domain: Accident Reporting

Web sites are increasingly replacing the dissemination of accident reports through conventional, paper-based documents. Unfortunately, most investigation authorities have insufficient resources to best exploit the visualisation and presentation opportunities of the new media. They simply provide electronic versions of the text-based document. Occasionally hypertext links are provided within single reports. There are, as yet, no on-line examples of accident reports that contain hypertext links between incidents. This is a significant limitation because many people have argued that designers must have a clear understanding of common causes between multiple failure if they are to prevent future accident and incidents [22]. This analysis has been confirmed by the initial results of the first international survey into the usability of on-line accident reports [27]. Accident reporting, therefore, satisfies an important criteria for our application domain: the effective presentation and visualisation of critical information is hindered by the costs of formatting and tailoring the data held by the organisation.

Accident reports can be supported by a wide range of visual styles of communication, each giving different views of an accident. Four of these are given below. Figure 1 illustrates how differing views can be presented of the same accident. On the left we have a Virtual Reality Markup Language (VRML) model of the events leading to the sinking of the Herald of Free Enterprise. We can see three dimensional views of timelines specific to particular entities in the accident. Users can explore the model so that by looking horizontally across the other lines it is possible to get an overview of the events that affected other agents in the accident at that moment in time. On the right we have a different view of the accident, showing the locations of key physical features on an image map of the ship[1]. Figure 2 gives examples of more detailed analyses of accidents. The diagram on the left is a fault tree that has been augmented to capture human error as well as system failure [17]. A limitation with this approach is that fault trees provide an extremely poor representation of concurrent events. This is significant if visualisations are to represent the ways in which different actors simultaneously contribute to major accidents [9]. The diagram on the right presents a Petri Net model of the Kegworth air traffic accident; this notation was specifically developed to capture properties of concurrent and distributed systems. In summary, multiple views are possible of each accident report - each selective in the detail which it presents and in the features which it emphasises - but information is shared (and should be cross-referenced) between these views.

2.2.2 Automated Synthesis from Problem Descriptions

The view of synthesis taken by the project is best explained using a simplified example. Figure 1 shows two views of the Herald of Free Enterprise accident. This section gives a flavour of how these two views could be synthesised from a problem description, using much simplified methods to allow us to give detail in a short space. First we need a formal language for describing the essential features of the accident. This must talk about propositions which hold at particular times (or intervals of time) involving types of object. A sample of these (with accompanying English text) is:

[1] Note that the Herald had rudders and propulsion at both ends so what might look like stern gear is really on the bow, as shown

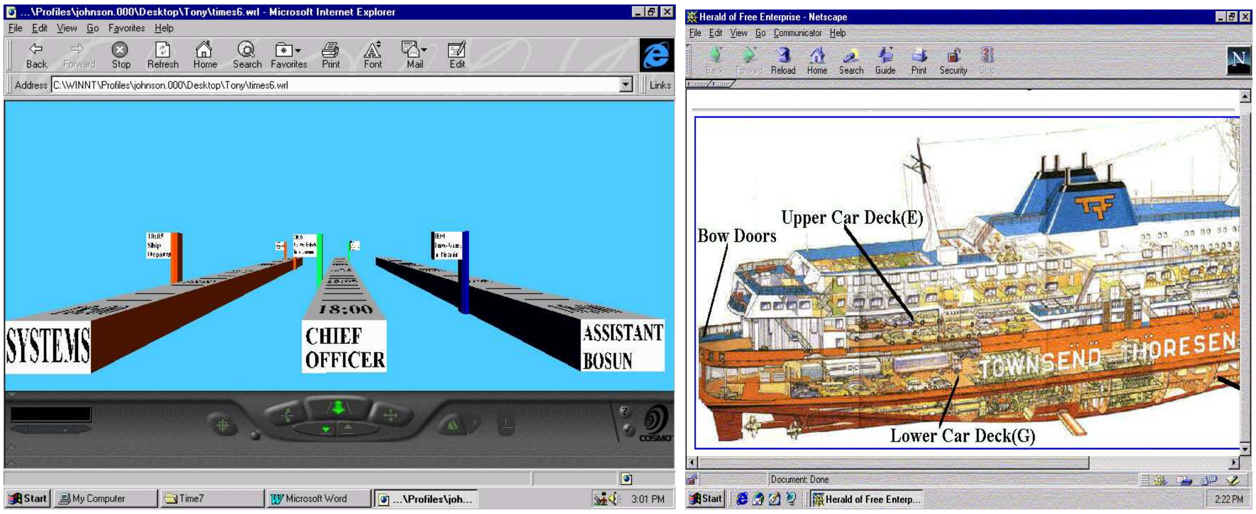

Figure 1: Two views of the Herald of Free Enterprise Accident

Proposition	Informal description
$left\_open(bow\_doors)@i(-\infty, \infty)$	The bow doors were left open.
$passed\_outer\_mole(herald)@18.24$	The Herald passed the outer mole at 18.24.
$passed\_outer\_mole(herald)@18.23$	The Herald passed the outer mole at 18.23.
$left\_for(chief\_officer, bunk(chief\_officer), bridge)@18.24$	The chief officer left for the bridge from his bunk at 18.24.
$thrown(assistant\_bosun, bunk(assistant\_bosun))@18.28$	The assistant bosun was thrown from his bunk at 18.28.
$stopped(bridge\_clock)@18.28$	The bridge clock stopped at 18.28.

Where $P@T$ denotes that a proposition, P, holds at time T and T can be either a particular time point or an interval, $i(T_1, T_2)$, with the symbol $-\infty$ being the earliest possible time and ∞ the latest. Having defined our problem in this formal notation we now need to utilise this for synthesising views of the accident. Taking the left-hand one from Figure 1 our aim is to generate time lines for stipulated entities (or entity groups) in the accident. We have two stipulated individual entities ($chief\_officer$ and $assistant\_bosun$) and one stipulated entity group ($system$) which includes the individual entities bow_doors, $bridge\_clock$ and $herald$. A basic algorithm for generating the timelines (stated here in English but easy to define formally) is the following:

- Construct the timeline for each stipulated entity/group by:
  - Finding all the propositions relevant to the entity/group.
  - Finding the subset of these which have more than one given time of occurrence ($passed\_outer\_mole$ is like this in our example). These will be denoted by "floating" flags on the timeline.
  - The remainder of flags are "non-floating".
  - Construct a 3-D timeline with floating and non-floating flags ordered along the bar.
  - For each flag, remember the correspondence between formal proposition and flag for possible later cross-referencing.
- Lay the timelines side by side in the 3-D visualisation.

The algorithm above allows automatic generation of the timeline visualisation (given default assumptions about shape and orientation of timelines in VRML). However, the right-hand view of Figure 1 cannot be generated completely automatically - we can only provide support in marking up location information in the image-map of the ship by generating a set of all the $system$ components mentioned in our formal description and prompting the designer to add them. When the designer adds components, our support system remembers

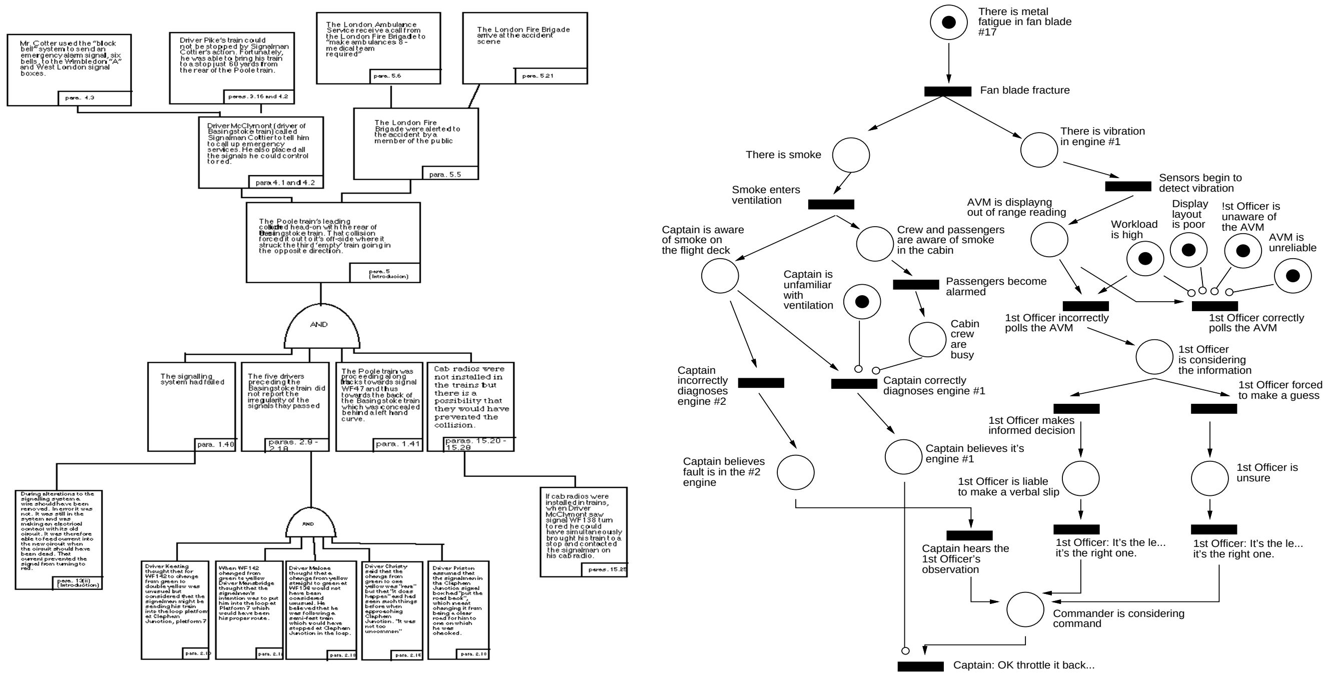

Figure 2: Two different models of causality

the correspondence between areas of the image and the objects in our formal problem description. This leaves us with the two views of Figure1.

This form of synthesis has a direct benefit to designers by automating part of design. However it has a number of indirect benefits and these are the main focus of our grant proposal. Four of these are:

- Cross referencing between views can be done automatically for those parts of the views which were derived from the formal problem description. For instance we could automatically install hyperlinks between the *system* components in the image-map view and the flags on the timeline which refer to those components (these would be inserted for each of the flags on the *system* timeline and one of the flags on the $chief\_officer$ timeline).

- Some types of change to the visualisations can be done very efficiently by altering the formal problem description and reconstructing (automatically) the views. For example, if we want to change the name of the ship from *herald* to $herald\_of\_free\_enterprise$, to avoid confusion with other ships in the "Herald" series, we can do that by altering the problem description and propagating the changes to flags (in the timelines) and labels (in the ship diagram) rather than the error-prone business of editing each view by hand.

- We can add more sophisticated algorithms to our synthesis mechanisms which guarantee that key information is available to those navigating the sites. For instance, we might want people looking at the timelines view to be able to stipulate a time interval and generate, on demand, a slice across the time lines for that interval, with appropriate cross referencing to other views. This sort of navigational support is difficult to engineer, and even more difficult to maintain, without an underlying formal model.

- The design of new Web sites is often spurred by seeing an existing one. If the existing one is derived by synthesis from a formal description then, as in maintenance, reuse may be supported by adapting the problem description rather than by hacking the visual presentation of the site directly. For example, if we want to re-use the timelines from our example to a new type of accident then we first ask "Can the new accident be described formally in the notation used in our original example?". If it can, then re-use is most efficiently done by writing a new formal description (in the existing language) and re-applying the existing generation mechanism.

It does not seem contentious that these would be significant benefits for Web site designers and that they introduce a novel way of thinking about design and maintenance in this domain. However, there are unresolved human factors (as well as knowledge representation) issues in obtaining these benefits. In particular, it is likely that numerous forms of presentation would be possible for the same problem description. These would show different areas of the problem (as in the two pictures of Figure 1 and differing levels of detail (compare Figure 1 to Figure 2). Site designers will need guidance in selecting and applying synthesis methods, based on the properties known from the problem description and compatible with the current (partially) synthesised site design.

2.2.3 Analytical Experiments

Our analysis of the generation method and its outputs is based on resources and expertise derived from previous projects at Glasgow and the HCRC at Edinburgh. These cover three basic aspects of the problem: access to domain experts; theories predicting desirable properties of visualisations; and design of Web-based evaluations based on established mehtods. We summarise each of these below.

The Glasgow group already possesses connections with safety assessment experts, including the UK Health and Safety Executive, the European Air Traffic Control Agency and the US Chemical Safety and Hazard Investigation Board, who have expressed willingness to evaluate the synthesised systems. Prior HCRC experience on the ILEX project [3] (where museum curators assessed the effectiveness of automatically generated web based descriptions of a gallery of museum exhibits) suggests that maximal benefit of evaluations by domain experts will be gained through the use of structured interviews, carried out by a specialist.

Evaluation of the appropriateness of particular styles of communication for given communicative tasks and audiences forms an ongoing theme of research at HCRC. Of particular relevance are cognitively based theories [28, 5] which predict the suitability of differing representational styles. A suitable representation is one which accurately and completely captures the structure of the domain and task for which it is intended. Varieties of languages (including textual, formal and diagrammatic) have certain *intrinsic* properties, either through form or conventions of use, and effective representations will exploit these inherent properties in capturing the structure of a represented artifact. We envisage these theories being applied to both inform the automated synthesis and to suggest experimental hypotheses concerning the relative effectiveness of alternative synthesised representations.

A series of empirical experiments will test hypotheses comparing the effectiveness of both alternative synthesised web documents, and of synthesised versus existing accident reports. In the first instance these hypotheses will be investigated through readily measurable properties, such as time to access information and ability to collate distributed knowledge. The HCRC has experience in this form of evaluation from several projects, including ILEX (The intelligent labelling explorer), GRACE (Graphical Communication in HCI) and VICAR (The Vicarious Learner). On ILEX natural langauge generation techniques were used to produce web documents and comparisons were made [20, 1] between dynamically and statically generated versions of such documents. For GRACE discourse theories were applied to aid navigation in hypertext documents, and these were assessed in [7]. The VICAR project explored the use of web based tertiary information in education. As part of sthis project a study of the effectiveness of the assignment of information to differing media is reported in [16]. The results of the evaluations that we perform for this project will indicate the most effective synthesised web documents, which will then be further evaluated by domain experts. The Glasgow group has worked extensively on evaluation techniques for web based visualisations, including VRML [12, 13]. As mentioned, it has also conducted the first international survey into the usability of existing web based accident reports. It is, therefore, ideally placed to evaluate the effectiveness of both manually prepared web documents and the synthesised forms produced by this project. We already have provisional offers of assistance in this from AerLingus, NASA, and the HSE. The Glasgow group has also evaluated web development tools; this supports the higher level analysis of whether accident investigators might be able to exploit the tools being developed by the Edinburgh group.

2.3 Relevance to Beneficiaries

Our results are of relevance to the following communities:

- The visual reasoning community has deepened our knowledge of the semantics of individual forms of communication but has less to say about the way in which these may be combined - a necessity in distributed systems like the Web. Our work will help to strengthen this aspect of academic understanding.
- A number of projects have studied the generation of visualisations from formal descriptions. These range from automatic site generation systems (from logic at Edinburgh, by more traditional means at sites such as `http://www.otal.umd.edu/Olive`) to text generation systems (such as the Edinburgh ILEX project mentioned earlier) to generating high-coherence documents from formal specifications (such as the work of Thimbleby and Ladkin). Our work complements these by studying more closely the engineering methods used in this kind of system and the issues of multi-modality which arise when more than one visualisation of a formal description is possible.
- One of the most difficult practical problems for Web site designers is the separation of content (the information presented) from form (the HTML, XML or VRML code describing what those accessing the information will see). We know of no commercial tool which comes close to the separation between problem description, specification and visualisation which we envisage. Nor, we believe, is industry likely to follow this route unless clear demonstrations of its practicality are produced by academia. We aim to provide these.

2.4 Dissemination and Exploitation

The project will begin with a workshop that is intended to forge links between researchers in the areas of visualisation and problem synthesis, and companies that are involved in accident analysis. This will build upon the results of a series of similar inter-disciplinary workshops that Johnson has organised with Prof Nancy Leveson (MIT). The initial project meeting is primarily intended as an elicitation exercise. It will also establish contacts to be maintained throughout the duration of the research. The exploitation of the work will be supported by these personal contacts.

We will also disseminate our work through a number of national and international conferences. We will target meetings that attract both academic and commercial participation, such as the bi-annual Interact conference. Our programme includes three main deliverables. The first paper will describe the application of our approach to examine the visualisation of the Sheen report into the Herald of Free Enterprise disaster. The second paper will be used to determine whether our approach can be successfully applied to provide alternaive views of the communications problems that exacerbated the Kings Cross accident. We will also produce a final report that performs a comparative analysis and identifies directions for future research. Copies of our progress reports will also be distributed to the regulatory and commercial organisations that attended the initial meeting.

We will host a workshop in the last year of the project. This forum will be used to distribute information about our inter-disciplinary techniques for accident reporting. A final report will be produced to record the more general findings of our research. This will be supplied to the participants as a means of encouraging the exploitation of our work.

2.5 Justification of Resources

The project requires a mixture of software synthesis and human factors skills which it is rare to find in a single researcher. Therefore we are applying for a research assistant at Edinburgh and two research students at Glasgow. To maintain a shared view of the project will require monthly project meetings. We are also asking for computing equipment to support video-conferencing and Email between sites, as well as supporting the development of synthesis, visualisation and Web server facilities essential to the project. Some specialist software is required to support synthesis (SICStus Prolog) and visualisation (a VRML editing tool and Java developers kit). We intend to present our results at relevant national/international conferences (see Section 2.4) and plan two workshops centred on the project (see Section 3). We have also applied for funds to cover the cost of computing and secretarial support at each site.

3 Workplan

A project workplan appears below, with tasks on the left run by Edinburgh and those on the right by Glasgow. Tasks spanning the diagram are run jointly because they require closer collaboration. The general aim is to use the first year to build the prototype synthesis mechanisms and work out, in theory, the visualisations which should be supported by them. In the second year we build the visualisation systems and evaluate the sites built using them. The results of this are used in the third year to build an improved synthesis system and perform further usability experiments.

The arrows from Glasgow to Edinburgh tasks indicate where results of Glasgow analyses feed into Edinburgh design. This and the joint work in the second year will require good channels of communication between the two sites. We shall maintain these by a combination of video-conferencing (now feasible at very low cost) and monthly whole-project meetings. Workshops at the end of the first and second years (organised along the lines of Johnson's BCS meetings described in Section 1) will provide points of contact with the wider HCI and software synthesis communities.

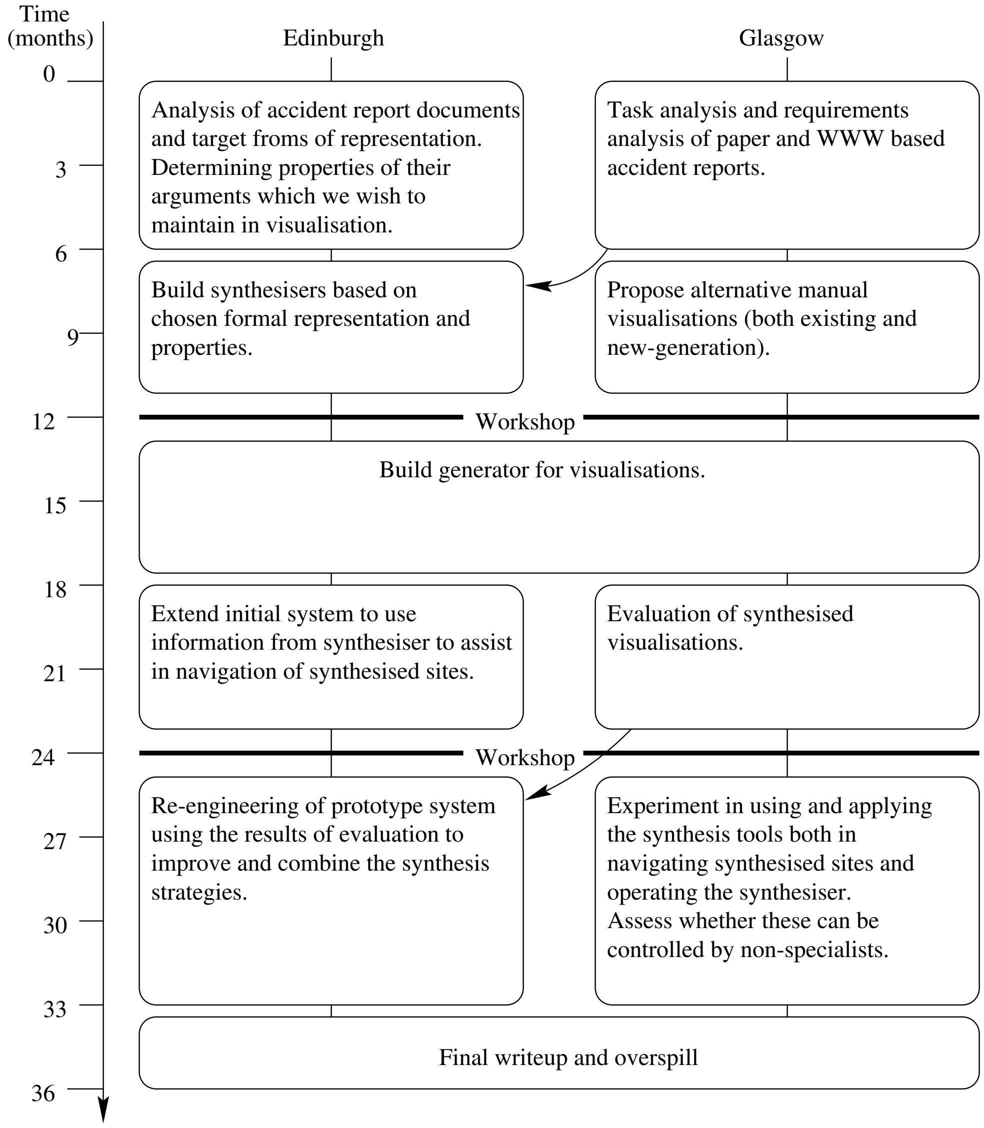

References

[1] E Amitay and J Oberlander. Convention says... In *Proceedings of the Flexible Hypertext Workshop (held in conjunction with Hpertext'97)*, Southampton, 1997.

[2] J. Chen-Burger and D. Robertson. Formal support for an informal business modelling method. In *Proceedings of the 10th International Conference on Software Engineering and Knowledge Engineering, San Francisco, USA*. Knowledge Systems Institute, Illinois, 1998.

[3] R Cox, M O'Donnell, and J Oberlander. Dynamic versus static hypermedia in museum education: an evaluation of ilex, the intelligent labelling explorer. In *Proceedings of the 9th International Conference on Artificial Intelligence and Education*, Le Mans, France, 1999. To appear.

[4] M.D. Dunlop, C.W. Johnson, and J. Reid. Exposing the layers of information retrieval evaluation. *Interacting with Computers*, 3(10):225–237, June 1998.

[5] C Gurr, J Lee, and K Stenning. Theories of diagrammatic reasoning: distinguishing component problems. *Mind and Machines*, 8(4):533–557, December 1998.

[6] J. Hesketh, D. Robertson, N. Fuchs, and A. Bundy. Lightweight formalisation in support of requirements engineering. *Journal of Automated Software Engineering*, 5(2):183–210, 1998.

[7] R Inder and J Oberlander. Using discourse to aid hypertext navigation. In *Proceedings of HCI International*, Yokohama, 1995. Elsevier.

[8] C.W. Johnson. Electronic gridlock, information saturation and the unpredictability of information retrieval over the world wide web. In P. Palanque and F. Paterno, editors, *Formal Methods in Human Computer Interaction: Comparison, Benefits, Open Questions*, pages 261–282. Springer Verlag, Berlin, Germany, 1997.

[9] C.W. Johnson. The epistemics of accidents. *Journal of Human Computer Systems*, 47:659–688, 1997.

[10] C.W. Johnson. The impact of utility and time on distributed information retrieval. In H. Thimbleby, B. O'Conaill, and P. Thomas, editors, *People and Computers XII: Proceedings of HCI'97*, pages 191–204, London, United Kingdom, 1997. Springer Verlag.

[11] C.W. Johnson. Ten golden rules for video over the web. In J. Ratner, E. Grosse, and C. Forsythe, editors, *Human Factors for World Wide Web Development*, pages 207–224. Lawrence Erlbaum, New York, United States of America, 1997.

[12] C.W. Johnson. On the problems of validating desktopvr. pages 326–338, London, United Kingdom, 1998. Springer Verlag.

[13] C.W. Johnson. Improving the presentation of accident reports over the world wide web. 1999.

[14] C.W. Johnson and M.D. Dunlop. Subjectivity and notions of time and value in information retrieval. *Interacting with Computers*, 10(4):67–75, 1998.

[15] R. Lecoeuche, C. Mellish, and D. Robertson. A framework for requirements elicitation through mixed-initiative dialogue. In *3rd IEEE International Conference on Requirements Engineering, Colorado Springs, USA*, pages 190–196. IEEE Computer Society, April 1998. ISBN: 0-8186-8356-2.

[16] J Lee, J McKendree, F Dineen, R Cox, T Mayes, and K Stenning. Modality, representation and educational technology. In *Workshop on Interpretation and Generation in Intelligent Multimodal Systems and Graphical Reasoning in Expert Systems, 4th World Congress on Expert Systems*, Mexico City, 1998. ITESM, Mexico.

[17] L. Love and C.W. Johnson. Using extended fault-trees in conjunction with traditional accident reports. In S. Howard, J. Hammond, and G. Lindgaard, editors, *Interact'97*, pages 102–103, London, United Kingdom, 1997. Chapman and Hall.

[18] M. Lowry and J. Van Baalen. Meta-amphion: Synthesis of efficient domain-specific program synthesis systems. *Automated Software Engineering*, 4:199–241, 1997.

[19] Luqi and J.A. Goguen. Formal methods: Promises and problems. *IEEE Software*, 14(1):73–85, 1997.

[20] J Oberlander, M O'Donnell, A Knott, and C Mellish. Conversation in the museum: experiments in dynamic hypermedia with the intelligent labelling explorer. *New Review of Hypermedia and Multimedia*. In press.

[21] J. Puigsegur, J. Agusti, and D. Robertson. A visual logic programming language. In *Proceedings of the 12th IEE Symposium on Visual Languages*, Colorado, 1996.

[22] J. Reason. *Managing the Risks of Organisational Failure*. Ashgate, Aldershot, United Kingdom, 1997.

[23] D. Robertson. Distributed specification. In *Proceedings of the 12th European Conference on Artificial Intelligence*, Budapest, Hungary, August 1996.

[24] D. Robertson and J. Agusti. *Software Blueprints: Lightweight Uses of Logic in Conceptual Modelling*. Addison Wesley/ACM Press, 1999. ISBN 0201398192.

[25] D. Robertson, A. Bundy, R. Muetzelfeldt, M. Haggith, and M Uschold. *Eco-Logic: Logic-Based Approaches to Ecological Modelling*. MIT Press (Logic Programming Series), 1991. ISBN 0-262-18143-6.

[26] M. Sage and C.W. Johnson. Interacting with haggis: Implementing agent based specifications in a functional style. In S. Howard, J. Hammond, and G.Lindgaard, editors, *Interact'97*, pages 126–133, London, United Kingdom, 1997. Chapman and Hall.

[27] P. Snowdon and C. Johnson. An international survey into the usability of accident reports. 1999.

[28] K Stenning, R Cox, and J Oberlander. Contrasting the cognitive effects of graphical and sentential logic teaching: reasoning, representation and individual differences. *Language and Cognitive Processes*, 10, 1995.