Journal of Universal Computer Science, vol. 3, no. 8 (1997), 899-928
submitted: 1/6/97, accepted: 10/8/97, appeared: 28/8/97

Negotiating the Construction and Reconstruction of Organisational Memories

Simon Buckingham Shum
(The Open University, U.K.
S.Buckingham.Shum@open.ac.uk)

Abstract: This paper describes an approach to capturing organisational memory, which serves to ground an analysis of human issues that knowledge management (KM) technologies raise. In the approach presented, teams construct graphical webs of the arguments and documents relating to key issues they are facing. This supports collaborative processes which are central to knowledge work, and provides a group memory of this intellectual investment. This approach emphasises the centrality of negotiation in making interdisciplinary decisions in a changing environment. Discussion in the paper focuses on key human dimensions to KM technologies, including the cognitive and group dynamics set up by an approach, the general problem of preserving contextual cues, and the political dimensions to formalising knowledge processes and products. These analyses strongly motivate the adoption of participatory design processes for KM systems.

Key Words: organisational memory, knowledge management, argumentation, participatory design, knowledge-based systems, collaborative systems

Categories: H, H1.2, H5.1, H5.2, H5.3

1. Introduction and Definitions

In order to operationalise the concept of Knowledge Management (KM), numerous disciplines are now trying to analyse the processes and products of organisational knowledge, in order to clarify what tangible representations future knowledge managers might work with. These representations of the domain facilitate viewpoints and analyses of particular information-types from particular perspectives. This paper describes one form of KM technology that has been developed over several years, which throws into relief a spectrum of human issues which are intrinsic to the process of designing and implementing KM representations -- computer-supported or otherwise instantiated. This is particularly germane to the application of artificial intelligence (AI) techniques to KM, currently one of the most strongly represented disciplines in KM research, since the success of such approaches rests heavily on finding appropriate representations for knowledge modelling, ontology design, knowledge-based system building, and the subsequent reasoning that these activities are intended to support.

Let us begin by unpacking the concepts in the title, since several potentially ambiguous terms have been used. Firstly, meaningful *memories* are not simply retrieved according to some database model, but are *reconstructed* in the context of who is asking, and for what purpose. Bannon and Kutti [Bannon 1996] present an excellent introduction to the need to shift from a passive 'storage bin' metaphor for organisational memory, to a more appropriate one of active reconstruction. We say different things to different people, varying the level of detail, emphasis, perspective,

and so forth. Moreover, what is sanctioned as reliable knowledge depends on the community of interested stakeholders, who confer significance on certain sources (e.g. people), whether explicitly or implicitly. Knowledge is in that sense also *constructed*, serving particular needs at a particular time. When attempting to create a shared information or memory resource, we should not be surprised to find that *negotiations* about what is included, how it should be organised, and who has access to it, become key processes. This resource will itself be *constructed* over time as contributions are added to the digital corpus, and as its form and role within the project evolve.

This paper introduces an approach to capturing organisational memory that takes into account the epistemological assumptions and collaborative processes implied by this framing of the problem. Teams use hypermedia groupware to construct graphical webs of argumentation and related documents as they discuss problems, recording aspects of their reasoning for future reference. This is a relatively mature approach which may be familiar to researchers in hypertext, computer-supported collaborative work (CSCW), groupware and software design rationale. The purpose of this paper is to contextualise it to the particular concerns of KM, and to use it to ground discussion of generic issues that KM technologies raise.

The paper starts in Section 2 by characterising the *context* of 'knowledge work' -- if 'knowledge workers' constitute an organisation's expertise, are there salient features of knowledge work that we can recognise? Section 3 introduces graphical argumentation as a candidate approach, with a particular niche in the design space of organisational memory systems. Section 4 introduces its representations for capturing group memory, Section 5 the appropriate supporting technologies, and Section 6 then characterises the kinds of knowledge that can be captured with this combination. Section 7 briefly surveys studies of the approach's application, moving into a discussion in Section 8 of the hands-on practicalities of using it, taking into account cognitive, social and organisational level issues. Particular attention is paid to the problem of capturing adequate context. Section 9 closes the paper by reflecting on the commitments that are made in adopting any representation, and the related issues of control and power that arise in managing knowledge about, and for, staff in an organisation.

2. Characterising Knowledge Work

The orientation of this research places a strong emphasis on the human dimensions to technologies for supporting organisational memory and expertise. The history of interactive computing shows repeatedly that it is the human issues which 'make or break' new methods and tools at work. If we use the analogy of a river to describe the 'work flow' at the level of an individual, team, or organisation, the designers of a new method or technology for organisational memory are placed in the role of 'river engineers' seeking to change the flow of the river in some way. What they want to do is tap into the deep currents of the river, channelling it in new, productive directions. The question is, do they understand the hidden currents, eddies, and dynamics of that river sufficiently? If not, the result can be destructive 'interference patterns' in the flow, or the force of the deeper currents may simply re-route around the changes.

It is, therefore, worth trying to clarify some of the salient properties of 'knowledge work', given our intention to enter and change this fast flowing 'river' with technologies. Two perspectives are considered: an empirical study of knowledge workers, and foundational work on characterising the properties of many real world problems.

2.1 A Study of Knowledge Workers

Firstly, on the basis of field studies of knowledge workers, Alison Kidd [Kidd 1994] has proposed several features which distinguish *procedural work* from *knowledge work*. All work is invariably a mix of the two, but increasingly, the procedural features are giving way to knowledge-based features. Kidd makes a number of distinctions, which are paraphrased below.

Knowledge workers are changed by the information in their environment, and they in turn seek to change others through information. Information is to be consumed, and once 'digested', is often of little further value. Information resources which may have longer term use are often left visible and uncategorised (hence the frequent untidy piles and whiteboards), so that they can be quickly referred to. This is the antithesis of more procedural work (e.g. a secretary or administrator), whose work requires a lot of *filing* into *inflexible* structures -- inflexible because the scheme is often standardised across the organisation, and because other staff also need to access those files.

Diversity and ad hoc behaviour patterns are common in knowledge work. New information is sought out, reused, and passed on in opportunistic ways, dependent on the changing context and interleaving of the worker's activities. In contrast, consistency of method and output is important in procedural work.

Communication networks are highly variable, with different patterns and use of media. Teams form and disband within the space of a day. The structure and job titles on an organisation chart are thus even less indicative than usual as to what someone does or with whom they work. Much of the knowledge exchanged is embedded in documents and email. Staff engaged in predominantly procedural work tend to have well-defined responsibilities and relationships, and the information flow that they maintain is more clearly defined.

These features provide a useful orientation to the domain of concern. They paint a picture of knowledge workers, and consequently their host organisations, as existing in continual flux as teams form and reform. In particular, the mobility of employees within and between organisations (coupled with 'out-sourcing' to external contractors) creates conditions that can more easily lead to the fragmentation of any persistent shared memory within a team about lessons learned in projects. Furthermore, keeping track of discussions, decisions and their rationale is made harder when teams form on a project-specific basis, proceed to work interdependently but with substantial autonomy, and then disband. Experiences are not commonly recorded in conventional documentation, remaining locked in individuals' memories -- individuals whose memories will fade, or who will take their expertise to other

jobs. These are both motivating factors for, and militating factors against, the development of organisational memory resources. Collaboration tools which do not impose rigid models of membership or role, and which are able to integrate many diverse media types would seem appropriate in such an environment, discussed further by Kidd.

2.2 Wicked problems

The second perspective on knowledge work comes from the formative work of Horst Rittel [Rittel 1972] [Rittel 1973]. Whilst the term 'knowledge work' was not in currency in the late 1970s, Rittel identified crucial features of intellectual work which are highly pertinent to current concerns. Rittel characterised a class of problem which he termed 'wicked', in contrast to 'tame' problems. Tame problems are not necessarily trivial problems, but by virtue of the maturity of certain fields, can be tackled with more confidence. Tame problems are understood sufficiently that they can be analysed using established methods, and it is clear when a solution has been reached. Tame problems may even be amenable to automated analysis, such as computer configuration design or medical diagnosis by expert system.

Wicked problems display a number of distinctive properties that violate the assumptions that must be made to use tame problem solving methods. Wicked problems:

- cannot be easily defined so that all stakeholders agree on the problem to solve;
- require complex judgements about the level of abstraction at which to define the problem;
- 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 -- these must be discovered;
- often have strong moral, political or professional dimensions, particularly for failure.

The connection between wicked problems and knowledge work should be apparent. Such problems are the typical challenges faced daily in, for instance, software design, government or social policy formulation, and strategic planning in organisations. It is also the case that wicked problems and lessons learned pose particular challenges for analysis and support by knowledge-based systems. What then is involved in supporting the capture of organisational expertise for such real world problems?

3. Negotiation, Argumentation and Knowledge Work

Let us develop the concept of negotiation, as introduced at the start. The claim is that knowledge work is dominated by communication, specifically *negotiation* and *argumentation.* There are several reasons for this.

- Firstly, much knowledge work is conducted in teams, and members have to communicate, increasingly distributed in space and time.

- Secondly, external factors often remove the control that a team has; the problem space is not stable. Goals, constraints and stopping rules are continually shifting. This demands a mode of working in which requirements, constraints and solutions must be regularly re-negotiated.

- Thirdly, Rittel concluded that wicked problems can only be tackled through what he termed an *argumentative* method [Section 4]. Understanding how to frame a wicked problem is the key to finding solutions: what are the key questions?; what are the key priorities?

- Fourthly, knowledge work is increasingly interdisciplinary. The different backgrounds, assumptions and agendas which members bring to a team can be extremely creative, but the inevitable conflict, debate, negotiation and compromise which is involved in reaching such creative solutions must also be acknowledged; this process can then be turned to the team's advantage.

In summary, an approach to capturing and representing organisational memory is required which is capable of supporting knowledge teams in:

- representing and reconciling multiple stakeholders' perspectives;

- re-negotiating project priorities in response to changed circumstances;

- communicating the rationale for decisions to others;

- recovering insights and solutions from past scenarios, to avoid 'reinventing the wheel'.

An organisational memory strategy which recognises the centrality of negotiation and argumentation in its employees' workflow (recalling the river metaphor) assumes from the start that the knowledge invested in a typical project is the product of much argument, compromise and the reconciling of different perspectives.

4. Visualising Argumentation

In *The Next Knowledge Medium* [Stefik 1986], Stefik proposes collaborative argumentation tools as one example of knowledge media. Such tools, "for arguing the merits, assumptions, and evaluation criteria for competing proposals" could provide "an essential medium in the process of meetings." "The languages provided by the tools encourage an important degree of precision and explicitness for manipulating and experimenting with knowledge", coupled with "augment[ing] human social processes." This conception of knowledge media lies at the heart of the representation and support technologies now proposed.

On the basis of his analysis of wicked problems, as introduced above, Rittel proposed the *IBIS* (Issue Based Information System) method, which encourages team members to actively discuss problems by raising *Issues* that need to be addressed, *Positions* in response to those Issues, and *Arguments* to *support* or *object-to* Positions. Conklin *et al.* [Conklin 1988] [Conklin 1991] developed a hypertext prototype called *gIBIS* (graphical IBIS) to support Rittel's IBIS method. In gIBIS, a team conducted its debates by building a graphical 'conversation map'. [Fig. 1] shows the gIBIS scheme, which illustrates how cumulative argument construction and critiquing can take place around a shared, graphical argumentation structure.

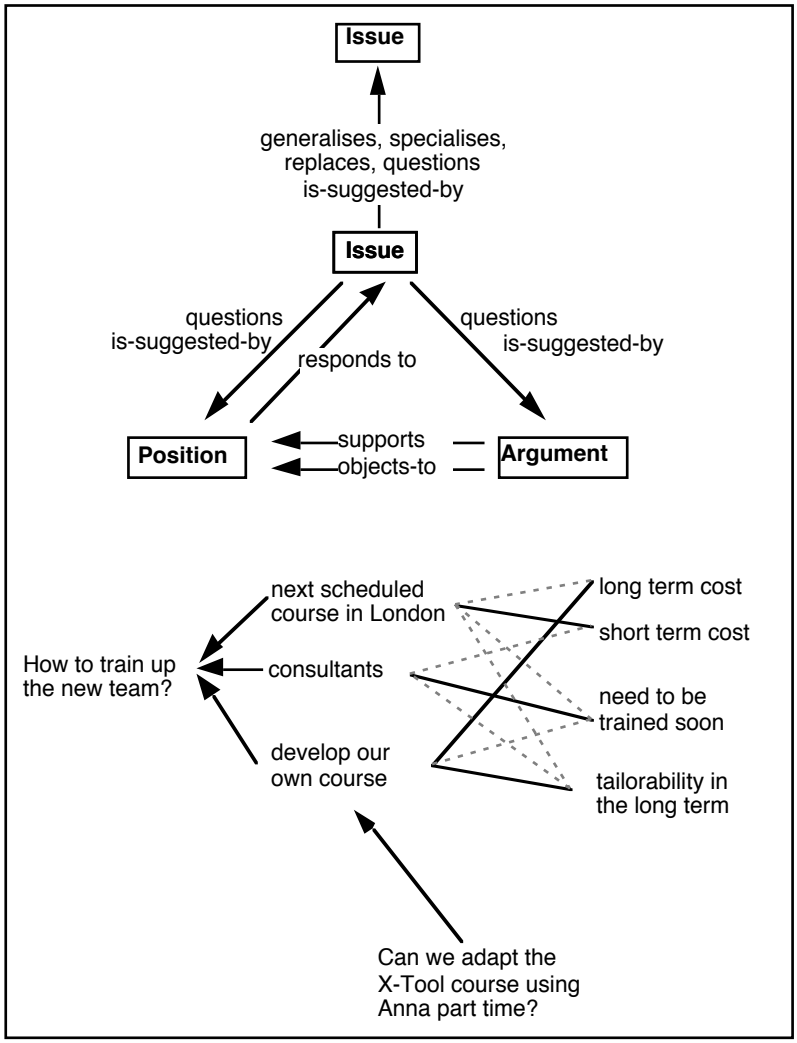

Figure 1: The graphical IBIS (gIBIS) notation [Conklin 1988] and an example, showing how it enables a team to cumulatively build graphical argument spaces.

Many others have since developed variations on gIBIS. The complexity of the notation, and its visual layout rules (which vary with different approaches), determine how large and elaborate an argument can be expressed. For instance, a more expressive argument schema is shown in [Fig. 2]. The *Decision Representation Language* [Lee 1991b] for supporting debate and qualitative decision making, introduces new constructs (e.g. the *Goal* node type), and allows participants to explore *Alternatives*, *Claims* backing them, and to contest through *Questions* and counter-*Claims* the relationships between these constructs.

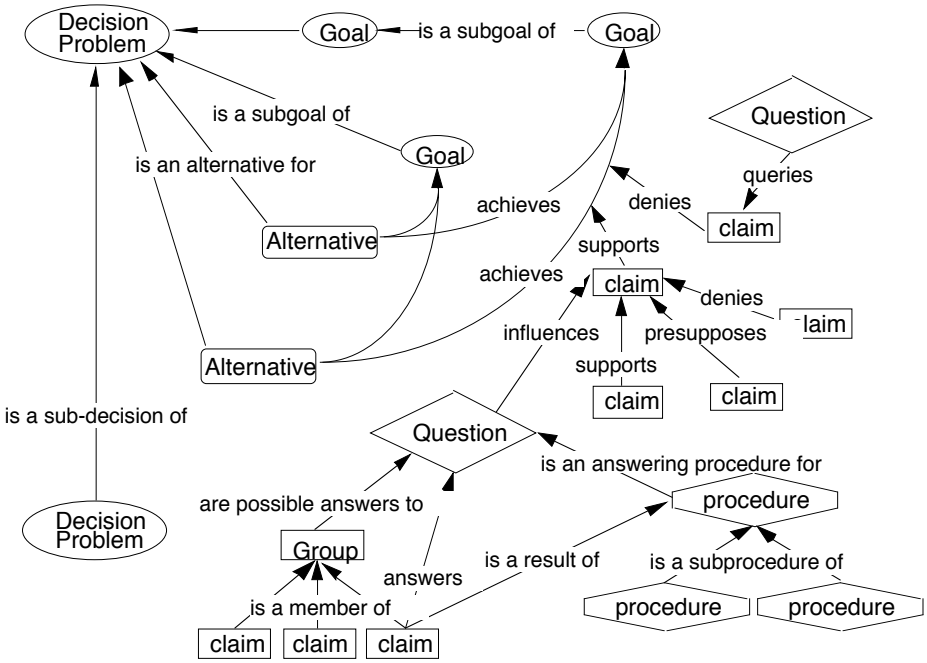

Figure 2: The Decision Representation Language, one of the most expressive notations for capturing collaborative arguments [Lee 1991b]. A support tool [Lee 1990] provides graphical and tabular views of the underlying argument network.

This paper focuses on notations like IBIS, which are 'lighter weight' than DRL, the emphasis being on suitability for quick and intuitive use during meetings. A similar notation to IBIS is *QOC* (Questions, Options and Criteria) [MacLean 1991], on which much of the usability evaluation work reported later has been based.

To summarise, having proposed that negotiation and argumentation are central to knowledge work, and having introduced the representation schemes which allow us to visualise such processes and products, let us now consider the technological support required. IBIS and QOC style representations have been used effectively with paper and pen, but computer supported argumentation is needed for easy editing, scalability and flexible linking, as discussed next.

5. Collaborative Hypermedia Infrastructure

Hypermedia is an ideal technology for capturing *informal knowledge types* with *inter-relationships which are hard to formalise*. This is in contrast to repositories that rely on more structured knowledge bases, requiring well-defined knowledge types and structures. The power that one gains from the latter comes at the cost of initial knowledge engineering effort, perhaps requiring a specialist. Moreover, as argued earlier, since the subject matter of most interest in knowledge work is often hard to formalise or continually changing, realistically, this encoding effort may be hard to justify even if it were possible in principle.

The evidence from cognitive studies of wicked problem solving points strongly to the importance of opportunistic ideas and insights. Hypermedia graphical browsers are well suited for linking together ideas without having to specify the precise semantics of their relations or roles (though see [Buckingham Shum 1996a] [Buckingham Shum 1997b] who reports that for certain types and stages of problem solving, even semiformal schemes can be too formal, impeding the creative flow).

Hypermedia is also well suited to organisational memory capture in a second essential respect: *media integration*. Debates, decisions and rationale do not exist in a vacuum, but in relation to ongoing work which relies on, and generates, many forms of artifact (e.g. faxes; email; reports; sketches; prototypes; simulations). It is crucial that these different artifacts can be integrated into the debates captured as semiformal argumentation. Hypermedia systems were designed precisely for this kind of media structuring, as exemplified in the the *QuestMap* hypermedia groupware system [Conklin 1993][GDSS 1996], shown in [Fig. 3]. This system is derived from the gIBIS research prototype described earlier [Fig. 1].

Finally, a review of the role of hypermedia cannot ignore the World Wide Web, the first truly global hypermedia system. In response to the need for tools to support asynchronous discussions between geographically dispersed participants, we are now seeing the emergence of Web systems to support argumentation of the sort illustrated above. One example is *HyperNews* [LaLiberte 1995], a system which supports discussions as textual threads through a combination of hierarchical indentation, augmented by icons which indicate whether a contribution is for example, an agreement, disagreement, or new idea. [Fig. 4] shows an example of argumentation on the Web (using a version of HyperNews), taken from an electronic journal peer review debate between an author and several reviewers, adopting an argumentation-based approach described in [Sumner 1996].

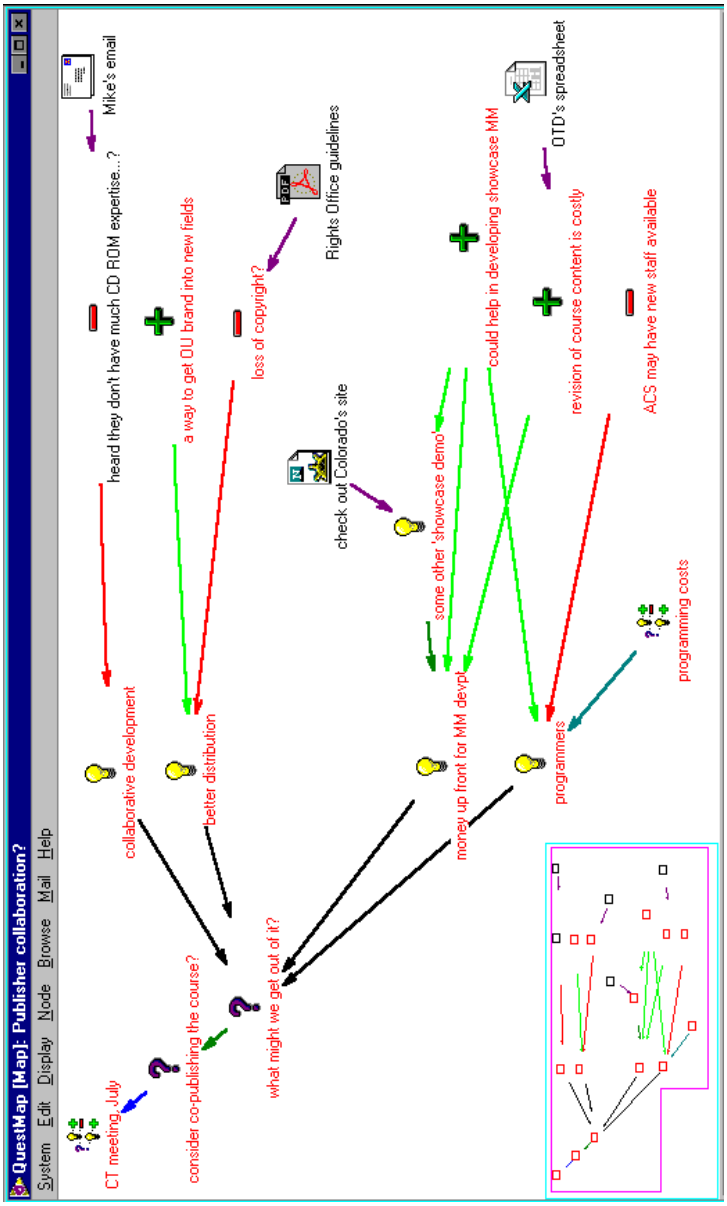

Figure 3: A screen from the QuestMap system [GDSS 1996]. Based on Rittel's IBIS argumentative model, this hypertext groupware system provides teams with a way to conduct synchronous or aynchronous debates. Ideas *are suggested in response to* Questions*, and their* Pros *and* Cons *traded off against each other. New* Questions *can be raised by any element of previous discussion. Other media can be integrated into the web of debate through* Reference *nodes (e.g. reports; spreadsheets; video; presentations; code).*

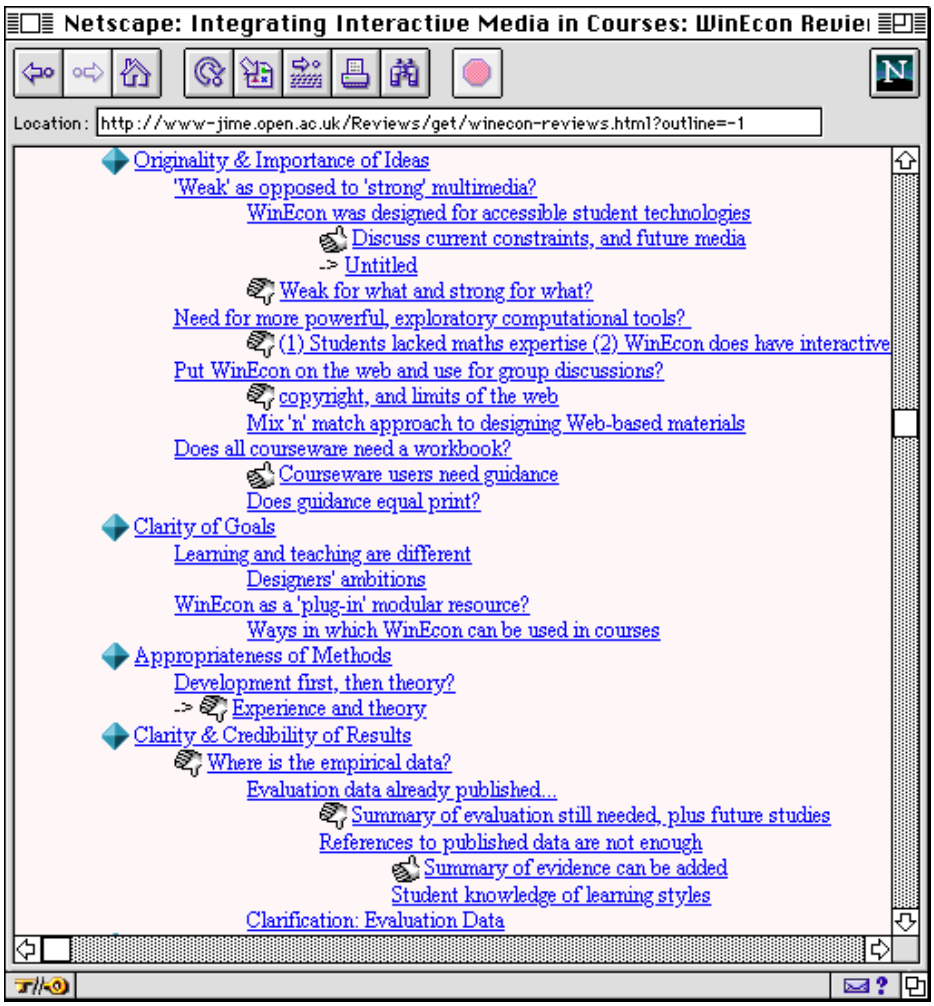

Figure 4: Web-based argumentation in the context of journal peer review [Sumner 1996].

Such systems represent first generation Web argumentation tools. A similar textual outline representation was used in one of the most significant design rationale case studies [Burgess Yakemovic 1990], summarised in [Section 7]. The Web is still a highly impoverished hypermedia system compared to many other systems, indeed, its simplicity is a major factor contributing to its explosive growth [Buckingham Shum 1997c]. However, with richer hypertext models [Bieber 1997], and the possibility of richer interactivity on the Web through developments such as Java and browser plug-ins, direct-manipulation graphical interfaces on the Web will become commonplace (e.g. [Kremer 1996]).

6. What Kinds of Knowledge are Captured?

The use of a tool like QuestMap [Fig. 3] allows teams to visualise their discussions, past and present. The following scenario may help to concretise how this might work in practice:

> In June 1995, a meeting agenda is circulated specifying the *Questions* to be resolved. Over the network and in their own time, the multidisciplinary team members prepare by tabling their *Ideas* , beginning to critique these with *Pros* and *Cons*, linking in relevant reports, costings etc. In the meeting, the debate is projected onto a large wall to track the strengths and weaknesses of each idea as it is explored; following the meeting, team members reflect on the decisions made, and continue to discuss them, updating the map as new results and ideas come in. This map is emailed to others who were not present, who can quickly see what issues were discussed, which ideas were rejected, what decisions made, and on what basis. In September, several issues debated in June suddenly become critical. The relevant part of the map is retrieved, and it is realised that several Ideas rejected then are now valid. Moreover, links were created in June's meeting back to a previous discussion in May 1994, when a similar problem had been elegantly resolved. This provides a clue to the team as to how to resolve the current issues.

This scenario illustrates the affordances of an organisational memory resource coupling hypertext with argumentation. Firstly, it *supports the process* of discussion and negotation between multidisciplinary stakeholders. Secondly, it *captures the products* of those negotiations, providing the basis for an organisational memory. A team using such a tool builds for itself a form of intellectual trace which they can then draw upon. A group memory based on such a trace can help find answers to the following kinds of question:

- Have we faced problems similar to this before, and what was done?
- Who identified this problem/suggested this solution?
- What solutions were considered, but rejected, and why?
- If we change this decision, what might be affected?
- What led to this document being changed?
- What were the main criteria taken into consideration when that decision was made?

A resource based on this kind of approach clearly cannot represent all classes of organisational expertise; it should be seen as occupying one niche in the design space of tools to capture and maintain different organisational knowledge types. Some types of organisational expertise are without a doubt amenable to storage in more conventional databases, such as patents, procedures, employee qualifications, reports, etc. 'Intellectual auditing' [Brooking 1996] can help to identify this kind of intellectual capital.

However, a strength of the approach described here (discussed further by Conklin [Conklin 1996]), is that the knowledge is captured *collaboratively,* and *in situ*, during the meeting or asynchronous debate, in the immediate context of one's work. Knowledge is represented, stored and indexed in relation to the real activities by which one's work is accomplished (as well as through some more abstract indexing system if so desired). Discussing *through* the medium of collaborative, graphical

argumentation *eases the transition* from the messy, changing, contextualised, social, multimedia world, to their abstracted entry in an organisational memory system. As entries are made in the organisation's long term memory, they bring with them (in the form of the web of discussion and work artifacts) important elements of the context in which they arose. Such cues are frequently used to recover memories [Eldridge 1992].

7. Argumentation in Use

Collaborative, hypermedia argumentation has been tested since the mid-1980s to support knowledge work in a variety of contexts. Most of the earlier work on argumentation was taking place in research labs on the leading edge of the emerging technology of hypertext, for which graphical argumentation became something of an experimental 'white rat' for testing technological flexibility. However, more recent research has placed an increasing emphasis on application to real, small-medium scale projects. This section points interested readers to more detailed reports of such studies. More detailed reviews of the research cited below can be found in [Buckingham Shum 1994] [Buckingham Shum 1996b].

Firstly, and not surprisingly, there has been a longstanding interest in the contribution that collaborative argumentation can make to complex, intellectual work where the quality of reasoning and accessibility of rationale for decisions are particularly important. Experimental fields of application have included government policy formulation [Conklin 1988] [Rittel 1973], scientific reasoning [Smolensky 1987] [VanLehn 1985], and legal analysis [Newman 1991]

As hypertext matured as a technology, some of the most significant design disciplines began, and continue, to look at collaborative argumentation as a way to capture project/organisational memory, and manage the kind of changing environment and competing agendas described earlier. *Argumentative design rationale* is attracting substantial interest in *Human-Computer Interaction* [Carroll 1991] [MacLean 1989] [Moran 1996], *Software Engineering* [Conklin 1989] [Jarczyk 1992] [Lee 1991a] [Potts 1988] [Potts 1994] [Ramesh 1993], *Knowledge Engineering* [Stutt 1995] [Vanwelkenhuysen 1995], and *Knowledge-based Design Environments* [Fischer 1991] [Garcia 1992].

Thus far, the only financially costed benefits of this form of organisational memory come from a software engineering case study which introduced a textual version of IBIS argumentation, similar in form to the outline view provided by the HyperNews Web system [Fig. 4]. This was used by a team working on a large commercial system development [Burgess Yakemovic 1990]. The study reports the discovery of eleven design flaws during the conversion of argumentation from outline to graphical form. The time savings gained for the project as a result were estimated at between three and six times greater than the time cost of converting the argumentation formats. It is evident that, as with any new tool, the success of IBIS in this case owed much to the enthusiasm of the team using it, in particular the maintainer of the issue base. Organisational practices and cultural differences in other teams were obstacles that prevented the uptake of the approach more widely (see [Section 8.3]. The availability

of tools like QuestMap [Fig. 3] helps to make the approach more widely available, and in time should clarify the strengths and weaknesses of this particular approach in the context of different organisational cultures.

8. Hands-On Practicalities

In this section, attention focuses on the practicalities of using argumentation schemes. It is all too easy to propose new tools which should work in principle, only to find that insufficient account has been taken of the actual demands that they make in real work settings (borrowing our earlier metaphor, the force of the 'river' may be underestimated).

8.1 The Cognitive Costs and Benefits

Organisational memory of any sort comes at a cost -- someone must construct, index, and maintain it. There is no way for a knowledge capture enterprise to avoid this cost-benefit tradeoff. It is a question of how to negotiate it. Thus, minimal capture effort initially (e.g. video-record every meeting and store every document), simply shifts load downstream (how to recover the relevant records from memory?). In turn, the initial investment of knowledge encoding/engineering effort provides computational services subsequently.

Midway between these two extremes, the semiformal hypertext approach described here enables knowledge workers (not knowledge engineers) to structure their deliberations using a high level, reasonably intuitive vocabulary (e.g. Questions, Ideas and Arguments). What are the overheads introduced by such schemes?

Analysis of the hands-on practicalities of using such a scheme [Buckingham Shum 1996a] [Buckingham Shum 1997b] has highlighted four key cognitive tasks:

- *Unbundling* -- teasing apart ideas which have been expressed together. A typical example would be when in one utterance someone raises a problem, and proposes a solution plus supporting reasons. Much time is wasted in meetings because a disagreement with one element in an argument is taken to be a dismissal of the whole argument. Graphical argumentation can clarify the different elements and hidden structure.

- *Classification* -- deciding whether a contribution is, e.g. a Question, Option or Criterion. This is not always as simple as it sounds, because Options and Criteria may initially be expressed as Questions, or Criteria as solutions. A Yes/No Question can be asked about a particular Option, rather than clarifying the implicit problem to which that Option is one candidate solution. The task here is to cut through the surface form and recognise the 'deeper content.'

- *Naming* -- how to label the new contribution succinctly but meaningfully. It can often be difficult to articulate ideas succinctly. 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 also dependent on the anticipated future use of the record, 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 taking over from you? [Section 8.4]

- *Structuring* -- how a new element relates to other ideas. Many 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 sufficiently similar to another in a different context, or should a new Question be introduced? Has this Criterion already been used elsewhere under a different name?

There is evidence that the intellectual rigour that this process encourages (e.g. being encouraged to ask 'what really is the key Question here?') can focus team meetings about complex, wicked problems [Buckingham Shum 1997b]. There is also evidence that when a problem is not in fact wicked, structured argumentation may not be helpful, slowing down discussion unproductively. It is therefore a case of choosing the right tool for the job; argumentation integrates well with certain cognitive and group workflows, but obstructs others. We have sought to alert practitioners to these hands-on issues when training them [MacLean 1992-94].

8.2 Modes of Groupwork

How can collaborative argumentation be used in a meeting? What role should it play in the project? There is a range of roles, depending on how committed a team wishes to be to capturing its intellectual investment in this way (see next section for factors that may militate against this). [Fig. 5] shows various points along a continuum which illustrate options which a team can adopt according to their work patterns.

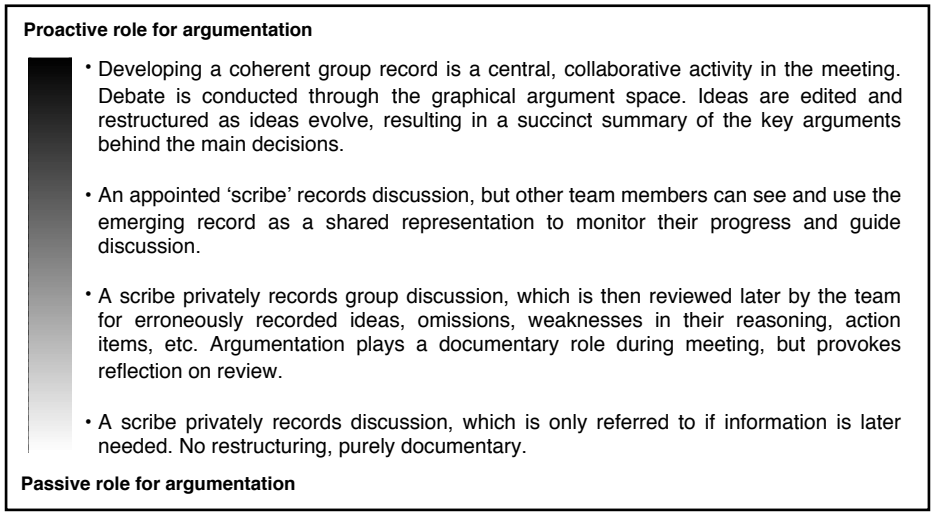

Figure 5: Graphical argumentation can play a proactive or passive role in team deliberation. The more a team learns to interact via the graphical argument space, the more transparent it becomes -- construction of the group memory becomes increasingly a co-product of the deliberation process, jointly owned by the team, and a living resource on which to build subsequently.

A team will in fact move back and forth along this continuum for different kinds of meetings, and indeed *within* a meeting depending on the kind of problem that is being discussed (see previous section). We expect organisations, and within them individual teams and team members, to adapt these generic representations and tools to their own priorities and work patterns. Almost invariably, a new method or tool will be used in ways never originally envisaged by its developers; for instance, an innovative use of QuestMap for business modelling is described by [Selvin 1997].

8.3 Organisational Culture

Understanding the human dimensions to a work representation cannot be restricted to the impact on cognition or group dynamics, critical though these are. As discussed in [Section 9.1], representations take on political dimensions as soon as they are introduced into a workplace [Bowers 1991]. Collaborative argumentation requires the adoption of a relatively open, transparent mode of communication, negotiation and accountability. Such an approach contrasts sharply with the harsh realities of some cultures, where there is distrust between employees and managers, and where efforts to improve meeting process, listen to all stakeholders', and make rationale more explicit are alien. [Grudin 1996] and [Conklin 1991] have suggested that employees might, for instance, refuse to document who made a particular decision and why, for fear of recriminations in the event of an error. Moreover, certain stakeholders may perceive such approaches as undercutting their power, since their arguments will be represented and treated on a more equal footing with other team members' views.

Once displayed in the argument space, an idea is less tied to its owner, and more vulnerable to rationale critique. Conversely, for some stakeholders, this will be empowering.

Ultimately, we cannot escape the fact that organisational memory, certainly of the sort described here, requires a compatible working culture. There can be little doubt that even for team members who know each other well, there is a process of negotiating mutually acceptable conventions for maintainng the group memory [Berlin 1993]. This must take place on a correspondingly larger scale to prevent an organisation-wide memory from dying through neglect or subversion, as seems to be the fate of so many new methods and tools which do not sufficiently appreciate the organisational dynamics they seek to change.

One may hypothesise that current excitement within the organisation and business literature about the shift to 'learning organisations' will create work cultures who will look favourably on collaborative argumentation tools. One may also hypothesise that the dynamic of change is two-way, and that in the hands of a committed team able to demonstrate its relevance to the organisation's business, collaborative argumentation tools could work from the bottom up as agents of change.

8.4 Negotiating the 'Context Paradox'

Information becomes useful knowledge once its significance in its original context is understood; divorced from its context, information is open to misinterpretation. In engaging in the enterprise of constructing organisational memory, therefore, we are faced with the challenge of effectively capturing sufficient context to accompany entries in the information base. What can be termed the "context paradox" is the possibility that more context will be needed to interpret whatever contextual information has already been provided. Attempts to provide richer, more extensive contextual information through, for instance, audio/visual multimedia commentaries, or more complex hypertext webs of information are still prey to the reinterpretation problem. A related irony is that the more contextual background there is to digest, the less likely it is that busy staff will do so.

The degree to which additional context is needed to interpret information correctly clearly depends on who the recipient of this information is. In creating what is intended to be a reusable resource, careful thought needs to be given to the user groups one is serving. For instance, colleagues who are co-present in a meeting have established a rich context in that time and place for intepreting each others' contributions. A video recording may help an outsider recover important elements of this, although not everything is captured on camera, and of course, prior knowledge of the context of the meeting may be critical to make sense of it. Tools are now being developed to assist in capturing important moments in meetings, and managing that corpus of material [Moran 1997].

As the intended user base of a group memory system expands from the core team, to encompass wider circles of staff, the common ground which can be assumed decreases, thus increasing the amount of implicit knowledge that needs to be made explicit. One way to think about this process is as the evolution from a group memory

for *unstable, provisional information* kept for the core team's own use, to a memory for more *stable, consensus information*. This corresponds to shifts from implicit to explicit knowledge, from being a private to a public resource, and from being a one-off entry (e.g. to facilitate a single meeting), to being a reusable resource of wider interest. Berlin *et al.* [Berlin 1993] have also described how the group's process must adapt when they commit to maintaining a group memory, even for themselves, as individual styles of entry must be held in tension with establishing agreed conventions.

How does the context paradox translate with respect to the particular approach presented in this paper? Graphical argument/document networks of nodes are quite terse compared to textual documents. They capture the essence of discussions, leaving the original participants to 'fill in the gaps' with their own memory -- the network is a resource to cue them. There is some empirical evidence that outsiders can have difficulty in making sense of someone else's graphical argument structures when they have not been involved in the original discussions [Bellotti 1995] [Shum 1993]. As emphasised earlier (based on evidence such as these studies), one solution is to tightly integrate the argumentation with the relevant documents, making it very easy to bring up a relevant document. Open hypermedia systems (e.g. Microcosm [Multicosm 1996a]) make it easy to link from point to point in any desktop document running in Microsoft Windows, and Webcosm extends this to web documents [Multicosm 1996b].

Another approach is to enrich the argumentation with expert commentary from one of the original team, who can introduce the discussion, much as a colleague might set in context some documents that they are handing over. With off-the-shelf products such as ScreenCam [Lotus 1996] for instance, one can easily record commentary to accompany a visual walkthrough of a map to introduce a particularly complex analysis, and for instance, bring out nuances behind particular arguments that are invisible. Subsequent users would play this guided tour first, to get an overview of the discussion they are about to step through in detail.

The key information design task is to design for different user populations, and to use different representations of context appropriately. Graphical argument structures have different cognitive affordances to time-based media. The latter can be very effective in conveying subtle information that is hard to express in graphical/textual summary form, whilst the latter provides an overview of the discussion space, and as a shared representation, supports collaborative reasoning about a problem. Detailed analysis of the individual and group cognitive affordances of graphical argumentation in a design context is presented in [Buckingham Shum 1997b].

To conclude this section, as the context paradox emphasises, efforts to supply richer context are still open to misinterpretation, and unless carefully designed, may be ignored due to information overload. If well designed, however, fewer people will lack important context, since the circle of readers who now share key common background knowledge has been widened. It is worth re-iterating that if a group memory is successful in providing contextualised information, what the reader will come to share with the team is not only an understanding of what they did and why, but also an appreciation of the tensions and trade-offs that set the context for those decisions.

9. "Knowledge (Management) is Power": Ethical and Representational Issues

This paper has intentionally focused on technologies embedded in contexts of use, seeking to elaborate scenarios of organisational memory usage as a way to highlight future possibilities, and to identify obstacles to uptake. This foregrounding of the human dimensions to knowledge technologies is extended in this final section to issues of power and control over *what* gets represented and *how*, *by whom*, and *for what purposes*. Our starting point is the fundamental issue of *representation.*

9.1 The Politics of Formalisation

> In selecting any representation we are in the very same act unavoidably making a set of decisions about how and what to see in the world. [...] A knowledge representation *is* a set of ontological commitments. It is *usefully* so because judicious selection provides the opportunity to focus attention on aspects of the world we believe to be relevant. [...] In telling us what and how to see, they allow us to cope with what would otherwise be untenable complexity and detail. Hence the ontological commitment made by a representation can be one of the most important contributions it offers. [Davis 1993]

> Classification systems provide both a warrant and a tool for forgetting [...] The classification system tells you what to forget and how to forget it. [...] The argument comes down to asking not only what gets coded in but what gets read out of a given scheme. [Bowker (in press)]

The above two quotes, the first from knowledge engineers, and the second from an ethnographer of organisational memory, draw attention to the filtering function that a representation provides, and the problem that through the process of simplifying a domain in order to describe it within a formal scheme, we may also be systematically factoring out certain classes of critical information simply because they are hard to formalise.

Whenever an authoritative body (e.g. corporate management, or a research funding council) declares an interest in certain concepts, it is inevitable that its dependents (e.g. managers, or researchers seeking grants) will seek to align their activities with these concepts in order to maintain a presence. The first point, therefore, is that the introduction of systematic KM (whether or not technology is involved) creates a new *economy of knowledge* and a *knowledge vocabulary.* Any group and their work will remain invisible and thus unresourced unless they can represent themselves within this new economy, using the right language. Bowker presents an illuminating analysis of the impact of 'professionalisation' -- systematic classification of skills and courses of action, and management of these via technology -- on a profession in which expertise takes the form of hard to codify tacit knowledge and craft skill, in this case nursing:

> One of the main problems that [...] nurses have is that they are trying to situate their activity visibly within an informational world which has both factored them out of the equation and maintained that they should be so factored -- since what nurses do can be defined precisely as that which is not measurable, finite, packaged, accountable. [Bowker, (in press]

This illustrates clearly the political dimensions to formal classification. The names and labels one uses unavoidably emphasise particular perspectives (see also [Suchman 1993] on the politics of computational categories in CSCW).

Knowledge-based systems require the systematic decomposition and classification of expertise; a knowledge-base unavoidably 'holds' an ontological view of the world (ontology with a small 'o'). More recent knowledge-sharing initiatives and other research devoted to formal Ontologies make more explicit the issues faced in knowledge modelling, independent of any particular symbolic implementation as a system. One question that the ontology community may be able to help answer is how to manage the inevitable incompletenesses and inconsistencies in an organisational knowledgebase, due to uncodified, or uncodifiable knowledge. If ontology building is to form part of AI's contribution to KM (as some argue), how can we ensure that areas of uncertainty or incompleteness are made explicit in the ontology, *and* carried through to the implementation and user interface of any KM system based on that ontology? If the KM system is to be used by the organisation's managers, then they must be sensitised to the limitations of the tool's ontology, as a check and balance to the seductive sense of control that manipulating clean computational abstractions offers. What training is required in order to wield such tools intelligently?

I have argued elsewhere [Buckingham Shum 1997a], that some of the most robust forms of knowledge sharing and communication that we know occur in organisations are socially based, and their content is extremely hard to formalise. These include the discussions that endow documents with significance [Brown 1996), the informal recounting of technical 'stories' to colleagues to pass on new insights [Orr 1986], and the importance of dedicated knowledge analysts to maintaining knowledge resources, and both persuading and assisting staff to access them [Davenport 1996]. [Fig. 6] schematically illustrates these three processes.

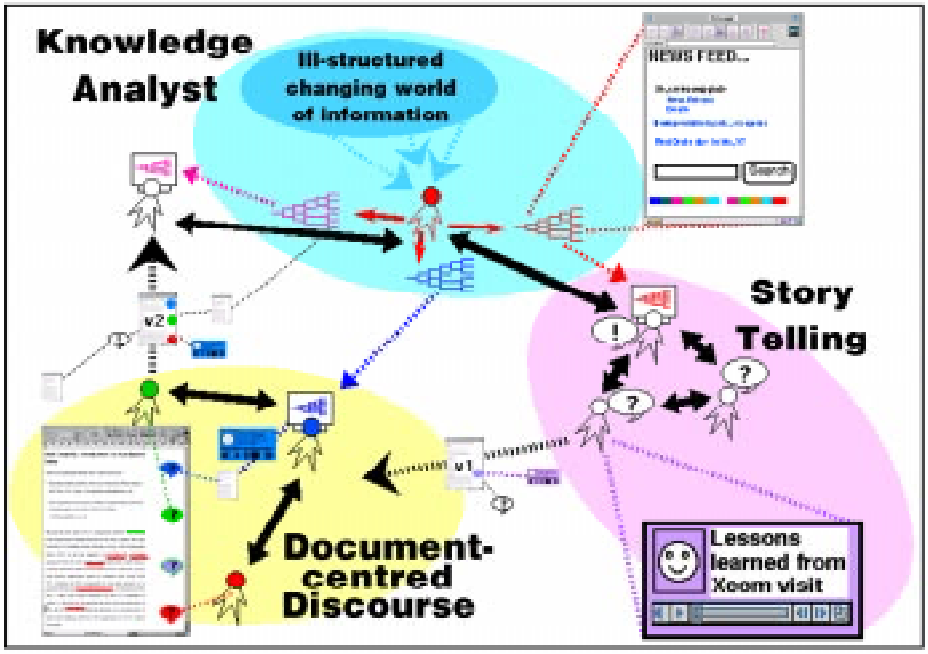

Figure 6: Pro-active knowledge analysts, technical 'story-telling' amongst staff, and document-centred discourse are three ways in which knowledge is shared within organisations. Media that are now emerging within many organisations to support these processes are illustrated -- Web intranets integrated with agents and broadcast media, desktop audio/visual recording tools, and document-discussion environments. Their integration with AI techniques is discussed more fully in [Buckingham Shum 1997a].

As illustrated, the representations and technologies that should be considered for such processes may well be rather different to the knowledge-based technologies with which we are currently familiar. (As an aside, there appeared to be a strong sense at a recent symposium on AI's role in KM [AIKM 1997] that formal representation of knowledge seems to have a limited role to play in organisational knowledge management, with the emphasis shifting to supporting the social, coordinated *processes* through which knowledge is constructed.)

To conclude this section, the representations we use shape the world we can see through them. All representations are simplifications; the question is are they *over*-simplifications? The baseline assumption in the argumentative approach is that there rarely is one correct view of the world to begin with; the first step is to take seriously the different viewpoints, and to then seek ways in which these can be expressed and resolved. As discussed above in relation to the context paradox, however, no representational scheme is immune from the danger that it becomes too simplistic, too terse to be useful, or too decontextualised to support meaningful interpretation.

9.2 'Participatory KM' Based on Stable, Sanctioned Knowledge

> *Dear Employee,*
>
> *In order to maintain and increase KnowTech's competitiveness, an intellectual audit is to be conducted on your department in the coming month, as part of a corporate wide strategy. This will provide Strategic Planning with a better understanding of your skills, communication networks and contributions to KnowTech's business. This will enable them to ensure that you are receiving the right information at the right time, and that we make the most of your valued expertise.*
>
> *The Management*

Software design is the process of moving from vague requirements to executable, computational models. Participatory design approaches to interactive system development emphasise the many stakeholders in a system development project, and the need to involve the system's end-users in order to co-design software and work practices. This final section draws on the participatory design perspective to examine the particular challenges that KM technologies face if they are to be collectively 'owned' by the staff whose knowledge is being managed. I return again to the foundational theme of representation that runs through this paper, identify the stakeholders that a participatory approach should involve, and then propose a heuristic measure for deciding when to commit to formally representing knowledge processes.

Knowledge-system design, as a particular form of software design, is the construction of computer-manipulable representations of domain knowledge. The process of *formal representation* raises a host of issues, some of which this paper has considered. From a participatory design perspective, three of formalisation's most significant features in a KM context are as follows:

1. *Representations can become less flexible,* that is, as layers are added, dependencies on old structures increase, and the whole structure becomes harder to change in response to changes in understanding, or of the domain being modelled. Representations tend also to become less tolerant of incompleteness, inconsistency, or ambiguity. This is of course useful for highlighting weaknesses in an organisation's KM, but it may also be a significant limitation, since the models that different parties hold of a domain may be equally valid, but shaped by competing priorities. It may not be possible to satisfy these with one elegant representation. The cost of formalising too early, even semiformally as hypertext, is that it may be too much effort to revise a representational scheme that turns out to be wrong, so it is left as part of the system. Clearly, the art is in knowing when to formalise.

2. *Representations become less comprehensible* to staff who are not knowledge-engineers. One of the consequences of formalisation is that the contents become increasingly inaccessible to the majority of stakeholders. It is of course common that a profession's language and representations are opaque to outsiders, but extra care needs to be taken in KM-system design, due to the

legitimate interests of different stakeholders in knowing what is to be encoded in the system, and what role this is playing in management decision making.

3. *Representations support automated analysis.* Clearly this is the main purpose of formalising, so why should this be a problem? Problems arise when the processes of decomposition and abstraction, required to create a representation capable of supporting automatic analysis, result in models which strip out important contextual details which are in fact critical to understanding the domain (see [Section 8.4] on the problem of 'capturing' context). Models of employees' skills, work processes and interdependencies may not adequately express the true nature of their expertise and coordination of work. If the representation is too incomplete (it will always be incomplete to some degree), then the most powerful manipulations and analyses are meaningless. This of course is not a novel insight, but organisational dynamics are particularly difficult to model.

It is rare to find knowledge modelling papers that explicitly recognise the informal and social knowledge processes in the organisations (real or imagined) for which they are designing (though see [Euzenat 1996] [van Heijst 1996] [Vanwelkenhuysen 1994] for exceptions). Combining social and computing disciplines in this way is surely a fruitful area for further multidisciplinary work, as exemplified by [Fischer 1995]. The formality and accessibility of knowledge representations are central to a participatory KM approach.

Who are the main stakeholders in a KM initiative, and what are their concerns? Obviously, management want to know how can they make the most of their investment in quality staff and hope that systematic KM will give them views and benchmarks on the organisation's state. For a company's information technologists, this may represent an opportunity to rationalise and upgrade the IT infrastructure. For the personnel/human resource division, this may be the opportunity to move towards a more 'learning organisation' culture. As for the staff whose knowledge and expertise is so central to the whole enterprise, and who may be expected to participate in the capture and subsequent use of any technology, they may be hoping to reduce wasted time trying to get information from other groups (it will now be online), reduce the need to handle the same queries repeatedly, and benefit from innovations elsewhere that they never hear about. All of these perspectives are interdependent. None can be examined in isolation except in an artificial, decontextualised way.

There are a number of questions, set out below, that can be asked of any proposed approach to organisational knowledge capture and re-use. These draw attention to the interdependencies between economics, technologies, work practices, and the power and responsibility that controlling knowledge repositories brings. As such, they may help to pre-empt the development of approaches which privilege any single set of concerns to the neglect of the others.

1. *What classes of knowledge/expertise are addressed by this approach?*
   There are many different classes of knowledge and expertise residing in an organisation. Relevant dimensions include tacit-explicit, procedural-declarative, tame-wicked, cognitive-cultural. Obviously, these vary widely in the extent to which they can be made (i) explicit, and (ii) formalised and

structured in digital repositories. A central challenge for organisational knowledge is to develop a better understanding of the most appropriate media for different kinds of personal and organisational knowledge/expertise. It is likely that the knowledge represented by some points in this multidimensional space cannot be formalised, without in the process invalidating it.

2. *What representational scheme is proposed, enabling what kinds of analysis and computation, with what justification?*
   What computational services over these repositories are proposed, in order to solve what kinds of problems? How does the repository reflect the changing world? Does analysis of such representations make idealised assumptions which do not hold in the real world embodiments of the knowledge/expertise being modelled? Such justification is needed when the contents of the repository relate to staff and their work practices.

3. *Who are the stakeholders? How will knowledge encoding and re-use impact their work practices?*
   Who is responsible for entering information into the repository -- a knowledge engineer; each staff employee? Does one have control over one's own area, e.g. one's 'skills profile'? Is it mandatory for all staff to keep their areas up to date; if so how is provision made for this? How does the system change inter-departmental relationships, since one's knowledge profile in the repository is now public, and therefore social? Do staff trust the system? If not, on what basis can the management?

Elsewhere [Eisenstadt 1996], we illustrate how these questions can be used to critique a system. If one takes seriously the complexity of modelling knowledge processes and products, one will approach the task of constructing organisational memory, or for that matter any KM resource, with some caution. As a heuristic approach which translates this caution into appropriate action, let us consider two related principles which can be summarised as:

KM technologies should formalise only knowledge which is stable, and sanctioned.

Stability refers to the rate of change in the domain being modelled, relative to the speed with which these changes can be detected (either by knowledge analysts, or automatically by the KM system), and the underlying knowledge representation then updated. Thus, as organisational structures change, as teams change, as individual's skills change, how will these be reflected in the KM system? This relative notion of stability implies that in principle, as advances in the flexibility of knowledge representation are made, the linkage between the model and the domain being modelled (organisational, group and individual cognitive processes) could become tighter, so that more dynamic classes of knowledge can be managed; the domain will be relatively more stable in relation to what the KM system can cope with.

Work practices become stable because they are *sanctioned* -- sustained by the relevant stakeholders. How can stable, sanctioned knowledge be identified? There is a relevant urban-planning practice to call upon here: after laying a fresh area of grass, wait for the main paths to be trodden down; it is then that one builds proper paths to

bear the heaviest traffic. In other words, in domains where consensus is unclear, formalisation should wait until the daily practices and routines of the organisation -- some of which may be too complex to predict in advance -- reveal the important, stable patterns that are in most need of support. These might include: regular transformations of knowledge from one medium to another; transfer of knowledge from one party to another; filtering functions; interdependencies between two or more schedules; checklists of action items that always need to be addressed whenever a certain event occurs.

The concept of sanctioning knowledge not only emphasises the right to know about and participate in any modelling of one's work domain, but also the right to know how one is represented in the KM system that results. This might take a number of forms, varying in the strength of the 'right to know' policy, and the technical complexity of implementing it:

- the right to know the form and content of one's entry in the knowledgebase (e.g. skills; networks; workflows; responsibilities);

- the right to know if automatic analysis or inferencing by the KM system forms the basis for management policy (appropriate questions can then be raised if there are concerns about the sufficiency of the representation or reasoning);

- the right to view, or even update knowledge stored about oneself (accessible user interfaces are required here), or to transform knowledge in one medium to another (e.g. from a video story to a textual summary, or vice-versa);

At this early stage, it is hard to predict the implications of a truly established 'knowledge economy' [Stefik 1986] operating within and between organisations. It is proposed that participatory KM design is a promising perspective to adopt: it involves all the relevant stakeholders in the complex business of modelling people's work practices and skills; it is appropriately cautious in recommending that representations be used only for stable, sanctioned knowledge processes; it emphasises the conflicts and interdependencies between the different agendas that the move towards systematic KM raises, in particular the political dimensions to controlling knowledge repositories and the legitimate concerns that this raises.

10. Conclusion

Dialogue between the AI community and other relevant disciplines such as human-computer interaction, collaborative computing, workplace ethnography and organisational learning is essential, in order to begin developing the detailed organisational scenarios of use that are at present conspicuous by their absence. From there, the first design iteration needs to be completed with empirical evidence of the success or failure of knowledge management technologies in action. Some might respond that it is too early in this field to see serious inter-disciplinary dialogue; each discipline is still struggling to formulate its own views on what The Knowledge Management Problem is. Historically, however, the evidence from more established domains of interactive system design is that the relationship between computing,

human and organisational disciplines is complex, and that each is changed through its dialogue with others. This paper has tried to illustrate how the human and computing sciences can productively engage with each other to analyse the domain, develop appropriate representations and technologies, and reason about scenarios of use from the many perspectives that interactive knowledge management technologies require.

Acknowledgements

I am grateful to Geof Bowker, Enrico Motta and Tamara Sumner for feedback on earlier drafts. The ideas in this article also benefited from discussions with delegates at the 1st International Conference on Practical Aspects of Knowledge Management, Basel 1996, and the AAAI Spring Symposium on Artificial Intelligence in Knowledge Management, Stanford 1997.

11. References

[AIKM 1997] AAAI'97 Spring Symposium on Artificial Intelligence in Knowledge Management; (Mar. 24-26), Stanford University, Palo Alto, CA 1997) [http://ksi.cpsc.ucalgary.ca/AIKM97/].

[Bannon 1996] Bannon, L. J. & Kuutti, K.: "Shifting Perspectives on Organizational Memory: From Storage to Active Remembering"; Proc. HICSS'96: 29th Hawaii International Conference on System Sciences, Hawaii (Jan. 1996), IEEE 1996)

[Bellotti 1995] Bellotti, V., Buckingham Shum, S., MacLean, A., & Hammond, N.: "Multidisciplinary Modelling In HCI Design...In Theory and In Practice"; Proc. CHI'95: Human Factors in Computing Systems, Denver, Colorado (May 7-11 1995), ACM Press: New York 1995), 146-153.

[Berlin 1993] Berlin, L. M., Jeffries, R., O'Day, V. L., Paepcke, A., & Wharton, C.: "Where Did You Put It? Issues in the Design and Use of a Group Memory"; (Ed.), Proceedings of ACM/IFIP INTERCHI'93: Conference on Human Factors in Computing Systems; 1993).

[Bieber 1997] Bieber, M., Vitali, F., Ashman, H., Balasubramanian, V., & Oinas-Kukkonen, H.: "Fourth Generation Hypermedia: Some Missing Links for the World Wide Web"; World Wide Web Usability: Special Issue of Int. J. Human-Computer Studies. Buckingham Shum, S. & McKnight, C. (Eds.), 47, 1, 31-65

[Bowers 1991] Bowers, J.: "The Politics of Formalism"; In M. Lea (Ed.), Contexts of Computer-Mediated Communication; Harvester Wheatsheaf,1991).

[Bowker, (in press)] Bowker, G. C.: "Lest We Remember: Organizational Forgetting and the Production of Knowledge"; Accounting, Management and Information Technologies, [http://alexia.lis.uiuc.edu/~bowker/paper.html].

[Brooking 1996] Brooking, A. & Motta, E.: "A Taxonomy of Intellectual Capital and a Methodology for Auditing It"; 17th Annual National Business Conference, McMaster University, Ontario, Canada (24-26 Jan.) 1996) [http://kmi.open.ac.uk/~simonb/org-knowledge/ic-paper.html].

[Brown, J. S. & Duguid, P. 1996] Brown, J. S. & Duguid, P.: "The Social Life of Documents"; First Monday, 1 [http://www.firstmonday.dk/issues/issue1/documents/].

[Buckingham Shum 1994] Buckingham Shum, S. & Hammond, N.: "Argumentation-Based Design Rationale: What Use at What Cost?"; International Journal of Human-Computer Studies, 40, 4 1994), 603-652

[Buckingham Shum 1996a] Buckingham Shum, S.: "Analyzing the Usability of a Design Rationale Notation"; In T. P. Moran & J. M. Carroll (Ed.), Design Rationale: Concepts, Techniques, and Use; Lawrence Erlbaum Associates, Hillsdale, NJ 1996).

[Buckingham Shum 1996b] Buckingham Shum, S.: "Design Argumentation as Design Rationale"; The Encyclopedia of Computer Science and Technology (Marcel Dekker Inc: NY), 35, 20 1996), 95-128

[Buckingham Shum 1997a] Buckingham Shum, S.: "Balancing Formality with Informality: User-Centred Requirements for Knowledge Management Technologies"; AAAI'97 Spring Symposium on Artificial Intelligence in Knowledge Management (Mar. 24-26 1997), Stanford University, Palo Alto, CA.
[http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-39-abstract.html].

[Buckingham Shum 1997b] Buckingham Shum, S., MacLean, A., Bellotti, V., & Hammond, N.: "Graphical Argumentation and Design Cognition"; Human-Computer Interaction, 1997, in press) [http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-25-abstract.html].

[Buckingham Shum 1997c] Buckingham Shum, S. & McKnight, C. (Eds.): "World Wide Web Usability": Special Issue of Int. J. Human-Computer Studies; 47 (1), 1-222.

[Burgess Yakemovic 1990] Burgess Yakemovic, K. C. & Conklin, J.: "Report on a Development Project Use of an Issue-Based Information System"; Proceedings of CSCW'90: Computer Supported Cooperative Work, Los Angeles, CA (7-10 Oct. 1990), ACM: New York, 105-118.

[Carroll 1991] Carroll, J. M. & Moran, T. P.: "Special Issue on Design Rationale"; Human-Computer Interaction Journal, 6 (3 & 4) 1991) 197-442.

[Carroll 1994] Carroll, J. M., Alpert, S. R., Karat, J., Deusen, M. S. V., & Rosson, M. B.: "Raison d'Etre: Capturing Design History and Rationale in Multimedia Narratives"; (Ed.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems; ACM Press, New York 1994).

[Conklin 1988] Conklin, J. & Begeman, M. L.: "gIBIS: A Hypertext Tool for Exploratory Policy Discussion"; ACM Transactions on Office Information Systems, 6, 4 1988), 303-331

[Conklin 1989] Conklin, J.: "Design Rationale and Maintainability"; Proceedings 22nd Hawaii International Conference on System Science, Los Alamitos: IEEE Computer Society Press, Vol. 2 1989), pp. 533-539.

[Conklin 1991] Conklin, J. & Burgess Yakemovic, K. C.: "A Process-Oriented Approach to Design Rationale"; Human-Computer Interaction, 6, 3&4 1991), 357-391 [Reprinted in: T.P. Moran and J.M. Carroll (Eds.) *Design Rationale: Concepts, Techniques, and Use,* (pp. 393-427). Hillsdale, NJ: Lawrence Erlbaum Associates 1996].

[Conklin 1993] Conklin, J. & Yourdon, E.: "Groupware for the New Organization"; American Programmer, Sept. 1993, 3-8.

[Conklin 1996] Conklin, J.: "Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy"; Group Decision Support Systems (1996). [http://www.gdss.com/DOM.htm].

[Davenport 1996] Davenport, T.: "Some Principles of Knowledge Management"; Graduate School of Business, University of Texas at Austin 1996). [http://knowman.bus.utexas.edu/pubs/kmprin.htm].

[Davis 1993] Davis, R., Shrobe, H., & Szolovits, P.: "What Is A Knowledge Representation?"; AI Magazine, 14, 1 1993), 17-33 [http://medg.lcs.mit.edu/ftp/psz/k-rep.html].

[Eisenstadt 1996] Eisenstadt, M., Buckingham Shum, S., & Freeman, A.: "KMi Stadium: Web-based Audio/Visual Interaction as Reusable Organisational Expertise"; Workshop on Knowledge Media for Improving Organisational Expertise, 1st International Conference on Practical Aspects of Knowledge Management, Basel, Switzerland (30-31 Oct. 1996) [http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-31-abstract.html].

[Eldridge 1992] Eldridge, M., Lamming, M., & Flynn, M.: "Does a Video Diary Help Recall?"; Proceedings of the HCI'92 Conference on People and Computers VII, Cambridge University Press, Cambridge. 1992).

[Euzenat 1996] Euzenat, J.: "Corporate Memory Through Cooperative Creation of Knowledge Bases and Hyperdocuments"; Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems, Banff, Canada (November 1996).
[ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/22euzenat.ps.Z].

[Fischer 1991] Fischer, G., Lemke, A. C., McCall, R., & Morch, A. I.: "Making Argumentation Serve Design"; Human-Computer Interaction, 6, 3&4 1991), 393-419 [Reprinted in: T.P. Moran and J.M. Carroll (Eds.) *Design Rationale: Concepts, Techniques, and Use,* (pp. 267-293). Hillsdale, NJ: Lawrence Erlbaum Associates 1996].

[Fischer 1995] Fischer, G., Linstaedt, S., Ostwald, J., Stolze, M., Sumner, T., & Zimmerman, B.: "From Domain Modeling to Collaborative Domain Construction"; DIS'95: 1st ACM Conference on Designing Interactive Systems, Michigan, ACM Press: New York 1995).

[Garcia 1992] Garcia, A. C. B. & Howard, H. C.: "Acquiring Design Knowledge Through Design Decision Justification"; Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 6, 1 1992), 59-71

[GDSS 1996] "QuestMap". Group Decision Support Systems, 1000 Thomas Jefferson Street, NW, Suite 100, Washington, DC 20007, U.S.A. [http://www.gdss.com/OM.htm].

[Grudin 1996] Grudin, J.: "Evaluating Opportunities for Design Capture"; In T. P. Moran & J. M. Carroll (Ed.), Design Rationale: Concepts, Techniques, and Use; Lawrence Erlbaum Associates, Hillsdale, NJ 1996).

[Jarczyk 1992] Jarczyk, A., Loffler, P., & Shipman, F.: "Design Rationale for Software Engineering: A Survey"; Proceedings of 25th Annual Hawaii International Conference on System Sciences (8-10 January 1992)

[Kidd 1994] Kidd, A.: "The Marks are on the Knowledge Worker"; Proc. ACM CHI'94: Human Factors in Computing Systems, Boston, Mass (24-28 April 1994), ACM Press: New York, 186-191.

[Kremer 1996] Kremer, R.: "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, Canada (November 1996). [ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/77kremer.ps.Z].

[LaLiberte 1995] LaLiberte, D.: "HyperNews"; Report Number, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign 1995). [http://union.ncsa.uiuc.edu/HyperNews/get/hypernews.html].

[Lee 1990] Lee, J.: "SIBYL: A Tool for Managing Group Design Rationale"; Computer Supported Cooperative Work, Los Angeles, CA, ACM Press: New York 1990), 79-92.

[Lee 1991a] Lee, J.: "Extending the Potts and Bruns Model for Recording Design Rationale"; Proceedings of the 13th International Conference on Software Engineering, Austin, Texas, New York: ACM 1991), 114-125.

[Lee 1991b] Lee, J. & Lai, K.: "What's in Design Rationale?"; Human-Computer Interaction, 6, 3&4 1991), 251-280 [Reprinted in: T.P. Moran and J.M. Carroll (Eds.) *Design Rationale: Concepts, Techniques, and Use,* (pp. 21-51). Hillsdale, NJ: Lawrence Erlbaum Associates 1996].

[Lotus 1996] Lotus: "ScreenCam"; Lotus Development Corporation 1996). [http://www.lotus.com/screencam/]

[MacLean 1989] MacLean, A., Young, R. M., & Moran, T.: "Design Rationale: The Argument Behind the Artifact"; Proceedings of CHI'89: Human Factors in Computing Systems, ACM: New York 1989), 247-252.

[MacLean 1991] MacLean, A., Young, R. M., Bellotti, V., & Moran, T.: "Questions, Options, and Criteria: Elements of Design Space Analysis"; Human-Computer Interaction, 6, 3 & 4 1991), 201-250 [Reprinted 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 1992-94] MacLean, A., Buckingham Shum, S., & Bellotti, V.: "Design Rationale Tutorial"; Tutorials at HCI'92-94: British Computer Society HCI Conferences, U.K. 1992-94)

[Moran 1996] Moran, T. P. & Carroll, J. M. "Design Rationale: Concepts, Techniques and Use". (Eds.); Lawrence Erlbaum Associates, Hillsdale, NJ 1996). [ISBN 0-8058-1566-X].

[Moran 1997] Moran, T. P., Palen, L., Harrison, S., Chiu, P., Kimber, D., Minneman, S., van Melle, W., & Zellweger, P.: ""I'll Get That Off the Audio": A Case Study of Salvaging Multimedia Meeting Records"; Proc. CHI 97: Human Factors in Computing Systems, Atlanta, USA (March 22-27 1997), ACM: New York [http://www.acm.org/sigchi/chi97/proceedings/paper/tpm.htm].

[Multicosm 1996a] "Microcosm". [http://www.webcosm.com/].

[Multicosm 1996b] "Webcosm"; [http://www.webcosm.com/].

[Newman 1991] Newman, S. & Marshall, C.: "Pushing Toulmin Too Far: Learning from an Argument Representation Scheme"; Technical Report Xerox Palo Alto Research Center 1991).

[Orr 1986] Orr, J.: "Narratives at work: Story telling as cooperative diagnostic activity"; Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '86), Austin, Texas, ACM Press (Year), 62-72.

[Potts 1988] Potts, C. & Bruns, G.: "Recording the Reasons for Design Decisions"; Proceedings of 10th International Conference on Software Engineering, Washington, DC: IEEE Computer Press 1988), 418-427.

[Potts 1994] Potts, C., Takahashi, K., & Anton, A.: "Inquiry-Based Requirements Analysis"; IEEE Software, March 1994), 21-32

[Ramesh 1993] Ramesh, B.: "Supporting Systems Development by Capturing Deliberations During Requirements Engineering"; IEE Transactions on Software Engineering, 18, 6 1993), 498-510

[Rittel 1972] Rittel, H. W. J.: "Second Generation Design Methods"; Interview in: Design Methods Group 5th Anniversary Report: DMG Occasional Paper, 1, 5-10. Reprinted in: Developments in Design Methodology, N. Cross (Ed.) 1984, pp. 317-327, J. Wiley & Sons: Chichester

[Rittel 1973] Rittel, H. W. J. & Webber, M. M.: "Dilemmas in a General Theory of Planning"; Policy Sciences, 4 1973), 155-169

[Selvin 1997] Selvin, A.: "An Issue-Based Hypertext Approach to Collaborative Process Redesign"; Report Number, NYNEX Science &Technology Technical Memorandum, White Plains, NY 1997).

[Shum 1993] Shum, S.: "QOC Design Rationale Retrieval: A Cognitive Task Analysis & Design Implications"; Technical Report EPC-1993-105, Rank Xerox Cambridge EuroPARC 1993).

[Smolensky 1987] Smolensky, P., Bell, B., Fox, B., King, R., & Lewis, C.: "Constraint-Based Hypertext for Argumentation"; (Ed.), ACM Hypertext'87 Proceedings; ACM: New York,1987).

[Stefik 1986] Stefik, M.: "The Next Knowledge Medium"; AI Magazine, 7, 1 1986), 34-46

[Stutt 1995] Stutt, A. & Motta, E.: "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, Canada (26 Feb-3 Mar 1995).

[Suchman 1993] Suchman, L.: "Do Categories have Politics? The Language/Action Perspective Reconsidered"; 3rd European Conference on Computer-Supported Cooperative Work, Milan, Italy (13-17 Sept. 1993), Kluwer Academic Publishers, 1-14.

[Sumner 1996] Sumner, T. & Buckingham Shum, S.: "Open Peer Review & Argumentation: Loosening the Paper Chains on Journals"; Ariadne (Bi-Monthly Magazine of the UK Electronic Libraries Programme), Issue 5 (Sept.) 1996)
[http://www.ukoln.ac.uk/ariadne/issue5/jime/].

[van Heijst 1996] van Heijst, G., van der Spek, R., & Kruizinga, E.: "Organizing Corporate Memories"; Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems, Banff, Canada (November 1996).
[ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/25vanheijst.ps.Z].

[VanLehn 1985] VanLehn, K.: "Theory Reform Caused by an Argumentation Tool"; Technical Report, Xerox Palo Alto Research Center 1985).

[Vanwelkenhuysen 1994] Vanwelkenhuysen, J. & Mizoguchi, R.: "Maintaining the Workplace Context in a Knowledge Level Analysis"; JKAW'94: Proceedings 3rd Japanese Knowledge Acquisition for Knowledge-Based Systems Workshop, Hatoyama, Japan 1994), 33-47.

[Vanwelkenhuysen 1995] Vanwelkenhuysen, J.: "Embedding Non-Functional Requirements Analyses in Conceptual Knowledge Systems Designs"; Proceedings 9th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada (26 Feb - 3 Mar. 1995), 45.1-45.15.