Knowledge Media Institute

# Graphical Argumentation and Design Cognition

*Simon J. Buckingham Shum, Allan MacLean, Victoria M.E. Bellotti & Nick V. Hammond*

**KMI-TR-25**

**May, 1996; updated February, 1997**

*Human-Computer Interaction (in press)*

[http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-25-abstract.html]

The Open University

Buckingham Shum, S.J., MacLean, A., Bellotti, V.M.E., & Hammond, N.V. (in press)
Graphical Argumentation and Design Cognition. To appear in: *Human-Computer Interaction*

# Graphical Argumentation and Design Cognition

**Simon J. Buckingham Shum**
Knowledge Media Institute, The Open University,
Walton Hall, Milton Keynes, MK7 6AA, U.K.
Tel: +44 1908-655723
Fax: +44 1908-653169
S.Buckingham.Shum@open.ac.uk

**Allan MacLean**
Rank Xerox Research Centre, 61 Regent Street,
Cambridge, CB2 1AB, U.K.
maclean@cambridge.rxrc.xerox.com

**Victoria M.E. Bellotti**
Apple Research Laboratories, Apple Computer Inc.,
One Infinite Loop, MS: 301-4A, Cupertino, CA 95014, U.S.A.
bellotti@apple.com

**Nick V. Hammond**
Department of Psychology, University of York,
Heslington, York, YO1 5DD, U.K.
nvh1@york.ac.uk

## ABSTRACT

Many efforts have been made to exploit the properties of graphical notations to support argument construction and communication. In the context of design rationale capture, we are interested in graphical argumentation structures as cognitive tools to support individual and collaborative design in real time. This context of use requires a detailed understanding of how a new representational structure integrates into the cognitive and discursive flow of design, that is, whether it provides supportive or intrusive structure. This paper presents a use-oriented analysis of a graphical argumentation notation (QOC). Through a series of empirical studies, we show that it provides most support when elaborating poorly understood design spaces, but is a distraction when evaluating well constrained design spaces. This is explained in terms of the cognitive compatibility between argumentative reasoning and the demands of different modes of designing. We then provide an account based the collaborative affordances of QOC in group design meetings, and extend this to discuss the evolution of QOC argumentation from short term working memory to long term group memory.

Simon Buckingham Shum is a cognitive ergonomist with interests in the interactive affordances of multimedia, and the human dimensions to organisational memory technologies. He is a Research Fellow at The Open University's Knowledge Media Institute.

**CONTENTS**



## 1. INTRODUCTION

### 1.1 Graphical argumentation: notations and tools

Since the early 1980s, many efforts have been made to exploit the properties of graphical notations, supported by hypertext tools, to try and improve the quality of argument construction and the clarity of its communication. The strategy of graphically representing arguments emphasises that asking the right questions and negotiating trade-offs are fundamental processes underpinning the solving of ill-structured, real world problems. Not surprisingly, most interest in the potential of such tools has been in their contribution to complex, intellectual work where the quality of argumentation plays a central role, such as writing (Hashim, 1991; Schuler & Smith, 1990; Streitz, Hanneman & Thüring, 1989), theoretical reasoning (Smolensky, Bell, Fox, King & Lewis, 1987; VanLehn, 1985), legal analysis (Newman & Marshall, 1991) and policy formulation (Conklin & Begeman, 1988; Rittel & Webber, 1973).

As we have documented elsewhere (Buckingham Shum & Hammond, 1994; Buckingham Shum, 1996b), the driving force behind this work derives from roots in the work of Toulmin (1958), Englebart (1963) and Brown (1986). The common thread running through their work is the conviction that the intellectual demands of visualising the structure of an argument can provide insight into its strengths and weaknesses, thus facilitating its more rigorous construction. Furthermore, they argued that exposing an argument's structure in this way facilitates its subsequent communication since important relationships can be more easily perceived and analysed by others.

As will be explained shortly, there is good reason to believe that valuable contributions could be made by such techniques to support the software design process, particularly with respect to requirements engineering and the design of human-computer interaction. Our goal is to provoke more reflective designing at appropriate points, through the introduction of argumentation structure into design. We are construing the problem, therefore, in terms of understanding argumentation structures as cognitive tools which need to be integrated into the different work processes they are intended to support. Thus, we have been investigating how argumentation structure interacts with the individual and group cognitive dynamics of design.

To clarify what is meant by graphical argumentation, let us consider a simple example (Figure 1) which uses a notation loosely based on Toulmin's (op cit.) analysis of argument structure. Suppose one morning, I conclude that *I can't get to work today* because *I got up late* (Step A). Someone might query why getting up late should necessitate missing work (B). My response might be that: *There is only one early bus each day* (C). My protagonist might challenge this, pointing out that *there is a bus at 11am*, and *a train at 10.30am* (D). I might then respond that in fact, *I have an exam today* (E). The debate could obviously unfold further, as this reason was probed. The expressiveness of the notation, and its visual layout rules (which vary with different approaches), determine how large and elaborate an argument can be expressed. Figure 1 serves to illustrate the kind of cumulative argument construction and critiquing which can take place around a shared, graphical argumentation structure (we are not proposing that design teams adopt this particular argumentation scheme).

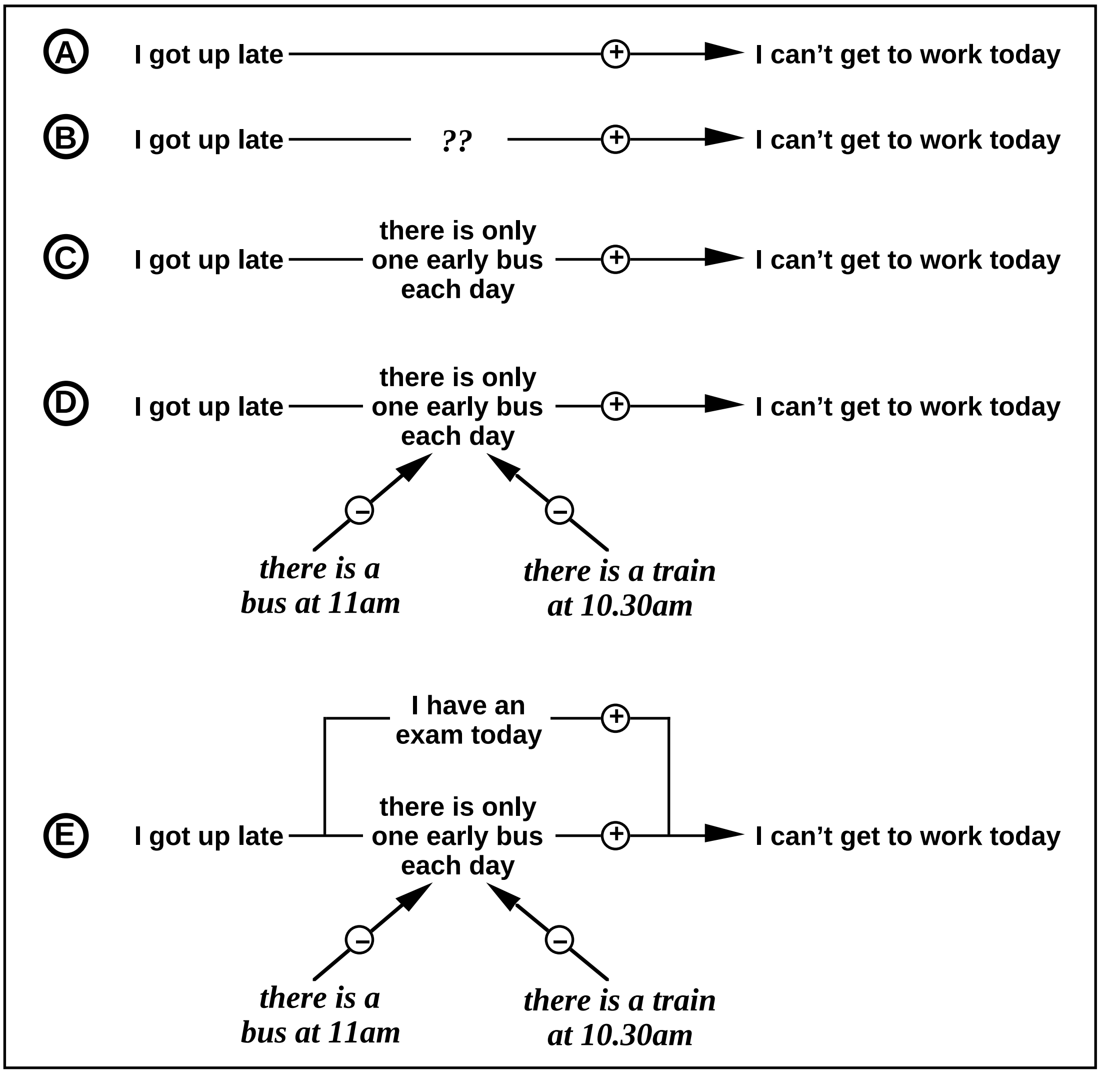


***Figure 1.*** **A graphical argument unfolds (A-E)**

Graphical argumentation on any scale requires software tools to display and manage the networks. Numerous prototype hypermedia argumentation tools have been developed in recent years (Arango, Bruneau, Cloarec & Feroldi, 1991; Fischer, Lemke, McCall & Morch, 1991; Lee, 1990a; Marshall, Halasz, Rogers & Janssen, 1991; Oinas-Kukkonen, 1996; Potts, Takahashi & Anton, 1994; Ramesh, 1993; Rein & Ellis, 1991; Schuler & Smith, 1990; Streitz *et al.*, 1989; Vanwelkenhuysen, 1995). Based on the gIBIS research protoype (Conklin & Begeman, 1988; Conklin & Burgess Yakemovic, 1991), the *QuestMap* collaborative hypermedia system (CMSI, 1993) was developed, a screen from which is shown in Figure 2. Conklin (1996) and Selvin (1997) have reported the use of this system in business contexts.

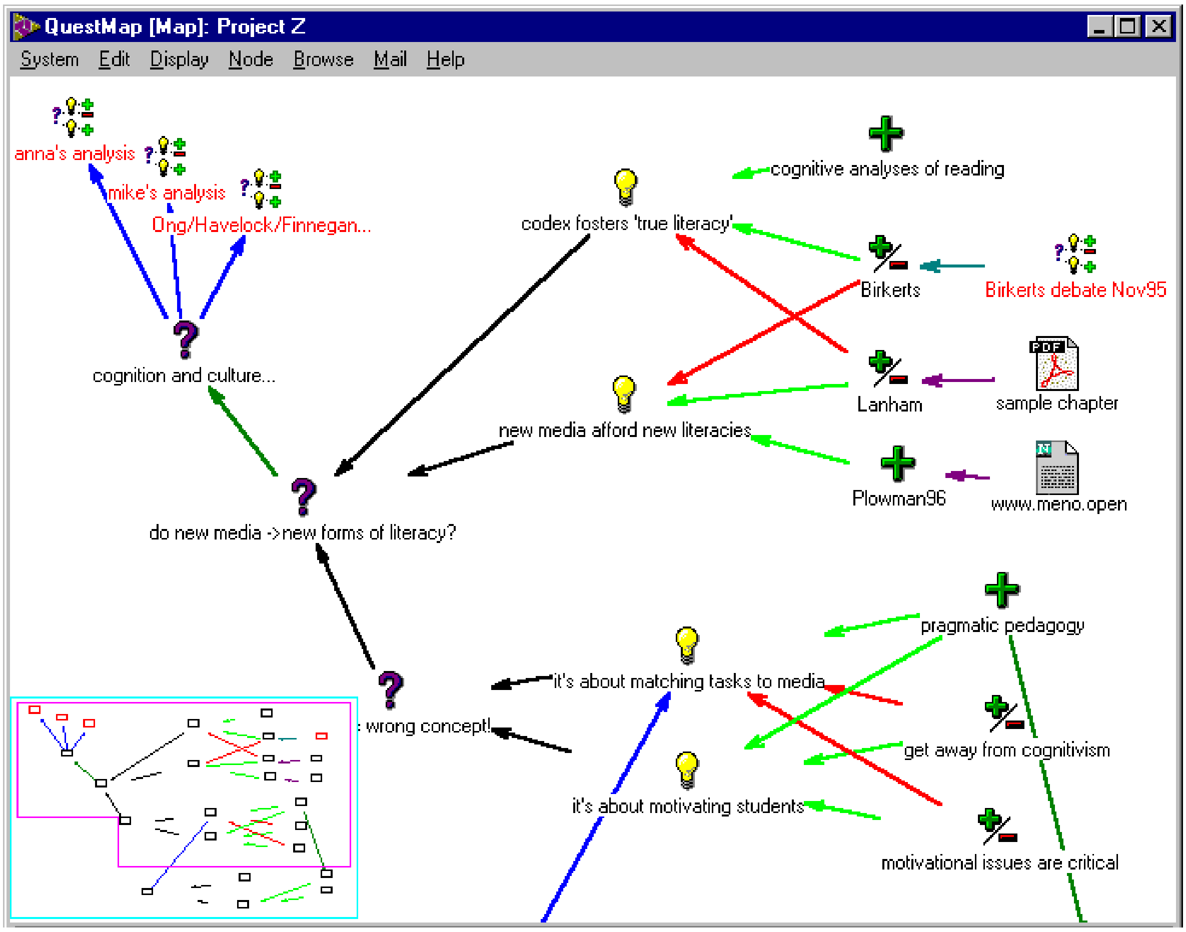


***Figure 2.*** **Illustration of the *QuestMap* argumentation tool as applied to a theoretical analysis of unresolved issues (CMSI, 1993), based on the IBIS argumentative model (see §1.2) and research into the gIBIS prototype (Conklin & Begeman, 1988). *Questions, Ideas,* and supporting or objecting *Arguments* are used to visualise discussions, track unresolved issues, and qualitatively assess the strengths and weaknesses of different positions. Also shown are links to relevant documents, and embedded maps which can encapsulate resolved problems, or contain more detailed analyses as backing for a particular node.**

In terms of scale and ubiquity, the World Wide Web is a significant development in the hypermedia field, and we are now witnessing the emergence of systems which can be regarded as ‘intellectual descendants’ to IBIS. Discussion environments such as HyperNews (LaLiberte, 1995) and Open Meeting (Hurwitz & Mallery, 1995) support structured discussions with tagged contributions such as *agree*, *disagree*, *new idea*. These outline-based schemes are very similar to the “indented-text-IBIS” outline view of Issues, Positions and Arguments used in a large design rationale field study (Burgess Yakemovic & Conklin, 1990). As Web user interfaces approach the quality of interactivity which we now expect from conventional applications, graphical argumentation over the net should become tractable (Gaines & Shaw, 1995; Kremer, 1996). Many applications of freely distributed systems like HyperNews are to support casual discussion similar to that found in newsgroups, in contrast to detailed argumentation. However, their careful adaptation to particular fields such as scholarly peer review for journals (Buckingham Shum, Sumner & Laurillard, 1996) may provide the opportunity for such systems to genuinely augment (i.e. improve, or make more efficient) intellectually intensive work, as originally envisaged.

To summarise, the structure of the argumentation in a debate can be represented in a variety of textual/graphical forms. Structured argumentation, particularly in graphical form, has been proposed as a way to support work requiring complex intellectual reasoning, and numerous hypertext systems have been built to support this. We now consider the specific domain of software design, in which graphical argumentation for *design rationale* has been proposed as a solution to a number of pressing problems in system development, particularly for novel and/or large scale interactive systems.

### 1.2 Design rationale, argumentation, and wicked problems

There has been a convergence of interest in recent years by several design research communities, on the representation of *design rationale* (DR) underlying software design decisions. DR is serving as a focal point for research in Human-Computer Interaction (Carroll & Moran, 1991; MacLean, Young & Moran, 1989; Moran & Carroll, 1996), Software Engineering (Conklin, 1989; Jarczyk, Loffler & Shipman, 1992; Lee, 1991; Potts & Bruns, 1988; Potts *et al.*, 1994; Ramesh, 1993), Knowledge Engineering (Stutt & Motta, 1995; Vanwelkenhuysen, 1995), and Knowledge-based Design Environments (Fischer *et al.*, 1991; Garcia & Howard, 1992). Since DR seeks to provide *rationalisation* for *design*, DR formalisms are also being used to link HCI theory and model-based analyses to design decisions (Bellotti, 1993; Bellotti, Buckingham Shum, MacLean & Hammond, 1995; Carroll & Kellogg, 1989; Carroll, Kellogg & Rosson, 1991).

In the domain of HCI, the extent to which design knowledge and reasoning can be formalised (which for present purposes means made machine-interpretable) is more limited than in more established engineering design domains. This is due in part to the relatively recent emergence of interactive system design with its emphasis on the user interface and understanding human work processes, and due in part to the necessarily craft nature of designing artifacts for human interaction. A design team’s understanding of the work context and of user requirements is invariably incomplete; this requires a design process in which requirements, constraints and design solutions must be regularly re-negotiated, and creates a difficult situation for subsequent designers since the rationale behind decisions may often be unclear, for instance, to someone maintaining the user interface in future versions of a system. Furthermore, keeping track of discussions, decisions and their rationale is made harder in organisations where teams form on a project-specific basis (perhaps ‘out-sourcing’ to external contractors), who proceed to work interdependently but with substantial autonomy, and then disband. Each of these factors may mean that team members do not know each other, or share less common background and culture; such conditions do not facilitate communication, and can make the later recovery and appropriate re-use of that reasoning hard or impossible.

The recruitment of argumentation schemes continues the core interest in supporting reflective, collaborative problem solving of earlier argumentation research, but draws additional motivation from concerns to find solutions to these particularly pressing problems in system design. In this context, there is evidence that argumentation-based design rationale can assist in tackling problems such as:

- clarifying vague requirements, and tracking the rationale for their inevitable evolution (Potts *et al.*, 1994);
- representing multiple stakeholders’ viewpoints, including that of end-users in participatory design (Sjöberg & Timpka, 1995);

- negotiating trade-offs between multidisciplinary analyses, such as software and user criteria (Bellotti *et al.*, 1995);
- maintaining consistency in decision-making, e.g. through propagating changes through networks (Lee, 1990);
- communicating rationale to other designers (McKerlie & MacLean, 1994);
- building cumulative design knowledge, through systematic re-use of rationale (Carroll & Rosson, 1991).

The use of argumentation to support design rationale assumes from the start a ‘dialectic, collaborative’ model for design knowledge, that is, the knowledge invested in a particular project is the product of more than one individual, and often beyond any individual’s grasp. According to this view, the processes of articulating and reconciling different perspectives are central to design, and should be recognised and supported.

This conception of design is supported by analyses showing the centrality of conflict and negotiation in system development (Goldkühl, 1991; Matthiassen, 1987). Of particular importance is the formative work of Rittel (1972; Rittel & Webber, 1973) which has strongly influenced design rationale research and proved to be key in interpreting the results of the studies reported here (see Conklin & Burgess Yakemovic, 1991 and Fischer *et al.*, 1991, for more detailed accounts of Rittel’s work). Rittel proposed that a particular class of problem should be recognised, which he termed *wicked problems*. Wicked problems possess a number of distinctive properties that elude design methods which assume that the problem is already understood sufficiently for it to be analysed using automatic tools, or top-down methods. Specifically, wicked problems:

- cannot be easily defined so that all stakeholders agree on the problem to solve;
- have no clear stopping rules;
- have better or worse solutions, not right and wrong ones;
- have no objective measure of success;
- require iteration—every trial counts;
- have no given alternative solutions—they must be discovered;
- require complex judgements about the level of abstraction at which to define the problem;
- often have strong moral, political or professional dimensions which cannot be easily formalized.

On this basis, Rittel concluded that wicked problems can only be tackled through what he termed an *argumentative* method, which took seriously (rather than finessing as many design methods and tools do) the inevitable debate, negotiation and conflict that invariably arise. Rittel developed IBIS (Issue-Based Information System), a explicitly discussion-oriented method to encourage team members to articulate all kinds of *Issues*, *Positions* in response to those Issues, and *Arguments* to *support* or *object-to* Positions. When visualised graphically (e.g. Figure 2) a map of the discussion grows as new contributions are added to the network.

To summarise, argumentation-based design rationale is pursuing a two-pronged agenda: to support the *immediate process* of designing by supporting the reflective analysis of wicked problems, out of which evolves a *longer term record* of rationale as a reusable resource. This paper examines how this representational process can be negotiated.

### 1.3 Confronting cognitive overhead

To claim benefits on the basis of a formalism's representational properties is one thing; to demonstrate its practical usability for parties other than the formalism's originators is a more complex matter. Ironically, this is a problem that besets even user-centred design methods and tools (Buckingham Shum, 1995), and an analysis of the graphical argumentation literature confirms also that this aspect has been neglected (Buckingham Shum & Hammond, 1994). Specifically, graphical argumentation must confront the challenge facing all knowledge representation enterprises, namely, the initial 'capture problem'—recording and appropriately structuring DR has associated overhead. Semiformal notations still introduce significant representational overhead, reflected in reports within the hypertext literature of such user problems as 'cognitive overhead' and 'premature commitment to structure' (Conklin, 1987; Conklin & Begeman, 1989; Fischer, 1988; Halasz, 1988; Marshall, 1987; Shum, 1991).

Argumentation schemes are undoubtedly easier to use in an asynchronous mode of working, when one can classify and formulate a response to an argument at one's leisure. The evidence available from the success of systems like HyperNews is that most people can learn to use a simple notational scheme such as *Issues*, *Alternatives*, *Pros* and *Cons* to tag comments in asynchronous discussions. However, our particular interest in this paper is in the design of such notations for *concurrent use* in meetings (either face-to-face or via networks), when the time constraints which govern analysis, coordination and contributions to discussion are much tighter. Concurrent use of argumentation schemes during design has been proposed by most of the design argumentation researchers referenced above; however, there are some reports that the use of such schemes during design can obstruct rather than facilitate meetings (Fischer *et al.*, 1991; Rein & Ellis, 1991). Such reports cannot be ignored in the development of argumentation schemes which are meant to provide cognitive support for reasoning during meetings.

On reflection, reports of cognitive overhead should not be surprising. Closely related to argumentation schemes are graphical 'concept mapping' tools, which are being used to visualise conceptual relationships in teaching (JRST, 1990; Kommers, Jonassen & Mayes, 1992) and knowledge engineering (Gaines & Shaw, 1995; Vanwelkenhuysen, 1995). The basis on which these work is that deeper understanding of a domain comes through the *discipline* of expressing knowledge within a structural framework, working to articulate important distinctions and relationships. Similarly, *argumentation* schemes aim to clarify reasoning by encouraging parties to make explicit important assumptions, distinctions, and relationships as they construct and rationalise ideas. In other words, it should not be surprising if expressing DR as semiformal argumentation takes some effort.

The challenge is to provoke more reflective designing at appropriate points, through the introduction of argumentation structure into individual and group design. We are construing the problem, therefore, as one of understanding argumentation structures as cognitive tools which need to be integrated into the workflow they are intended to support. In short, how do argument structures interact with the individual and group cognitive dynamics of design?

The remainder of the paper is organised as follows. §2 introduces the QOC notation which serves as a representative formalism within the field of graphical argumentation; §3 reports the results of three empirical studies of QOC in use; §4 discusses these results from two analytical perspectives, and §5 summarises our conclusions.

## 2. QOC DESIGN SPACE ANALYSIS

MacLean *et al.* (1989) present an approach to representing DR which uses a graphical argumentation scheme called *QOC* (for *Questions*, *Options* and *Criteria*). QOC was designed to be used as a means of more systematically representing the 'design space' around an artifact, and as such forms part of the more general *Design Space Analysis (DSA)* perspective as proposed by (MacLean, Young, Bellotti & Moran, 1991).[1] In QOC, *Questions* are used to encapsulate key issues which shape the design, *Options* are alternative answers to Questions, and *Criteria* are appealed to in assessing one Option over another. In addition, *Assessments* are the relationships between Options and Criteria (at their simplest, *supports* or *objects-to*). These elements are summarised in Figure 3. Boxed Options indicate a decision, or at least a working commitment. The similarity between QOC and the argumentation schemes already described will be immediately apparent.

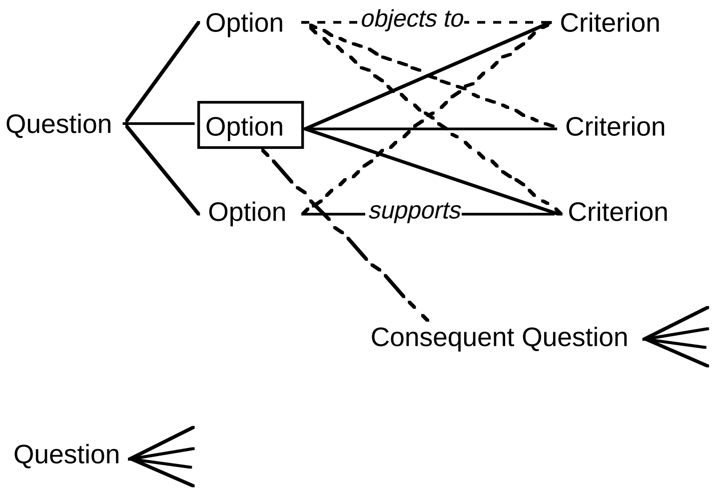


***Figure 3.* The vocabulary of the QOC notation, used to represent Design Space Analyses (MacLean *et al.*, 1991).**

The *design space* is the web of alternative Options for solving the Questions, and the trade-offs which are negotiated by choosing one Option over another. The design space can never be represented in its entirety, since one can continue to ask Questions *ad infinitum*, and from numerous perspectives (e.g. software engineering vs. human factors). DSA is the *process* of discovering the key Questions, exploring the local spaces of Options around these Questions, and justifying why one point in a local space is better than another, through Criteria and their assessments of Options. Extracts from two QOC argumentation structures are shown in Figure 4 to convey the kind of representation with which one works, and the Appendices provide further examples from Study C.

---

1 Except where we are discussing specific aspects of Design Space Analysis (DSA), for simplicity, we shall use the term ‘QOC’ to mean ‘QOC used within the DSA perspective’.

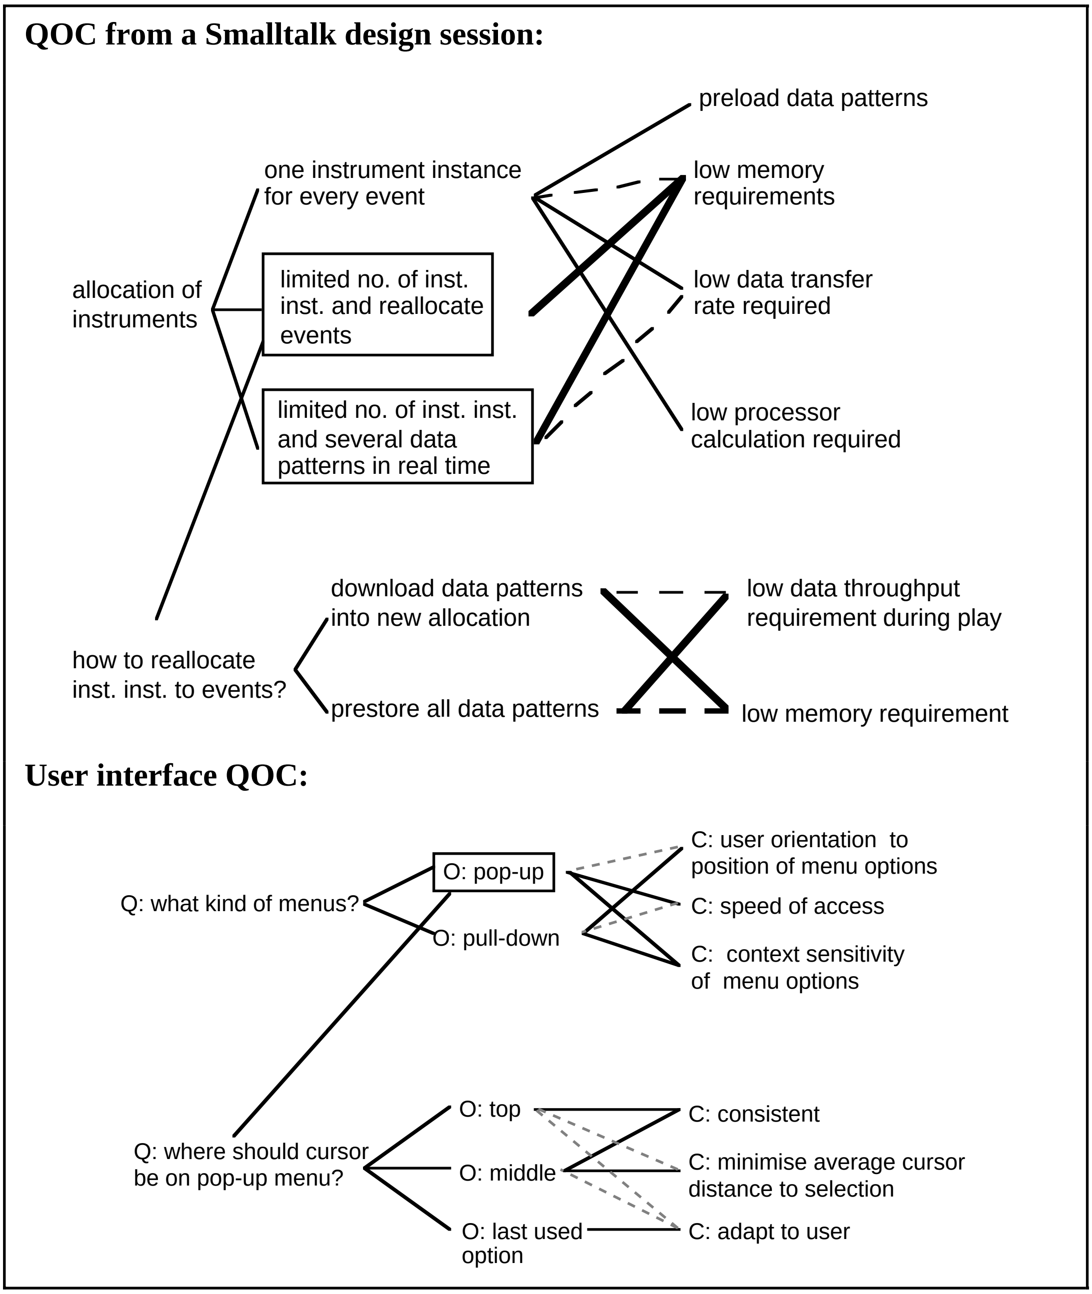


***Figure 4.* Two examples of QOC representing DR for very different kinds of design (note the informal weighting of Assessment links in the top example; important Criteria can also be emphasised if desired).**

## 3. QOC IN USE: THREE EMPIRICAL STUDIES

This section summarises the key findings of a series of empirical studies of QOC in use by designers. Study A was concerned to elucidate the meaning of 'cognitive overhead' more precisely. By studying pairs of designers learning QOC, it identified the core cognitive tasks involved. Study B was longitudinal case study of a single designer; its results began to clarify when and why QOC can be supportive. Studies A and B have been reported in more detail elsewhere (Buckingham Shum, 1996a), to which the interested reader is referred. Study C focused explicitly on group use of QOC, with different designers and design problems. Study C has not been reported before, and consequently is presented in more detail.

### 3.1 Data collection and analysis

All of the data reported in this paper are drawn from video-based observational analyses of designers using QOC whilst solving design problems. Details of the design problems are given under each study. Source transcripts can be obtained from the first author. In all of the studies, designers used pens and large sheets of paper as opposed to a software tool. Under these conditions, the authoring process could be studied with minimal interference from extraneous factors, whilst preserving or even enhancing properties of the representational medium such as display space, resolution, and ease of local editing. Whilst software tools are necessary for long term maintenance of large DRs, the core tasks of deciding how to express reasoning as structured argumentation remain essentially unchanged. Finally, in each study, the designers were debriefed, and given the opportunity to provide any feedback they wished on QOC's value in the session, and more broadly in relation to their work.

### 3.2 Study A: Core cognitive tasks in using QOC

Study A was the first in a series of investigations into QOC as a cognitive tool. Although we hoped that designers would find QOC useful, they were not using it for real design tasks taken from their own work. Twelve pairs of designers tackled the user interface design for an Automated Teller Machine (based on the problem used by (MacLean *et al.*, 1991)). Thus, the primary focus was on the immediate cognitive overhead that must be managed when using QOC after approximately an hour's training. Understanding what a designer is likely to experience on initial exposure to QOC allows us to improve training, in order to improve the chances of subsequent uptake.

Study A was judged successful in that it provided a detailed account of the "nuts and bolts" of using QOC. The study showed how designers must learn to manage four interleaving cognitive tasks in order to express ideas as QOC:

- **Unbundling[2]:** identifying and separating constituent elements of ideas which have been 'bundled together' when they were initially expressed, but which from an argumentation perspective need to be teased apart. A typical example would be someone who raises a problem and proposes a solution plus a supporting argument. The first step to avoiding fruitless argument is to recognise that there are several new elements being introduced, each of which can be responded to in different ways. For instance, much time can be wasted in meetings if a disagreement with one element in an argument is taken to be a dismissal of the whole argument. Teasing out discussion elements and structure pre-empts this kind of confusion.

- **Classification:** deciding whether a contribution is a Question, Option or Criterion. This is not always as simple as it sounds. A Question may ask whether to pursue a particular Option, rather than raising a broader issue to which that Option is but one candidate solution. Criteria may be proposed as solutions, if the Question asked about requirements, and so forth. None of these forms of discourse are 'wrong' in any sense, but they complicate the QOC representation process. The payoff from investing the effort to re-frame discussions is a succinct set of design spaces that support further reflection.

- **Naming:** labelling the new contribution succinctly but meaningfully. It can often be difficult to summarise an idea succinctly in a few words. The skill of doing so is nurtured over time, and the discipline involved can be helpful, although it can also be intrusive in a brainstorming mode of working. The overhead which naming creates is

[2] We are grateful to Jeff Conklin for drawing our attention to this task.

also dependent on the anticipated future use of the QOC, for instance, is it for colleagues present in the meeting, for a formal project review with a manager in three month's time, or for another team to whom you are handing over the design? We return to this issue in the discussion.

- **Structuring:** linking in a new element to other ideas. Meta-level representational and rhetorical decisions may arise at this point. For instance, what Question(s) does a new Option address? How does an Option trade-off against existing Criteria? Is this Question similar to another in a different context, or should a new Question be spawned? Has this Criterion already been used elsewhere under a different name?

An important conclusion from this study was that the above tasks can work against the facile capture of early ideas whose status or relationships were not yet clear. This motivated the concept of "rough QOC" as an important representational stage prior to the ability to formulate more rigorous argumentation (see Shum, 1991, for an analysis of cognitive dimensions which need to be taken into account). Buckingham Shum & Hammond (1994) discuss how argumentation tools could provide for rough DR such as this, and support its incremental formalization into more coherent argumentation, as understanding of the problem develops.

In summary, Study A provided a more precise understanding of the 'cognitive demands' of using an argumentation scheme. These four cognitive tasks in fact apply to a wide range of graphical entity-relationship notations.

### 3.3 Study B: Preliminary evidence of QOC's scope of application

Study B was a longitudinal case study, tracking a Smalltalk designer's use of QOC as he developed a digital music composition system. Three sessions were recorded (each lasting for about 1 ¼ hours) over a period of 3 ½ months.

In Session 1, it became clear that the designer had already invested much thought in the problem he had selected, with some ideas quite well developed, but others relatively vague. The main task to which QOC was put was to systematically lay out the space of alternatives (Options) being faced, in order to understand the high level trade-offs. The designer was very positive about QOC's role in this context—QOC assisted in drawing out existing but vague ideas, and clarified relationships between Options and Criteria which had up to this point remained obscured. Details and examples from Session 1 have been reported elsewhere (Buckingham Shum, 1996a; Buckingham Shum & Hammond, 1994).

In Sessions 2 and 3, however, serious difficulties were encountered in using QOC, and no explicit DR was constructed. The designer was focused on developing two Smalltalk data structures (sketching message passing hierarchies), describing his activity essentially as "gradual refinement" of the structures. It proved difficult to articulate useful Questions, Options and Criteria which offered any analytical leverage on the problem; Questions were either too general or specific to a particular design iteration; discrete Options were impossible to identify because the design was regarded as the gradual evolution of one Option over time; finally, Criteria remained useful only when expressed very generally at a global level of application, rather than differentiating alternatives to subproblems within the design space.

In summary, Study B found evidence that the argumentative mode of working required by and promoted through the use of QOC did support the design work brought by the designer to Session 1, but was incompatible with the work which dominated Sessions 2 and 3. An explanation for this is proposed following Study C, which provided corroborating evidence with a different set of designers and problems.

### 3.4 Study C: QOC in group design

*Data collection and analysis*

Study C involved two established teams of designers from a software company, and three new design problems. In addition, the DR training was extended to pre-empt some of the representational difficulties which had been observed in Studies A and B, by making more explicit elements of QOC ‘craft skill’, such as heuristics (MacLean *et al.*, 1991) and the role of rough, informal QOC (Shum, 1991).

After a morning’s training, the design teams each tackled two one-hour design problems. The first problem was provided by ourselves, the QOC tutors, and described a real prototype system that allowed users to share multimedia documents of different kinds. The design task focused on evaluating and improving two alternative user interfaces for finding out who on the system was using which documents (this was called the *People&Objects* problem). The designers were asked to use QOC to evaluate the designs, improve them if possible, and summarise their conclusions. This problem was chosen because it was sufficiently constrained for the one hour available, and provided numerous alternatives and criteria to consider. This was a relatively well-understood problem for the tutors, since QOC analyses had already been conducted (MacLean, Bellotti & Shum, 1993).

The second problem was chosen by the teams from one of their ongoing projects. One team (*NetGroup*) selected a user interface design problem, the problem being to display the global structure of a graphical network, and provide mechanisms for zooming in for more detailed views. The second team (*FileGroup*) were concerned with designing the optimal file format for storing data files in an application they were building. The designers had specified two clear issues to resolve in the exercise, which were to decide on a file format, and to design the file header which carried the information needed to interpret the file.

All data were collected on-site at the software company. Sketches, notes, and QOC were recorded on paper, and each session video recorded. All discussion and QOC was transcribed, recording the interweaving of sketches, notes, and QOC during discussion.

Analysis of the way in which the design teams used QOC for each problem was conducted by comparing their design activities against a process model for developing QOC design spaces which had been presented in the training. Essentially, this model outlined five main phases of activity through which one typically goes in using QOC, with the provision for opportunistic jumping between the phases (Figure 5). As we explain below, this was used as a yardstick to analyse how QOC was used. The day finished with a video-recorded debriefing in which the designers provided feedback on how they had found using QOC for the first time.

| | | |
|---|---|---|
| **Phase 1:** | **Organise available material** | |
| | Get relevant information down; get a feel for the main issues; classify ideas as QOC if their type is obvious | |
| **Phase 2:** | **Structure material into rough QOC** | |
| | Structure and make sense of the information available; find good Questions | |
| **Phase 3:** | **Flesh out design space** | |
| | Use current understanding of design to help generate new ideas; generate new Options and Criteria | |
| **Phase 4:** | **Reformulate design space to tidy it up** | |
| | Refine Questions (possibly leading to restructuring) | |
| **Phase 5:** | **Make design decisions/commitments** | |
| | Evaluate and select Options | |

***Figure 5.*** **QOC process model presented during training of Study C designers, drawing attention to five phases of activity involved in using QOC.**

*Patterns of QOC-use*

In the People&Objects design problem, the designers were able to construct appropriate QOC structures, and followed the process model quite closely. The exception was Phase 4, in which there was very little explicit restructuring of the first-pass QOC representation, and apart from modifications to names, there was very little reformulation before final decision making. Comparison of the QOCs from the two groups showed that they addressed a similar number of issues (Netgroup eight, FileGroup nine), at a similar level of detail, and topic. This is perhaps not surprising, given that this problem was chosen as a vehicle for studying QOC use, and both the designers' and the hypothetical users' tasks were well-defined. However, the fact that they were able to use QOC to conduct an analysis of the problem attests to a common, satisfactory level of proficiency with the notation.

As described above, the two design groups tackled very different kinds of problems in the second design exercise, and it was here that distinctive differences emerged in the use of QOC. *NetGroup* followed the process model closely. As recommended by Phases 1-3, they spent longer than before on recording key issues as unclassified notes, and classified and elaborated several as rough QOC. They elaborated the design spaces around three Questions, the second and third of which analyzed in detail an initial decision to display graphical representations (Appendix A). This refinement process corresponds to Phase 4 development of the QOC. They reviewed the QOC record at the end, and made their final decisions, as the process model suggested (Phase 5).

In contrast, *FileGroup's* pattern of working was characterised by a disconnection between representing QOC, and discussing and resolving the problem; the QOC was almost superfluous. In terms of the process model's activity phases, FileGroup's mode of working reflected lengthy periods of discussion punctuated by bursts of minimal QOC activity. Exploring alternative Options and reformulating Questions (Phases 3 and 4) occurred to a very limited degree (e.g. adding a Criterion to an earlier Question; changing an Option name when its meaning became clearer). All of the Questions asked were predominantly to document decisions they had already made, with the cursory addition of an alternative Option which was then 'rejected' (Appendix B).

*QOC acted as a 'brake' on discussion*

There was evidence from both design teams that they found QOC to be something of a burden to use. In the Study C debriefing, we received the following feedback:

- FileGroup designers felt that QOC had held back their normally fast-flowing design meetings;
- FileGroup's appointed QOC scribe was 'left out' as discussion moved on whilst he was recording the last decision and rationale;
- Both groups reported that they felt they were meant to step through each decision in turn, by enumerating possibilities at each stage which laboured things;
- Both groups reported that there were points when the way forward seemed to be clear, or when a good decision had been made, but the group had moved forward so quickly that they then had to wait for the QOC to be updated.

## 4. DISCUSSION

Our discussion of these results adopts two perspectives, each of which provides an account of when and why QOC can be supportive or obstructive to design. The first examines QOC's cognitive properties in relation to the design problem being tackled, focusing on the mapping between QOC's argumentative mode of working and the actual task demands. The second examines the additional cognitive properties that QOC can take on in a group design context. Moreover, we find that this second analysis leads to insights into the 'lifecycle' of argumentation structures if they are to support both the initial *process* of wicked problem resolution, as well as provide a coherent *product* which can serve as a useful design rationale for others.

### 4.1 QOC's cognitive properties in relation to design problem solving

When is graphical argumentation useful in design problem solving, and when does it distract from the task in hand? This is the key question which the developers of argumentative methods and tools must be able to answer.

Studies B and C provide a consistent response to this question concerning QOC's use. Let us start by considering the designers' own descriptions of how they were working on each problem, and how QOC fitted in. The Study B designer contrasted two different kinds of designing, drawing a distinction between Session 1, when QOC was useful, and Sessions 2+3, when it was not:

[Study B: Session 1 debriefing] *[explanatory comments are shown like this]*

> ...the activity I was doing here was different from what I was doing yesterday *[i.e. in work prior to session 1]* – the actual low level structures. This *[session 1]* is more about big decisions—policies—whereas this *[refers to own notes]* is about implementation. So this *[QOC]* is certainly very useful for strategies.

[Study B: Final debriefing]

> ... It was quite tricky to think in that way then *[i.e. tricky to use QOC in session 1]* but it was useful... but I think I'd find it almost impossible to wrench myself into thinking in that way *[for sessions 2 and 3]*; it would be unnatural almost, 'cos it's not that kind of path that you take when you're doing this sort of thing.

On two occasions he used analogies to describe how he had been working in the latter sessions:

> Have you ever watched someone design a circuit board?
>
> It's more... like painting a picture of something, and you ask why did you put the trees there and not there? There are so many Criteria and they all interrelate.

Consider now the Study C designers in *FileGroup*, describing the People&Objects user interface problem (provided by us, as QOC tutors) in contrast to their own file-format problem:

[Study C: Final debriefing]

D1: In our last problem *[i.e. file-format problem],* it was actually a problem about how you store things which has real physical constraints and measurable results from doing something in a very real sense. So we could sit there and say, well, we could do it in compressed binary or whatever, and then say, we won't do that because it gives you this factor of something, whereas in the People&Objects *[problem]* we were saying, well this would be nicer to look at, more accessible, in some more, some less quantifiable way...

D2: (Our) final problem ... was much more hard and fast and much better defined in some sense...

D1: It wasn't so much selecting options or even selecting criteria, as measuring the criteria; we spent a lot of time deciding that... deciding the real critical trade-offs was actually the whole point.

The Study B designer describes Session 1 as "strategic", in contrast to being "low level" in Sessions 2+3. Similarly, the Study C designers talk of having a "better defined" problem to work on whose focus was "deciding the critical trade-offs." The Study B designer reported mentally testing and revising the data structure, once he had made the 'strategic' decision to pursue that particular Option. This, we also believe, is what the FileGroup designer meant by "measuring the criteria" in the above extract. FileGroup was tackling a problem with more focused goals at the start of the session, and well established, quantitative metrics (Criteria) for evaluating alternatives.

In both of these studies, we propose that the designers were engaged in a form of problem solving which may be broadly characterised as *multiple constraint satisfaction*—the refining of a small part of a design, testing it against different scenarios, overcoming one small flaw after another.[3] Within our QOC Design Space Analysis approach, we understand this process to be design space *evaluation*. The space of relevant Questions, Options and Criteria is relatively well understood, but what remains is the iterative adjustment of the Options and Criteria to obtain the optimal balance of assessments.

Why is such a mode of working apparently incompatible with an argumentative design approach like QOC? We look to Rittel's analysis of problem solving at this point, as introduced earlier. Problems which are amenable to the multiple constraint satisfaction mode of working are essentially what Rittel termed 'tame' problems, in contrast to 'wicked' problems which cannot be solved algorithmically due to their open-ended nature. As a mode of working, multiple constraint satisfaction assumes by definition that the constraints are well defined, and that one can tell when they are satisfied. These are characteristics which Rittel identified as being conspicuously absent in wicked problems, where one must struggle to articulate the problem usefully, and identify key constraints.

In contrast to FileGroup, the NetGroup designers in Study C found themselves working with the craft principles of user interface design, with vague criteria for evaluating the quality of their interface design. As a result, they found themselves faced with a problem bearing many of the hallmarks of wicked problem solving. Consequently, laying out the issues using QOC caused far less disruption than experienced by the other team (FileGroup), because QOC supported the cognitive task at hand. Thus, in contrast to design space *evaluation*, there was far greater emphasis on design space *elaboration.*[4]

In summary, just as Rittel showed that design methods for tame problems are futile for tackling ill-structured, wicked problems, the results of Studies B and C demonstrate the problems which are encountered in the opposite situation, in which designers tried to apply

3 Gruber & Russell (1996) describe experimental design rationale tools which are able to provide simulation environments for well understood domains. Certain classes of rationale can be generated in response to queries by interrogating formal models of the domain and design.

4 We are grateful to Wendy Kellogg for this characterisation of the two modes of design space analysis.

a technique intended for wicked problems to tame problems. Whilst this might have been predictable in theory (being the converse of Rittel's principle), to our knowledge, this is not a problem which has been recognised in the argumentation literature to date. The (unforeseen, and fortuitous) selection by the Study B and FileGroup designers of problems which turned out to be predominantly tame in nature, has provided empirical evidence to alert us to this boundary condition on the applicability of QOC, and, we would argue, similar argumentation schemes. This in turn highlights the need to sensitise designers to the fact that they may find themselves attempting to do the same thing. No design problem is wholly tame or wicked, but changes in 'degree of wickedness' depending on the designers' expertise and what stage they are at in the design process. If designers can develop the 'meta-skill' of being aware of the kind of design in which they are engaged, they are more likely to analyse ideas using graphical argumentation at those points when it can offer greatest analytical leverage.

## 4.2 QOC's cognitive properties in group design

The Study C evidence from the videos and final debriefing is that managing concurrent representational work during group work is a significant factor. Since this was the designers' first exposure to structured argumentation, like any new language it would have been a bit cumbersome. Indeed, there is other evidence that reactions such as these are common on initial exposure to graphical argumentation schemes (Rein & Ellis, 1991). However, we believe that an important issue surfaced in the group design context of Study C that had not surfaced previously.

The previous section focuses on the cognitive mapping between a single designer, the QOC notation, and the task at hand; it does not inform us about QOC's collaborative affordances. However, software design is frequently a group activity; meetings punctuate, indeed often regulate, the development process in almost all design organisations. A notation assumes new properties in a group context, since it may now serve as a *shared* language and focus of attention. In order to better understand QOC's effect on group dynamics, it is necessary to adopt a different perspective. The vocabulary of discussion shifts now to concerns such as role distribution, and the importance of QOC structures as representations collaboratively developed and owned by the group.

In the following analysis, we draw on a framework for describing the shared workspace activity of small groups engaged in conceptual design. This makes more explicit the roles and commitment which are needed in a group to facilitate QOC's collaborative use. This is then used to analyse from a different perspective the disruptive effect which QOC had on the FileGroup design team.

*Shared representations: mediation, storage and expression*

Tang (1991) has proposed a framework for describing the shared workspace activity of small groups engaged in conceptual design. Based on studies of design teams using whiteboards and paper, he developed a matrix analysis of their activities in terms of designers' *actions* (textual listing, drawing, and gesturing), and the *functions* that these serve (storing information, expressing ideas, or mediating interaction). This characterisation of group design has since served as the conceptual basis for several prototype shared drawing tools (e.g. Greenberg, Roseman, Webster & Bohnet, 1992; Ishii, 1990; Lu & Mantei, 1991; Tang & Minneman, 1990; Tang & Minneman, 1991).

An important distinction is that drawn between recording notes and sketches for *information storage*, typically after explicit agreement, as opposed to the *development and expression of ideas*. The distinction is between documenting 'complete' ideas for later recall, and expressing 'incomplete' ideas to enable others to react to and build on them. Clearly, ideas recorded for later recall may at any point become a vehicle for developing further ideas (and Tang presents instances of interaction which bridge categories). The third function of workspace actions in the framework is to *mediate* interaction. Tang illustrates how all three action types are used by team members to coordinate discussion through the processes of turn-taking and directing the group's attention to different issues. Finally, all of this representational activity is set against a background of *talking* .

The introduction of QOC argumentation into group design adds a new dimension to our conception of group design. Firstly, given its inherently conversational nature, an argumentation scheme will inevitably impact the discussion that forms the background to Tang's three *Actions.* However, argumentation is also an explicit, representational activity, and as such can be introduced as a new kind of *Action*, serving a number of possible *Functions* (Figure 6).

| ***Action*** / ***Function*** | **Collaborative Argumentation** |
|---|---|
| **Store information** | QOC serving as a long term, reusable project memory |
| **Express ideas** | QOC supporting the emergence and debate of new ideas |
| **Mediate interaction** | QOC structuring and focusing discussion |

***Figure 6:*.Adding a *collaborative argumentation* column to Tang's (1991) framework, to clarify QOC's roles in group design.**

It is QOC's role in each of these functions—*mediating interaction, expressing ideas, and storing information*—that we now examine.

*QOC's mediation of group design*

Argumentation schemes are designed specifically to mediate group interaction. An essential part of the Procrustean discipline of using a constrained vocabulary is to encourage a cognitive dialogue with the representation, so that it can 'talk back' to the designer and expose weaknesses in thinking. QOC has certain representational affordances that draw participants' attention to particular aspects of the process. For instance, a Question with no Options, or Options with no assessing Criteria, demand attention at some point. More subtle patterns such as a Decision which has no objecting Criteria will suggest to more experienced users that they have not thought deeply enough about an issue, since there are rarely Options with no negative trade-offs.

Maintaining group momentum whilst recording QOC is another relevant issue. Faced with the task of documenting ideas, the aim is to record the information in the most efficient, timely manner possible. However, Tang observed that the delay this introduces creates a problem for a group striving to maintain momentum in a meeting, and suggested that a group can manage the delay in three possible ways:

- wait for the scribe to finish recording ideas;
- occupy the pause with individual work;
- move on to discuss another issue.

When we consider QOC use in Study C, this problem did indeed arise. The designers used the third strategy, and reported that consequently the QOC scribe was left out of discussion and had to catch up. In general, the second strategy did not occur, since a single scribe was appointed for recording QOC, and hence controlled the representational workspace. It is also worth noting that whilst it was common for designers to watch whilst QOC was recorded (the first strategy), the process of deciding how to summarise what had just been talked about often provoked further contributions, such as suggesting a missing Criterion, or a better name for an Option.

This is the kind of shared 'ownership' of the QOC which can smooth the potential disruption caused by making rationale explicit in a meeting—the extra effort required to create the QOC is managed by making it the group's responsibility, not just that of a dedicated scribe. The frequency with which QOC is recorded will be another factor determining loss of momentum. One group might prefer to record everything almost as it arises, whilst another might periodically reflect on what the most important ideas over the last hour, week or month have been. Naturally these two uses are not mutually exclusive, and will also depend on the kind of problem being tackled. The following sections examine in more detail the continuum between using QOC to proactively *express* ideas, in contrast to a more passive role in which it is used simply to *document* ideas retrospectively.

### *Idea expression using QOC*

Tang contrasts the processes of *storing* ideas with that of *expressing* ideas, as follows:

> [T]he goal of storing information is an *artifact* that records information for later recall. The *process* of creating that artifact is often troublesome, due to the time delay involved. However, when expressing ideas, the goal is to enlist the interaction of the group to develop ideas. Having the group experience and participate in the *process* of creating workspace artifacts is an integral part of expressing and developing ideas. (Tang, 1989, p.86, original emphasis)

Goel & Pirolli (1989) have also documented how representations such as sketches support the process of developing and refining design ideas, as distinct from the subsequent encoding of the products of that process:

> Within a single symbol system, he *[the designer]* constructs multiple representations of the artifact. In both cases, we want to note that these external representations are not for communicating something after the fact. They serve an indispensable role in the generation, evaluation, and decision-making process. Once decisions are made, symbol systems serve to record and perpetuate them. (Goel & Pirolli, 1989, p. 32)

The above quotes emphasize the importance of the *processes* involved in creating that record. Tang noted that artifacts used or created to aid the expression of ideas (such as single words, doodles, emphases on sketches or words) are often meaningless without knowing the context in which they were embedded, which would include gestures, preceding and ongoing discussion, and shared background knowledge. In other words, partially formed notational expressions of this sort will be unavoidable when a notational system plays an active role in synchronous group meetings.

Graphical argumentation is no exception. Quickly created, rough QOC is often unintelligible to outsiders, being dependent on those present at its creation for correct interpretation. When used for expressing and developing ideas, names are often impoverished, sufficient only for the group to understand; similarly, problem decomposition may be partial, again, useful only to the immediate design group as they explore the problem. There may be invisible links between elements which are 'seen' only by the group, or hidden significance in particular spatial clusters. Although these phenomena break the 'official' notational grammer of QOC, these are not 'wrong' uses of QOC which we should discourage designers from adopting. On the contrary, in QOC training courses we explicitly legitimize such exploratory representations in order to make designers aware that they should not expect to construct clean, coherent argumentation first go.

Clearly, there is a representational tension which must be negotiated. Argumentation schemes provide most analytical leverage when the effort is made to engage in the four cognitive tasks identified in Study A—unbundling, classifying, naming and structuring. It is these which distinguish graphical argumentation from other less structured forms of notetaking or discussion. There is a balance to find, therefore, between recognising the importance of rough, ill-formed QOC, and the value of taking the extra effort to create cleaner, more rigorous QOC. Ultimately, this is a practical skill which designers must develop for themselves in accordance with their individual and collective working styles. It will also depend on the organisational demands that will be made on the QOC record, as discussed next.

*Evolution from expressive to documentary QOC*

As a rule, the lifecycle of a design rationale cannot end at the stage of rough QOC. If it is to serve as a longer term project memory, or is to make sense to outsiders who cannot fill in missing details known only by the original group (who of course will themselves forget details as time passes), a QOC structure will need to be edited and integrated with other representations. We find that Tang's framework provides insight not only into the detailed dynamics of shared workspace activity, but that its concepts offer insight into the longer term 'macro-evolution' of design argumentation structures. This extension of Tang's framework is not required to account for the results reported in this paper (Study B, the one longitudinal case study reported, involved only a single designer). However, other QOC case studies cited below provide preliminary evidence of QOC's role as a longer-term, shared memory resource, and confirm the analysis that follows.

The distinction between *expression* and *storage* within Tang's framework describes very neatly the evolution from a representation for *unstable, provisional rationale*, to one for more *stable, consensus rationale*. This corresponds to shifts from implicit to explicit knowledge, from a private to a public resource, and from a one-off temporary representation to facilitate a single meeting or project, to being a reusable resource of wider interest. This evolution can be accomplished through two processes which we shall call *enrichment* and *conversion*.

Through *enrichment*, the expressive QOC is retained as the basic structure for a more complete, longer term storage QOC. For example, important assumptions may need to be declared, acronyms and abbreviated names of nodes expanded, and implicit spatial organization of the QOC made more explicit.[5] As new Questions, Options and Criteria are added, they are often concretised in the form of sketches and other design renderings; however, rather than leave these locked in the memories of the designers or fragmented in various places and media, they can be linked into the QOC. The design rationale itself will need to be linked to other relevant design documents to illustrate points or provide important background context (Shum, MacLean, Forder & Hammond, 1993; MacLean & McKerlie, 1995). Through this process, rough QOC, a team's short term *working memory*, is enriched to the point where it can provide a *long term memory* resource.

Enriching the existing structure means that designers can work within a single notational (probably hypermedia) environment, using the graphical browser as the primary working representation. No conversion to, and associated editing of, other formats is required. The main requirement on any user is that they are familiar with the QOC notation. Our experience in training designers (MacLean, Buckingham Shum & Bellotti, 1992-94) is that assuming one is familiar with the general problem domain, QOC is not a hard notation to grasp or read (but more skill is needed to construct it, as described in this paper). Currently, there is documented evidence that project members not involved in the construction of enriched QOC, but familiar with the notation and the problem, can understand and build on it (McKerlie & MacLean, 1994). Elsewhere, we have proposed some of the functionality

5 Work on providing computational support for detecting and formalizing meaningful spatial arrangements of ideas is described by Marshall & Shipman (1995).

which might be expected in an environment to support the enrichment of argumentation (Buckingham Shum & Hammond, 1994).

In the process of *QOC conversion*, the key ideas and arguments currently in expressive, graphical QOC form are distilled into a different form for storage. The most obvious example is to write a conventional textual report; this might inherit its structure from the QOC network, and could include relevant fragments of QOC argumentation.

Conversion obviously introduces new effort to create the new format document. However, developments in hypertext writing tools suggest that the basic structure could be generated automatically from the QOC network. One danger of conversion is that the new document becomes ‘frozen’ (unable to respond to new developments), rather than ‘living’ as a resource which can continue to grow as new discussion is added (one of the strengths of network-based notations). Nonetheless, conventional documents will be needed in organisations for the foreseeable future, and digital documents have the ability to dynamically update representations from one application embedded in another. The goal is to preserve the links to the argumentation which gave rise to the document, just as the rationale for other designed artifacts should be preserved and made easily accessible.

## 5. SUMMARY AND CONCLUSION

For many years, researchers have been striving to understand how computers can augment our ability to solve problems. One of the most obvious ways is to seek to augment the *debates* about problems—both in one’s own mind and with other minds—which underpin much intellectual work. The strategy of graphically representing the structure of, and links between, arguments is based on extensive evidence that articulating incisive questions, holding in tension conflicting perspectives, and negotiating trade-offs are fundamental processes underpinning the resolution of ill-structured, wicked problems. However, structuring reasoning and discussion comes at a cost, namely, it requires the integration of a new representational structure into the cognitive and discursive flow. The question which we have sought to address is whether this provides supportive or obstructive structure.

This paper has presented a use-oriented analysis of the QOC argumentation scheme in the context of design rationale capture. It has described the individual and group demands of using a semiformal argumentation scheme to support design reasoning, seeking to identify when, and clarify why, QOC supports particular kinds of design problem solving and obstructs others. We have argued that QOC provided most support when designers needed to elaborate poorly understood design spaces in order to clarify the key Questions, Options and Criteria. However, it was a distraction when they were trying to evaluate relatively well understood design spaces in which the main task was to satisfy multiple constraints through iterative testing and adjustment, until an optimal solution has been obtained. We have argued that these findings can be understood in terms of cognitive compatibility between the argumentative mode of working required and encouraged by QOC, and the different modes of work required by the problems.

We then drew on a framework for understanding the role of shared representations in meetings, to clarify factors particular to group use of QOC. Teams need to actively manage QOC recording during their work, and in this context QOC can play a proactive or passive role in mediating deliberation. This interacts with the role which it is expected to play in the future as a form of longer term project memory; since expressive QOC is often intelligible only to the group whose ideas it mediates, it will need to evolve into more publically intelligible documentary QOC. It is hard to see how this representational tension can be avoided, since the demands on a formalism to satisfy both private and public roles are mutually incompatible. What is needed are technologies that support the transition from expression to documentation, that is, from informal, incomplete, private rationale to more formal, complete and publicly intelligible rationale.

The studies reported here are grounded in the design domain. However, ill-structured, wicked problems are recognisable in many real world domains, all of which will involve the interplay between elaborating and evaluating solution spaces. On this basis, our analysis should inform efforts to provide representational support for teams tackling such problems, both in their argumentative process, and in constructing an intelligible group memory of this process as a resource for the future.

***Acknowledgements.*** We are indebted to Wendy Kellogg and Jack Carroll whose particularly detailed reviews helped to sharpen our thinking. Also to Jeff Conklin, Jean McKendree, Tom Moran, Tamara Sumner, John Tang and an anonymous reviewer for feedback on earlier versions. For Study C, Ian Daniel provided valuable technical assistance and David Elworthy transcribed video and QOC with great efficiency. We are grateful to the designers who participated from the Impact Project (Nestlé Rowntree, York), Logica Cambridge Ltd., University of York Music Technology Group, and The Warehouse Project, York.

***Support.*** Parts of this research were conducted whilst Simon Buckingham Shum was at the University of York and Rank Xerox Cambridge EuroPARC, and Victoria Bellotti was at Rank Xerox Cambridge EuroPARC. Funding is gratefully acknowledged from Science & Engineering Research Council CASE Award 88504176, Rank Xerox, and ESPRIT Basic Research Actions 3066 and 7040 (AMODEUS Projects 1 & 2).

## REFERENCES

Arango, G., Bruneau, L., Cloarec, J.-F., & Feroldi, A. (1991). A tool shell for tracking design decisions. *IEEE Software*, *March*, 75-83.

Bellotti, V. (1993). Integrating theoreticians' and practitioners' perspectives with design rationale. *Proceedings of the InterCHI'93 Conference on Human Factors in Computing Systems*, 101-106. New York: ACM.

Bellotti, V., Buckingham Shum, S., MacLean, A., & Hammond, N. (1995). Multidisciplinary modelling in HCI design... In theory and in practice. *Proceedings of CHI'95 Conference on Human Factors in Computing Systems*, 146-153. New York: ACM.

Brown, J. S. (1986). From cognitive ergonomics to social ergonomics and beyond: In D. A. Norman & S. W. Draper (Eds.), *User centered system design* (pp. 457-486). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Buckingham Shum, S. (1995). Practise what we preach: Making HCI design techniques usable. *Proceedings of STIMDI'95 Conference of the Swedish Interdisciplinary Interest Group for Human-Computer Interaction*, 1-13. Uppsala: Dept. Computer Science, Uppsala University, Sweden. [http://www.mrc-apu.cam.ac.uk/amodeus/abstracts/cp/cp89.html].

Buckingham Shum, S. (1996a). Analyzing the usability of a design rationale notation: In T. P. Moran & J. M. Carroll, *Design Rationale: Concepts, Techniques, and Use* (pp. 185-215). Hillsdale, NJ: Lawrence Erlbaum Associates.

Buckingham Shum, S. (1996b). Design argumentation as design rationale. *The Encyclopedia of Computer Science and Technology, 35,* 20, 95-128.

Buckingham Shum, S. & Hammond, N. (1994). Argumentation-based design rationale: What use at what cost? *International Journal of Human-Computer Studies*, *40*, 603-652.

Buckingham Shum, S., Sumner, T., & Laurillard, D. (1996). On the future of journals: Digital publishing & argumentation. *Adjunct Proceedings of HCI'96: 11th British Computer Society Conference on Human-Computer Interaction*, 48-50. Dept. Computer Science, Middlesex University, U.K. [http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-24-abstract.html].

Burgess Yakemovic, K. C. & Conklin, J. (1990). Report on a development project use of an issue-based information system. *Proceedings of CSCW'90 Conference on Computer Supported Cooperative Work*, 105-118. ACM: New York.

Carroll, J. M. & Kellogg, W. A. (1989). Artifacts as theory nexus: Hermeneutics meets theory-based design. *Proceedings of CHI'89 Conference on Human Factors in Computing Systems*, New York: ACM.

Carroll, J. M., Kellogg, W. A., & Rosson, M. (1991). The task-artifact cycle: In J. M. Carroll (Ed.), *Designing interaction: psychology at the human-computer interface*. New York: Cambridge University Press.

Carroll, J. M. & Moran, T. P. (1991). Introduction to this special issue on design rationale. *Human-Computer Interaction Journal*, 6, 197-200.

Carroll, J. M. & Rosson, M. B. (1991). Deliberated evolution: Stalking the view matcher in design space. *Human-Computer Interaction*, 6, 281-318. Also in T.P. Moran and J.M. Carroll (Eds.) *Design rationale: Concepts, techniques, and use*, (pp. 107-145). Hillsdale, NJ: Lawrence Erlbaum Associates, 1996.

CMSI (1993). *QuestMap* (software). Austin, TX: Corporate Memory Systems, Inc. [http://www.cmsi.com/info/].

Conklin, J. (1987). Hypertext: An introduction and survey. *IEEE Computer*, 20, 9, 17-41.

Conklin, J. (1989). Design rationale and maintainability. *Proceedings 22nd Hawaii International Conference on System Sciences* (Vol. 2, pp. 533-539). Los Alamitos: IEEE Computer Society Press.

Conklin, J. (1996). *Designing organizational memory: Preserving intellectual assets in a knowledge economy*. Corporate Memory Systems, Inc. [http://www.cmsi.com/info/pubs/desom/].

Conklin, J. & Begeman, M. L. (1988). gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Office Information Systems*, 6, 303-331.

Conklin, J. & Begeman, M. L. (1989). gIBIS: A tool for all reasons. *Journal of the American Society for Information Science*, *40*, 200-213.

Conklin, J. & Burgess Yakemovic, K. C. (1991). A process-oriented approach to design rationale. *Human-Computer Interaction*, 6, 357-391. Also in T.P. Moran and J.M. Carroll (Eds.) *Design rationale: Concepts, techniques, and use,* (pp. 393-427). Hillsdale, NJ: Lawrence Erlbaum Associates, 1996.

Englebart, D. C. (1963). A conceptual framework for the augmentation of man's intellect: In P. Howerton & Weeks, *Vistas in information handling* (pp. 1-29). London: Spartan Books: Washington, DC Also in Greif, I., Ed. (1987). Computer-Supported Cooperative Work: A Book of Readings, pp. 35-65. Morgan Kaufmann Publishers: San Mateo, California.

Fischer, G. (1988). Panel session position statement: A critical assessment of hypertext systems. *Proceedings of CHI'88 Conference on Human Factors in Computing Systems*, 223-227. ACM: New York.

Fischer, G., Lemke, A. C., McCall, R., & Morch, A. I. (1991). Making argumentation serve design. *Human-Computer Interaction*, 6, 393-419. Also in T.P. Moran and J.M. Carroll (Eds.) *Design rationale: Concepts, techniques, and use,* (pp. 267-293). Hillsdale, NJ: Lawrence Erlbaum Associates, 1996.

Gaines, B. R. & Shaw, M. L. S. (1995). Concept maps as hypermedia components. *International Journal of Human-Computer Studies*, *43*, 323-361.

Garcia, A. C. B. & Howard, H. C. (1992). Acquiring design knowledge through design decision justification. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, *6*, 1, 59-71.

Goel, V. & Pirolli, P. (1989). Motivating the notion of generic design within information-processing theory: The design problem space. *AI Magazine*, *Spring*, 18-36.

Goldkühl, G. (1991). Information systems design as argumentation: An investigation into design rationale as a conceptualization of design. *Proceedings of 14th IRIS Conference on Information System Research In Scandinavia.* Umeå University, Sweden.

Greenberg, S., Roseman, M., Webster, D. & Bohnet, R. (1992). Human and technical factors of distributed group drawing tools. *Interacting with Computers*, *4*, 364-392.

Gruber, T. R. & Russell, D. M. (1996). Generative rationale: Beyond the recall and replay paradigm. In T. P. Moran & J. M. Carroll, *Design rationale: Concepts, techniques, and use* (pp. 323-349). Hillsdale, NJ: Lawrence Erlbaum Associates.

Halasz, F. G. (1988). Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, *31*, 836-852.

Hashim, S. H. (1991). WHAT: An argumentative groupware approach for organizing and documenting research activities. *Journal of Organizational Computing*, *1*, 275-302.

Hurwitz, R. & Mallery, J. C. (1995). The Open Meeting: A web-based system for conferencing and collaboration. *Proceedings of the Fourth International Conference on The World-Wide Web*, Boston: World Wide Web Consortium [http://www.w3.org/pub/Conferences/WWW4/Papers2/349].

Ishii, H. (1990). TeamWorkStation: Towards a seamless shared workspace. *Proceedings of CSCW'90 Conference on Computer-Supported Cooperative Work*, 13-26. New York: ACM.

Jarczyk, A., Loffler, P., & Shipman, F. (1992). Design rationale for software engineering: a survey. *Proceedings of 25th Hawaii International Conference on System Sciences.* Los Alamitos: IEEE Computer Society Press

JRST (1990). *Journal of Research and Science Teaching* (Special Issue on Concept Mapping), *27*, 10.

Kommers, P. A. M., Jonassen, D. H., & Mayes, J. T. (Eds.). (1992). *Mindtools: cognitive technologies for modelling knowledge*. Berlin: Springer-Verlag.

Kremer, R. (1996). Toward a multi-user, programmable web concept mapping “shell” to handle multiple formalisms. *Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Calgary: Dept. Computer Science, U. Calgary, Canada [ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/77kremer.ps.Z].

LaLiberte D. (1995). *HyperNews* (software). Urbana-Champaign, IL: National Center for Supercomputing Applications. [http://union.ncsa.uiuc.edu/HyperNews/get/hypernews.html].

Lee, J. (1990a). SIBYL: A qualitative decision management system: In P. Winston & S. Shellard (Eds.), *Artificial intelligence at MIT: Expanding frontiers* (pp. 105-133). Cambridge, MA: MIT Press.

Lee, J. (1990b). SIBYL: A tool for managing group design rationale. *Proceedings of CSCW’90 COnference on Computer Supported Cooperative Work*, 79-92. New York: ACM.

Lee, J. (1991). Extending the Potts and Bruns model for recording design rationale. *Proceedings of the 13th International Conference on Software Engineering*, 114-125. New York: ACM.

Lu, I. & Mantei, M. (1991). Idea management in a shared drawing tool. *Proceedings of the 2nd European Confernce on Computer Supported Cooperative Work*,, 97-112.

MacLean, A., Bellotti, V., & Shum, S. (1993). Developing the design space with design space analysis: In P. F. Byerley, P. J. Barnard, & J. May (Eds.), *Computers, communication and usability: Design issues, research and methods for integrated services* (pp. 197-219). Amsterdam: Elsevier North Holland Series in Telecommunication.

MacLean, A., Buckingham Shum, S., & Bellotti, V. (1992-94). Design rationale tutorial. Tutorials presented at *HCI’92-94: British HCI Group Annual Conferences, U.K.* Available from first author.

MacLean, A. & McKerlie, D. (1995). Design Space Analysis and use-representations: In J. M. Carroll (Ed.), *Scenario-based design: Envisioning work and technology in system development* (pp. 183-207). New York: Wiley.

MacLean, A., Young, R. M., Bellotti, V., & Moran, T. (1991). Questions, options, and criteria: Elements of Design Space Analysis. *Human-Computer Interaction*, *6*, 201-250. Also in T.P. Moran and J.M. Carroll (Eds.) *Design rationale: Concepts, techniques, and use,* (pp. 53-105). Hillsdale, NJ: Lawrence Erlbaum Associates, 1996.

MacLean, A., Young, R. M., & Moran, T. (1989). Design rationale: The argument behind the artifact. *Proceedings of the CHI’89 Conference on Human Factors in Computing Systems*, 247-252. ACM: New York.

Marshall, C. C. (1987). Exploring representation problems using hypertext. *Proceedings of the Hypertext’87 Conference*, 253-268. ACM: New York.

Marshall, C. C., Halasz, F. G., Rogers, R. A., & Janssen, W. C. (1991). Aquanet: A hypertext tool to hold your knowledge in place. *Proceedings of Hypertext’91 Conference.* ACM: New York.

Marshall, C. C. & Shipman, F. M. (1995). Spatial hypertext: Designing for change. *Communications of the ACM*, *38*, 8, 88-97 [http://irss.njit.edu:5080/cgi-bin/bin/option.csh?mar-ship/marshall.html].

Matthiassen, L. (1987). Systems, processes and structures: A contribution to the theoretical foundation of information systems development and use: In P. Docherty, K. Fuchs-Kittowski, P. Kolm, & L. Mathiassen (Eds.), *System Design for Human Development and Productivity: Participation and Beyond*. Amsterdam: North Holland.

McKerlie, D. & MacLean, A. (1994). Reasoning with design rationale: Practical experience with design space analysis. *Design Studies*, *15*, 214-226.

Moran, T. P. & Carroll, J. M. (Eds.). (1996). *Design rationale: Concepts, techniques, and use.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Newman, S. & Marshall, C. (1991). *Pushing Toulmin too far: Learning from an argument representation scheme*. (Technical Report SSL 92-45). Palo Alto: CA: Xerox Palo Alto Research Center.

Oinas-Kukkonen, H. (1996). Debate Browser: An argumentation tool for MetaEdit+ environment. *Proceedings of the Seventh European Workshop on the Next Generation of CASE Tools*. [http://rieska.oulu.fi/~hok/ngct96.ps].

Potts, C. & Bruns, G. (1988). Recording the reasons for design decisions. *Proceedings of 10th International Conference on Software Engineering*, 418-427. Washington, DC: IEEE Computer Press

Potts, C., Takahashi, K., & Anton, A. (1994). Inquiry-based requirements analysis. *IEEE Software, March*, 21-32.

Ramesh, B. (1993). Supporting systems development by capturing deliberations during requirements engineering. *IEE Transactions on Software Engineering*, *18*, 498-510.

Rein, G. L. & Ellis, C. A. (1991). rIBIS: A real-time group hypertext system. *International Journal of Man-Machine Studies*, *24*, 349-367 Also in Greenberg S. (Ed.) *Computer supported cooperative work and groupware*, (pp. 223-241). Academic Press: London, 1991.

Rittel, H. W. J. (1972). Second generation design methods. *Design Methods Group 5th Anniversary Report: DMG Occasional Paper*, *1*, 5-10. Also in N. Cross (Ed.) *Developments in design methodology*, (pp. 317-327). Chichester: Wiley & Sons, 1984.

Rittel, H. W. J. & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, *4*, 155-169.

Schuler, W. & Smith, J. (1990). Author's Argumentation Assistant (AAA): A hypertext-based authoring tool for argumentative texts. In A. Rizk, N. Streitz, & J. André (Eds.), *Hypertext: Concepts, systems and applications*. Cambridge: Cambridge University Press.

Selvin, A. (1997). *An issue-based hypertext approach to collaborative process redesign* (Technical Memorandum TM 97-0007). White Plains, NY: NYNEX Science & Technology.

Shum, S. (1991). Cognitive dimensions of design rationale. In D. Diaper & N. V. Hammond (Eds.), *People and computers VI: Proceedings of HCI'91* (pp. 331-344). Cambridge: Cambridge University Press.

Shum, S., MacLean, A., Forder, J., & Hammond, N. V. (1993). Summarising the evolution of design concepts within a design rationale framework. *Adjunct Proceedings of the InterCHI'93 Conference on Human Factors in Computing Systems*, 43-44. New York: ACM [ftp://kmi.open.ac.uk/pub/simonb/interchi93.ps].

Sjöberg, C. & Timpka, T. (1995). Inside multidisciplinary design in medical informatics: Experiences from the use of an argumentative design method. *Proceedings of theMEDINFO'95 Tri-annual World Conference in Medical Informatics*, Vancouver. [http://www.mrc-apu.cam.ac.uk/amodeus/abstracts/cp/cp76.html].

Smolensky, P., Bell, B., Fox, B., King, R., & Lewis, C. (1987). Constraint-based hypertext for argumentation. *Proceedings of the Hypertext'87 Conference,* 215-245. New York: ACM.

Streitz, N., Hanneman, J., & Thüring, M. (1989). From ideas and arguments to hyperdocuments: Travelling through activity spaces. *Proceedings of the Hypertext'89 Conference* , 343-364. New York: ACM.

Stutt, A. & Motta, E. (1995). Recording the design decisions of knowledge engineers to facilitate re-use of design models. *Proceedings 9th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems.* Banff, Calgary: Dept. Computer Science, U. Calgary, Canada.

Tang, J. C. (1991). Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34, 2, 143-160. Also in S. Greenberg (Ed.) *Computer supported cooperative work and groupware,* (pp. 11-28). London: Academic Press, 1991.

Tang, J. C. & Minneman, S. L. (1990). VideoDraw: A video interface for collaborative drawing. *Proceedings of the CHI'90 Conference on Human Factors in Computing Systems*, 313-320. ACM: New York.

Tang, J. C. & Minneman, S. L. (1991). VideoWhiteboard: Video shadows to support remote collaboration. *Proceedings of the CHI'91 Conference on Human Factors in Computing Systems*, 315-322. ACM: New York.

Toulmin, S. (1958). *The uses of argument*. Cambridge: Cambridge University Press.

VanLehn, K. (1985). *Theory reform caused by an argumentation tool* (Technical Report ISL-11). Palo Alto, CA: Xerox Palo Alto Research Center.

Vanwelkenhuysen, J. (1995). Embedding non-functional requirements analyses in conceptual knowledge systems designs. *Proceedings of the 9th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems*. Banff, Calgary: Dept. Computer Science, U. Calgary, Canada.

## APPENDIX A: NETGROUP'S QOC ARGUMENTATION
## (Study C, Network-browser user interface design)

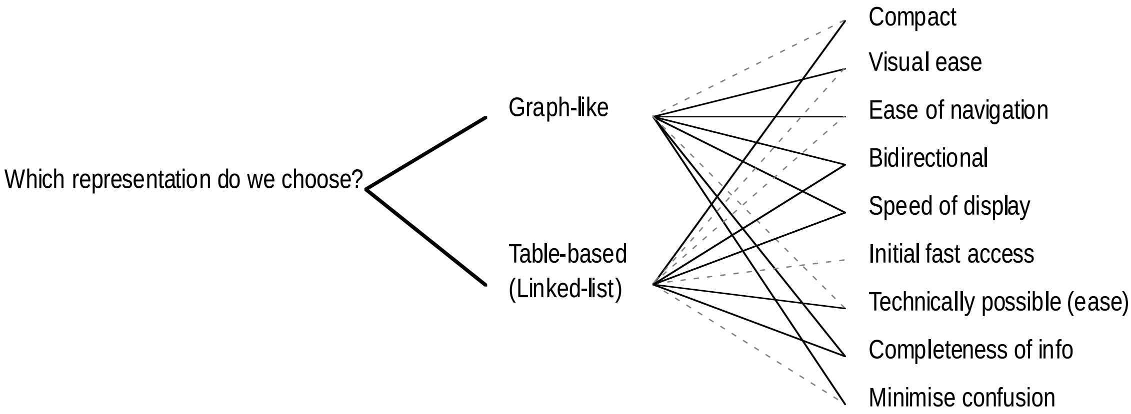


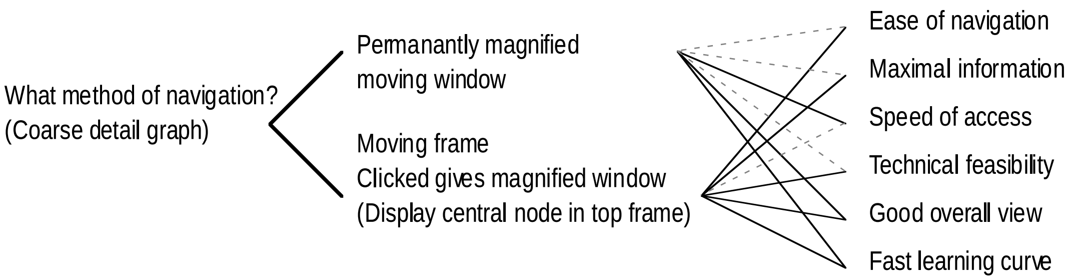


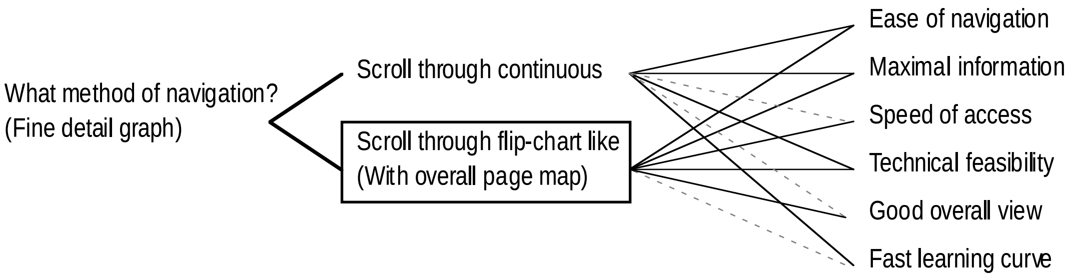

## APPENDIX B: FILEGROUP'S QOC ARGUMENTATION
## (Study C, File-header design)