Seventh International Workshop on Hypertext Functionality: Organizational Memory Systems & HTF (Helsinki, December 12-13, 1998) University of Oulu, Dept. Information Processing Science, Research Papers Series A29, Nov. 1999. ISSN 0786-8413

Supporting Granular Reuse of Knowledge Elements in an Organizational Memory System

Albert M. Selvin
Senior Member of Technical Staff
Verizon eBusiness
400 Westchester Ave.
White Plains, NY 10604 USA
(914) 644-2156
albert.m.selvin@verizon.com

Abstract
This paper presents an approach to creating and populating an HTF-based organizational memory resource in such a manner as to maximize both the efficiency of the population activities and the likelihood that elements of the OM resource will be retrieved and reused in multiple contexts. The approach is based on more than five years of experience in the creation and maintenance of such resources in an industry setting. The paper first presents our conceptualization of organizational memory and how we have approached supporting such memory, including an argument that the grain-size of elements that make up an organizational memory resource must be smaller than the document level. Next, we provide a brief introduction to the approach and present three mechanisms that have been used in multiple project settings, with real-life examples of each. Finally, we present our future research plans.

Introduction
This paper presents an approach to creating and populating an HTF-based organizational memory resource in such a manner as to maximize both the efficiency of the population activities and the likelihood that elements of the OM resource will be retrieved and reused in multiple contexts. The approach is based on several years of experience in the creation and maintenance of such resources in an industry setting (1, 2), specifically a systems development and process redesign organization located within a large telecommunications firm. The HTF aspects of the approach are currently based on leveraging and extending the hypertext functionality in an off-the-shelf commercial tool, which was reported at HTFII in 1996 (3).

The paper first presents our conceptualization of organizational memory and how we have approached supporting such memory, including an argument that the grainsize of elements that make up an organizational memory resource must be smaller than the document level. Next, we provide a brief overview of the approach and present three mechanisms that have been used in multiple project settings, with real-life examples of each. Finally, we present our future research plans.

Supporting Organizational Memory
Organizational memory (OM), like individual memory, is made up of ideas and of associations between ideas, or "knowledge elements". OM itself is a concept or phenomenon; it’s something that occurs and exists whether or not explicit attention is paid to it, or whether there is software support for it. It is important to make the distinction between OM as a phenomenon, and a carrier or resource for OM, which could be a computer system. Such an OM system (OMS) is best thought

of as a resource for constructive acts or events (4), moments where someone is actively making something happen or bringing something into being (5). OMSs provide a physical representation of (some) aspects of organizational and interpersonal communication, and make the representation of those aspects available for subsequent manipulation and use in the context of ongoing activities (6). Such systems are primarily of use when something is "remembered" in the service of doing something else. In other words, an OM resource is best thought of not as an archive of data, but rather as a resource to be brought to bear in the service of a current activity. This paper reports on a particular embodiment of an OMS that provides such support.

In our OMS work, we have been concerned with the *quality* of the constructive event, and in particular, the degree to which an OMS can support the quality of the event. We define an instance of an OMS-supported constructive event as a moment when a piece of information or knowledge stored in the system (a node, fact, idea, etc.) is associated with one or more other pieces of information in a new way, or in a new context. Given this criteria, a “good” OMS is one that makes pieces of the represented OM – what we call “knowledge elements” -- available for such new associations in as easy a manner as possible, and in such a manner that maximizes the likelihood that such associations will take place.

To meet this criteria, the carriers of knowledge elements in an OMS need to be as "open" as possible. That is, knowledge elements need to be "exposed" to inquiry and available for new associations. They need to be easily delinked from their surrounding context to be associated with something new, but without forgetting that original surrounding context. Several researchers have observed that there is little information available on re-use of information gathered in an OM resource (7, 8). It may be that the "health" (9) or "viability" of an OMS could be defined as the measure of how much and how often knowledge elements get reused (reappear) in new contexts.

Further, we have found that documents, in and of themselves, are insufficient carriers/repositories for OMSs. While interrelated documents (a la Lotus Notes or the WWW) can play an important role in an OMS, they are inefficient for OM support for two reasons.

First, they are expensive to create, maintain, and interrelate. Organizations that have adopted document-based knowledge management systems find that they require a large staff of librarians and editors to categorize and maintain relevance of the systems’ content. (10) Second, and more of issue to this paper, is that document-based systems are not open enough (given the definition of openness above). A document is simply too large of a grainsize. While an OMS might assist users in finding documents that match a stated area of interest, it is still up to the user to examine the document, parse its structure, locate particular items of interest, interpret them in their surrounding context, then manually extract them to relate them to their current activity.

Reusing knowledge elements from one physical context (the original document) to another (e.g. a new document) thus requires a good deal of user effort. To speak metaphorically, documents embed the atoms and molecules of interest within complicated compounds, making it an effort to find them and relate them to other atoms and molecules in the compound one is trying to create in the context of the present constructive event. To put it another way, by linking elements together in linear paragraphs, sections, and chapters, documents provide a thick forest that, in effect, obscures the trees. In addition, when an instance of constructive use of OM happens, people usually suddenly think of an old idea or fact in relation to the new context; they don't usually think of a whole huge structured ensemble of ideas (i.e. a document) in relation to the new context. Therefore, an OMS needs to provide support at a more granular level than linked documents.

Once a knowledge element of relevance to a user’s current constructive event has been discovered in an OMS, a user must be able to incorporate it into their current physical context (typically by copying and pasting). Such mechanisms should allow the user to create a link back to the specific original context, ideally without requiring additional manual work (such as writing a URL with specific location data, creating a bookmark, etc.). Open carriers of knowledge elements for OMSs provide access to surrounding structure and context in both original and current contexts, as well as making the relations between elements themselves explicit.

To aid in the above, an OMS should allow for the transformation of one representation of a knowledge element and its context to another formalism (11) without losing the semantic meaning implied by the original structure. For example, elements whose original physical representation were as the contents of adjacent cells in a spreadsheet should be "harvestable" as linked nodes in a concept map-type representation, in such a way so that they are available for subsequent linking in other concept maps in the same OMS. An OMS should provide simple, consistent, coherent ways of harvesting and recombining information that make it easy to input, easy to classify (formalize), easy to validate, reliable and accurate, unbundled and granular, and easy to combine and recombine without taking too much time and effort.

Project Compendium as an OMS

The Project Compendium approach provides several strategies for such open carriers, developed out of experience with creating knowledge repositories (defined as collections of knowledge elements, structures, and relationships) on the fly in the service of particular project activities, then reusing the elements of these repositories in multiple ways. A unique aspect of Project Compendium is the ability to support rapid and flexible *capturing* of knowledge elements in an OMS as well as the demonstrated potential for multiple reuse and multiple leveraging of the same content (elements) with minimal added "cost" (effort).

Project Compendium's HTF and other aspects have been extensively described elsewhere (12, 13). This section draws on the framework developed by Zimmerman and Selvin (14) to summarize the input and formalization mechanisms Project Compendium uses to create and maintain an open and reusable OMS.

Input

In Project Compendium, data input into the system can happen based on several strategies:

- Facilitated collaborative sessions. Meetings are held in which the system is used as a shared display and a facilitator captures discussion in full view of participants. Validation of the content happens through the presence of the subject matter experts at the moment of input.
- Individuals working on their own, either entering new knowledge elements (in the form of individual nodes) and associations between elements (in the form of associative and transclusive links), or editing and editing previously entered nodes and associations
- Harvesting data from pre-existing documents that have had their own validation cycle

Formalization

Project Compendium employs a variety of incremental formalization mechanisms (15), such as applying codes and keywords to knowledge elements in various tools (either the hypertext authoring environment itself or in other tools, such as MS-Word); hyperlinking of elements; and associating elements according to predefined templates. The templates correspond to either project-specific or a priori schema, such as the World Modeling framework developed by Maarten Sierhuis and Rob van der Spek.

Project Compendium formalization "costs" include the time and effort involved in setting up & maintaining categories and keywords, the change in work practices to add the categories and keywords to elements and to participate in facilitated sessions, the cost of facilitator involvement in face-to-face collaborative sessions, a commitment to consistently enter information and apply the formalization mechanisms, and the opportunity cost of lost data if formalization isn't applied. Benefits include: creating the ability to cross-reference and interlink project information elements, an evolved group vocabulary, and task definitions, making information elements available in multiple contexts, providing opportunities for knowledge discovery (users may find associations between elements they were not aware of in the process of formalization), the ability to generate up-to-the-minute reports of status of various aspects of the project, and reusability of data for other project teams.

Granular Reuse Mechanisms and Examples

This section presents three of the mechanisms for granular reuse that have been developed as part of the Project Compendium approach, and shows examples of how the mechanisms have been used in actual project settings. In each of the examples, the original context for populating knowledge elements and relationships into the OMS is presented, followed by a description of the subsequent context of reuse and an explanation of the "constructive event" that caused the subsequent manipulation of the original elements.

Mechanisms covered here include:

- Mechanism 1: harvesting ideas from source documents for other contexts
- Mechanism 2: extending analysis elements from one project team to another
- Mechanism 3: reuse over long term within same project team

Mechanism 1: Harvesting Ideas from Source Documents for Other Contexts

In this mechanism, users employ a software tool to extract knowledge elements and relationships from source documents and represent them within the OMS as typed nodes and links, available for subsequent linking and reuse in other contexts.

For example, the following source document was originally sent as an email as comments from an internal customer on a demonstration of a prototype:

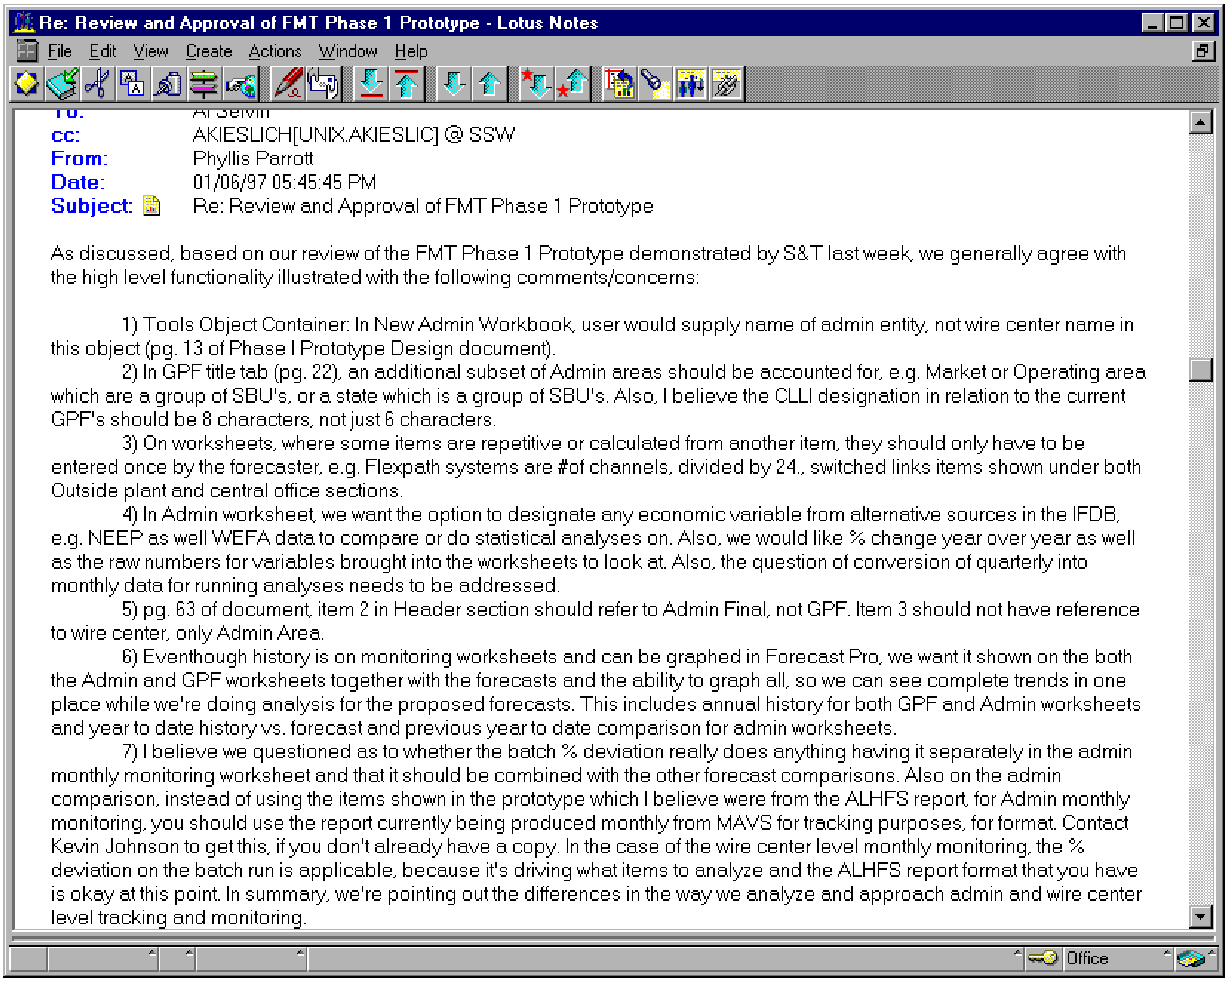

The project leader of the system development effort used a Project Compendium software tool to import and represent each paragraph in the email as a node in a map in the OMS:

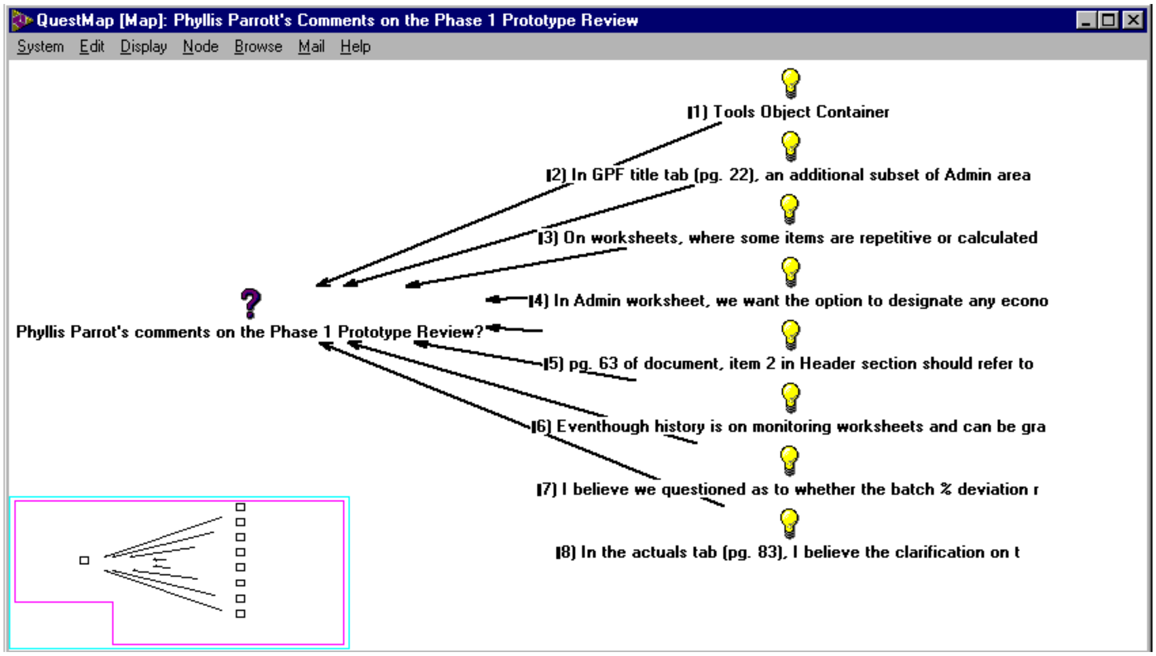

Subsequently, individual items from this email were reused in various contexts. In the following example, a member of the project team linked one of the nodes (representing a particular point from the original email) into a discussion of forthcoming work to be done on the project:

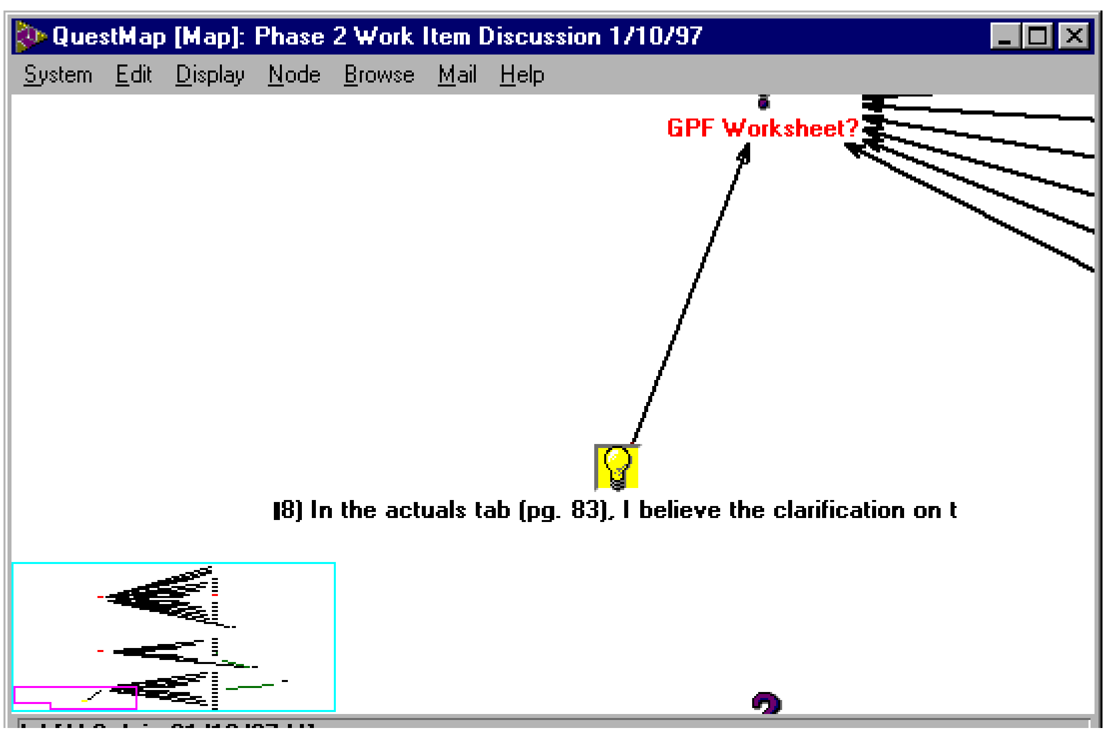

In addition, annotations added to the node at various times indicate that different members of the project team had considered the information in light of their concerns at the moment:

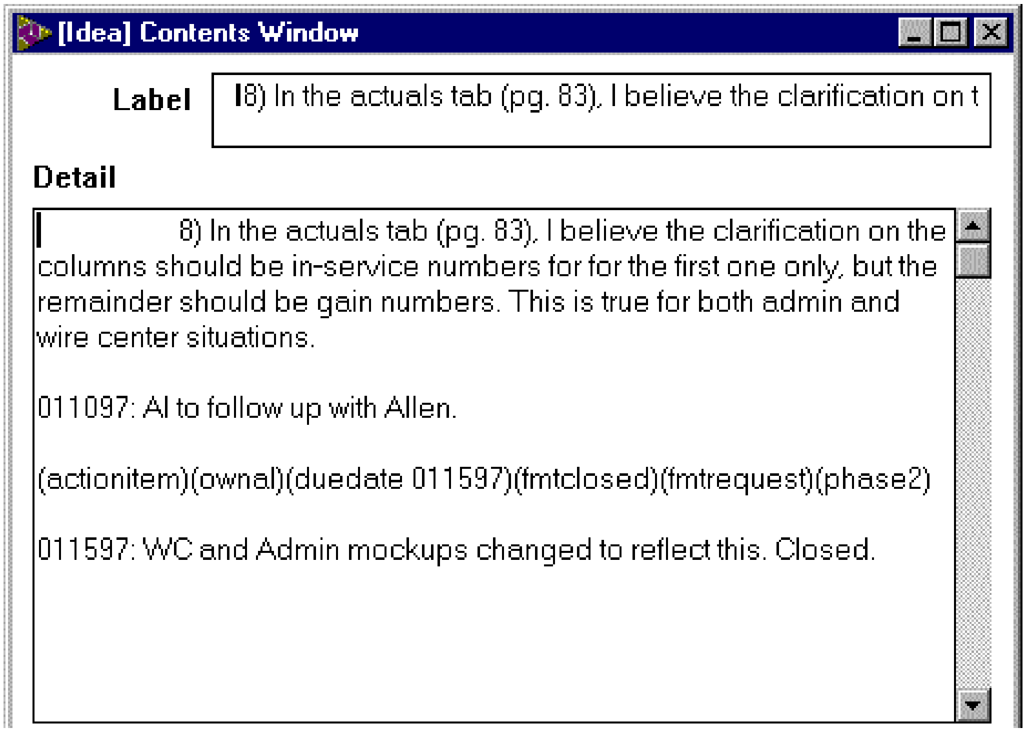

Mechanism 2: Extending Analysis Elements from One Project Team to Another

In this mechanism, following the Project Compendium formalization schemes outlined above allows preservation of knowledge elements in such a way as to allow delinking from the original context and reuse in a completely different context by different groups of people.

The following example shows knowledge elements originally developed in a 1995 effort to develop requirements for future residential broadband planning and engineering functions and systems. In this case, input of the original elements was accomplished via site visits and interviews followed by creation of formal analysis models. The models were initially created by two analysts working together, then individually by a single (new) user. As part of the 1995 project, the material was later retrieved and displayed (and further modified) at face-to-face validation meetings held with the original subject matter experts.

The map below shows a grouping of organization (role) models as performed by the original 1995 analyst team:

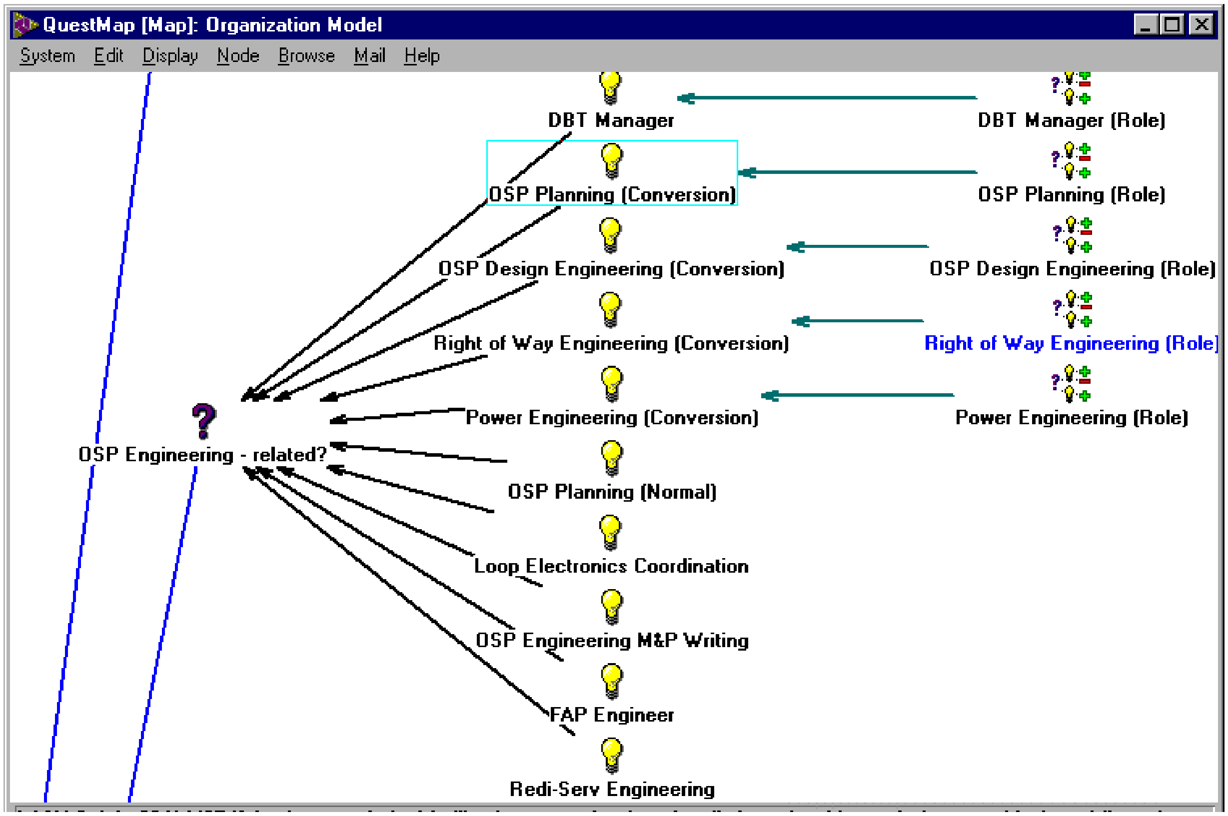

As part of the analysis effort described above, the analysts identified computer systems and other resources used by various roles within the organization. The map below shows an excerpt from one of these maps, showing which systems (LFACS, ECRIS, etc.) were used by the role named "OSP Planning":

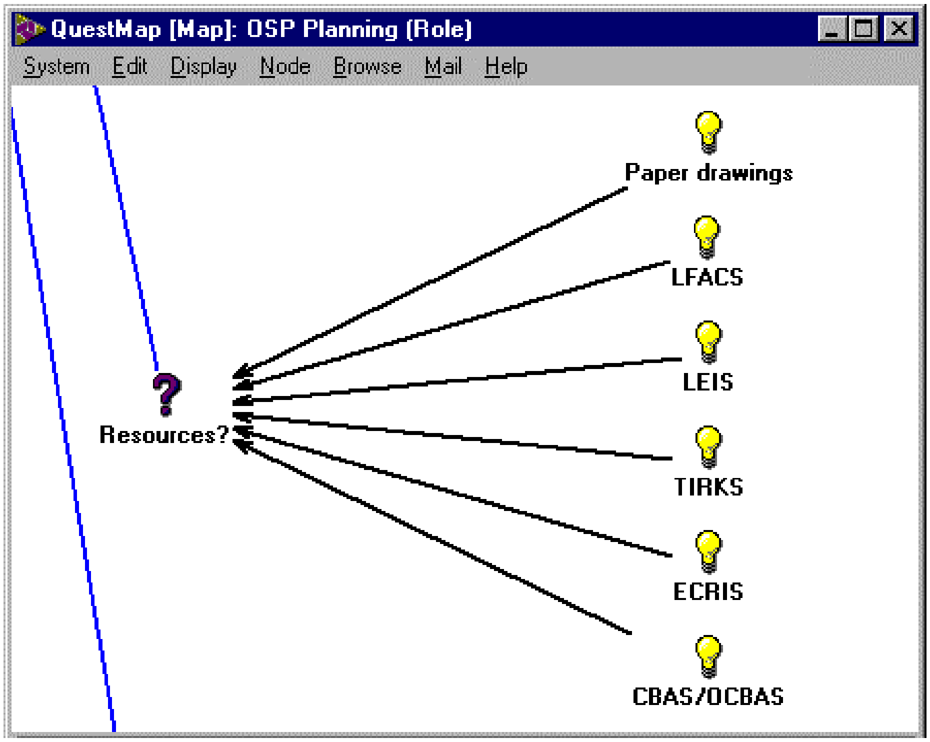

In 1998, a newly formed engineering systems organization with 150 systems in its portfolio set up several committees to get a handle on all the systems from various perspectives (financial, database architecture, participatory design). Participants in the committees were spread over many groups, cities, and work locations. The new committees took the nodes representing engineering systems and roles/functions from the 1995 database and applied their own templates and tagging schemas (formalization mechanisms) to them.

In the map below, members of one of the committees (Data Architecture) has begun doing an analysis of the data architecture aspects of one of the systems originally identified in the 1995 analysis (ECRIS). They have applied a template of issues (stewarded data entities, reference data entities, source systems, and system interfaces) of concern to them (the same template will be applied to all other systems in the group of 150):

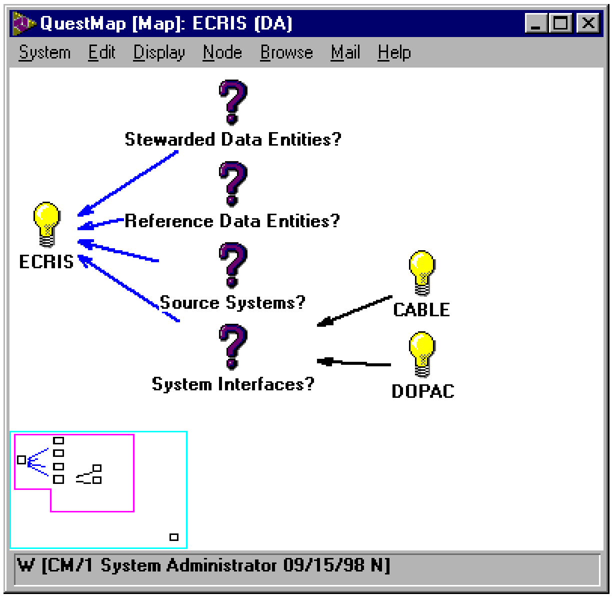

Another committee (Participatory Design) has done the same, applying their template of issues to the same system:

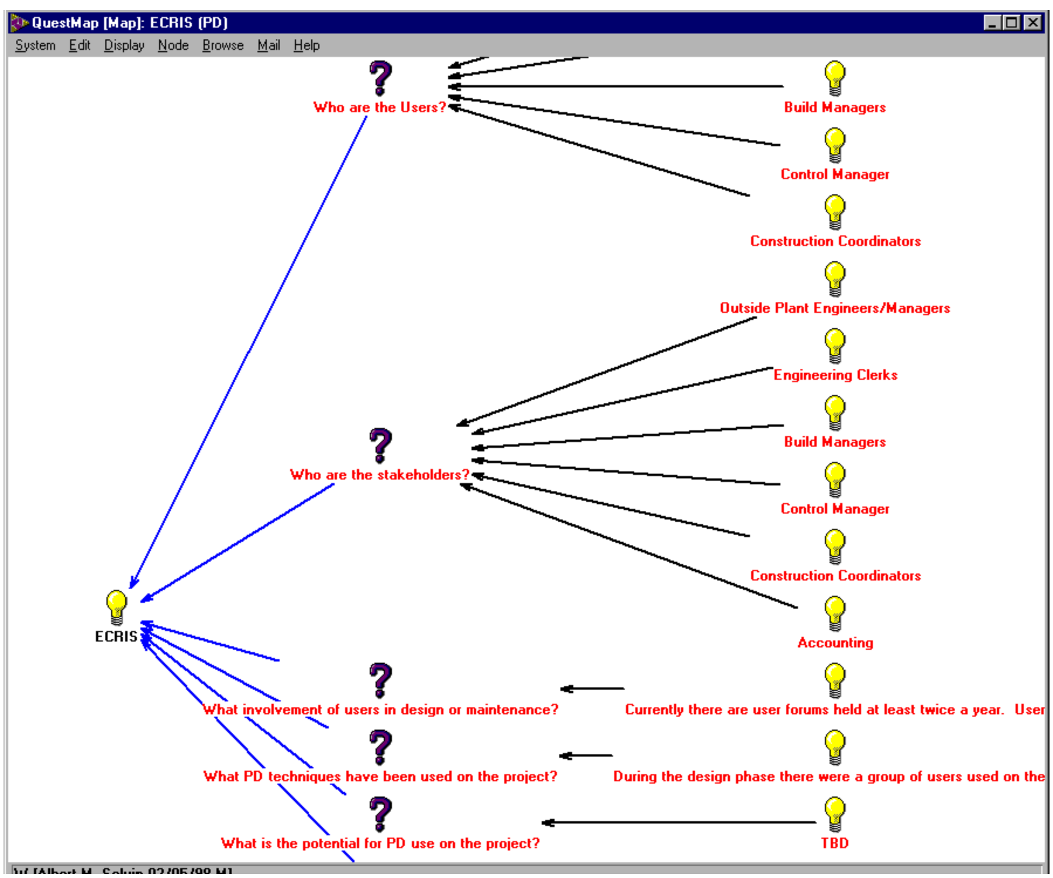

Mechanism 3: Reuse Within the Same Long-Term Project

In this mechanism, a particular team is able to reuse knowledge elements and relations created early in its lifecycle in subsequent contexts for different purposes.

The following example shows knowledge elements originally developed in late 1996-early 1997 as part of a requirements analysis effort for a systems development project. The first map in the series shows a high-level overview or collection of maps containing the requirements specifications for the various modules of the system:

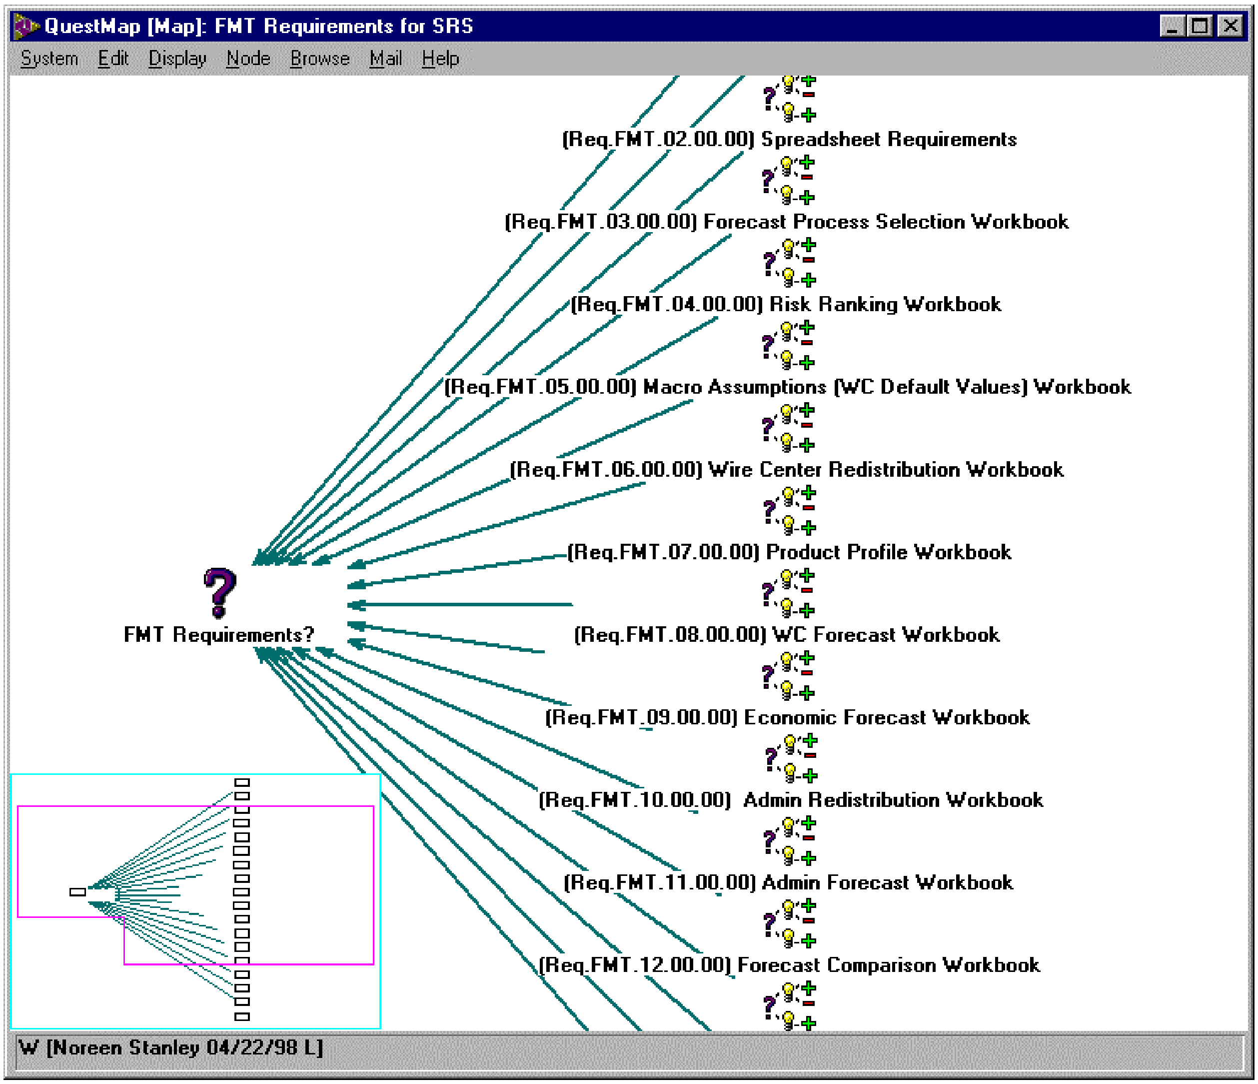

Each of the map nodes could be opened to reveal the formal requirements for that module. In the map below, formal requirements for the Admin Redistribution Workbook module are gathered for presentation. Each node could itself be opened to display the details of the particular requirement:

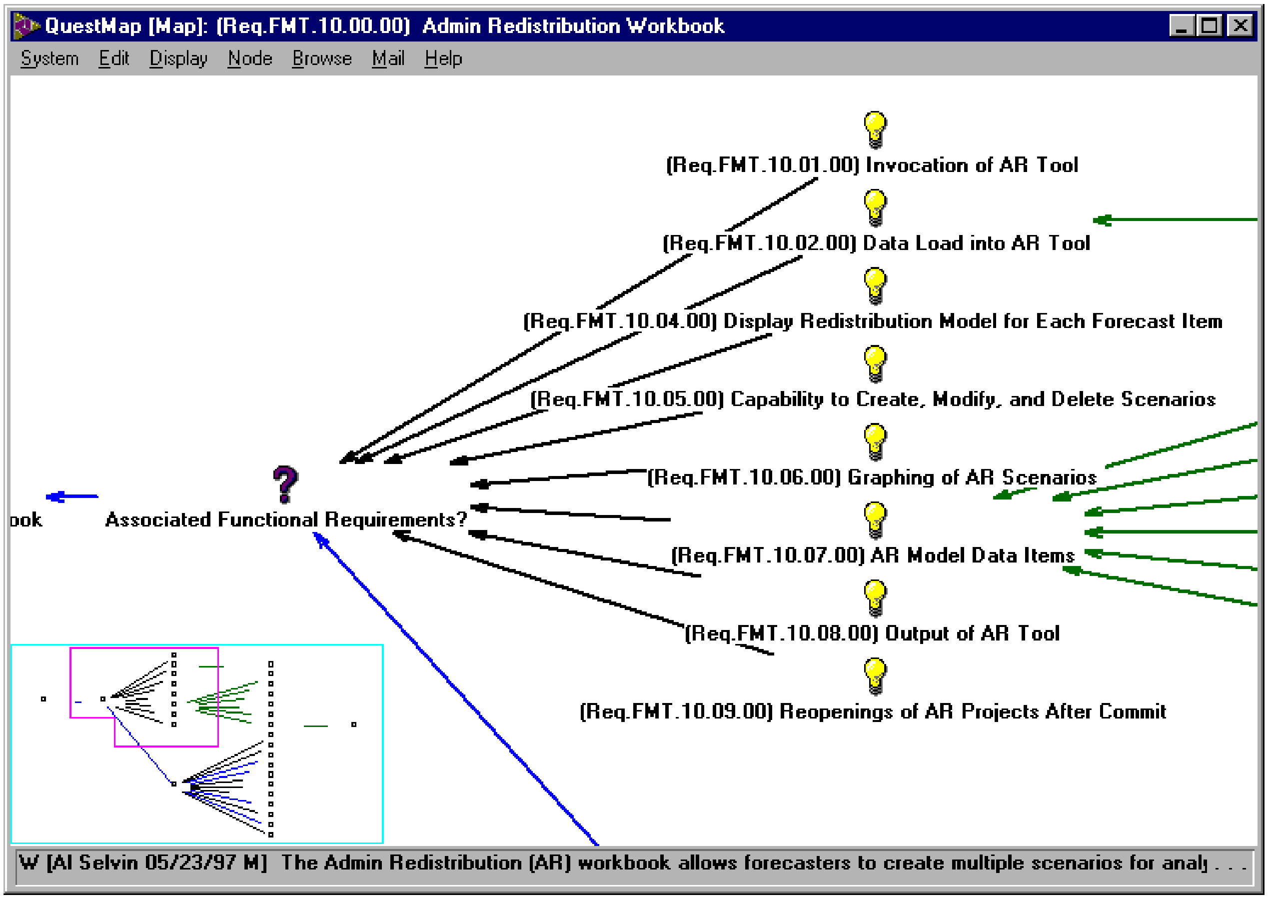

The project team used a Project Compendium software tool to generate a single MS-Word document containing all the requirements information from the various individual maps, for review and approval by the project's customer. Subsequent to this, members of the project team added new items (mostly in an informal manner) to the individual maps as new requirements or issues emerged over the next months. In the map below, new requirements and issues pertaining to the Admin Redistribution Workbook module have been placed in the map at various times:

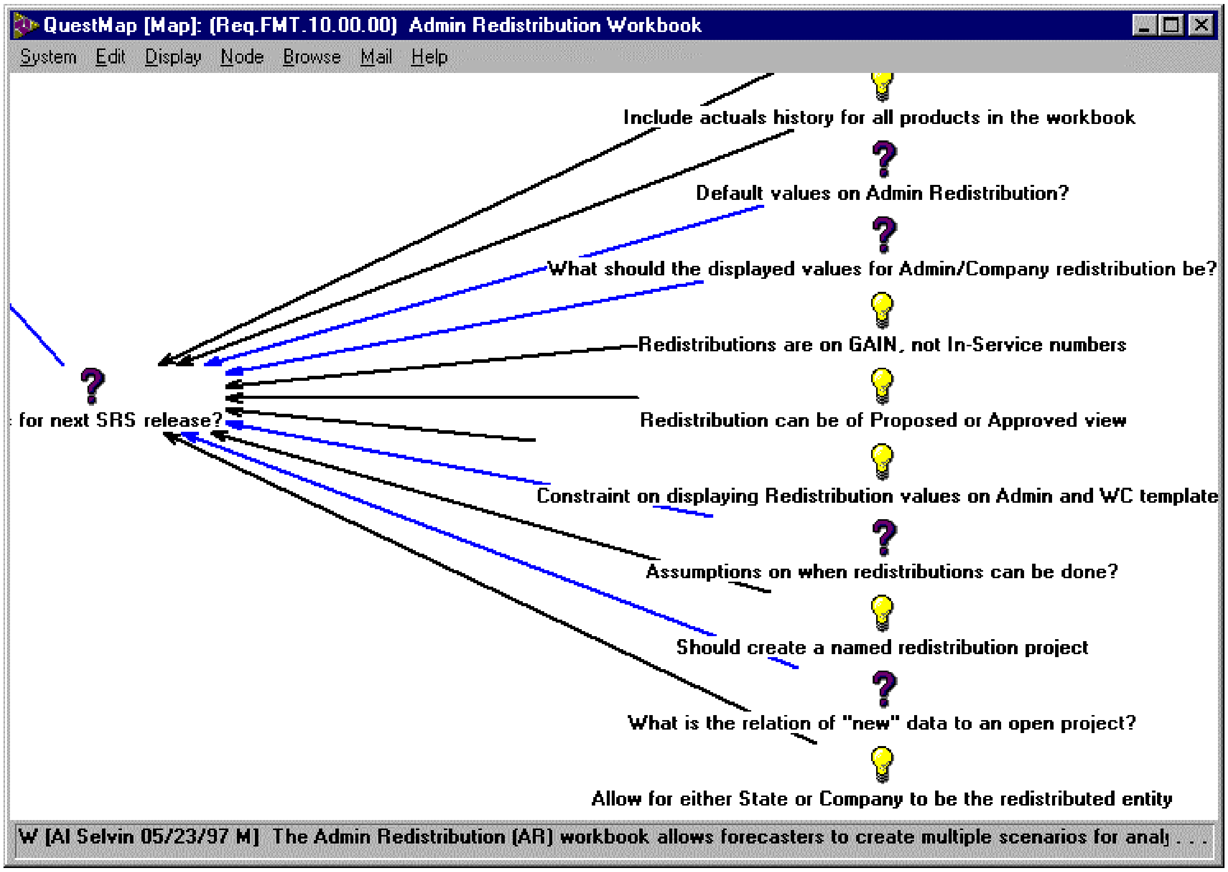

Eighteen months after the original requirements had been developed, the project leader gathered the map nodes representing the discrete groups of requirements on a new map. This was used in an interactive session with the project's "Core Team" of users and customers to prioritize development for the various modules of the system.

The map below shows the result of the work done in that session. While the map was displayed in front of the group with an LCD projector, the project leader manipulated the horizontal order of the icons representing the various modules in response to the discussion in the room. As part of the prioritization exercise, individual map icons were opened to display both the original requirements and the subsequent additions and modifications mentioned above.

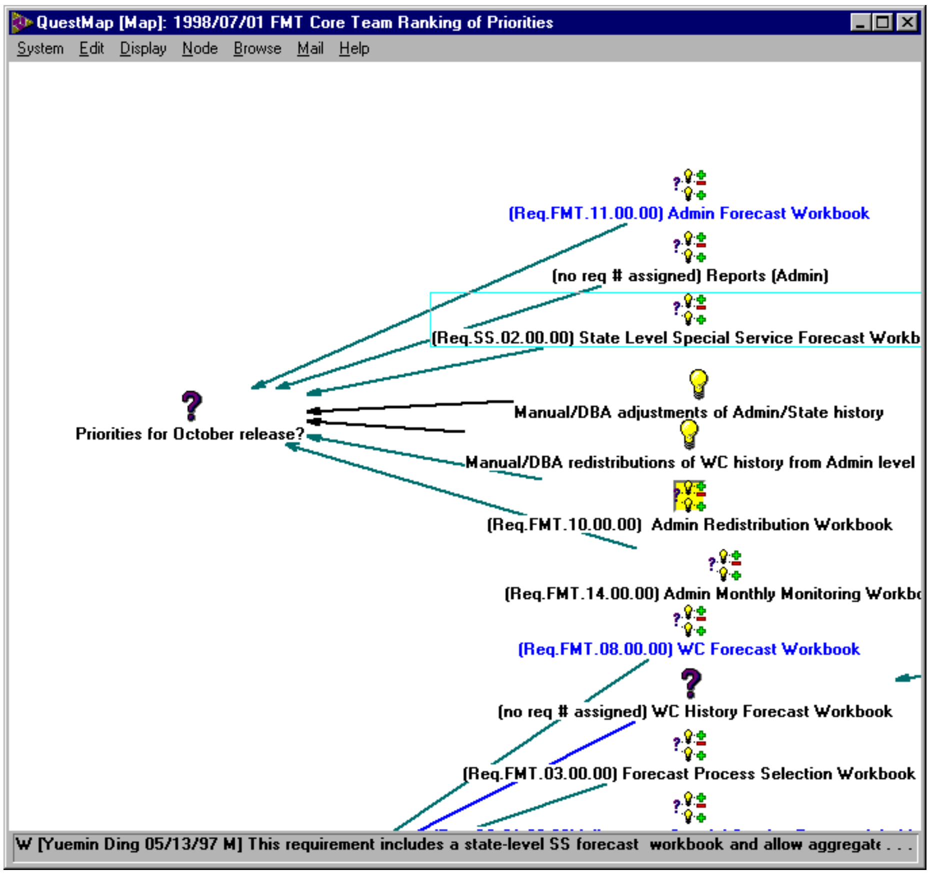

The window below shows a partial list of all the different contexts (views) that the map node representing the "Admin Redistribution Workbook" requirements collection has appeared in over the two years of the project:

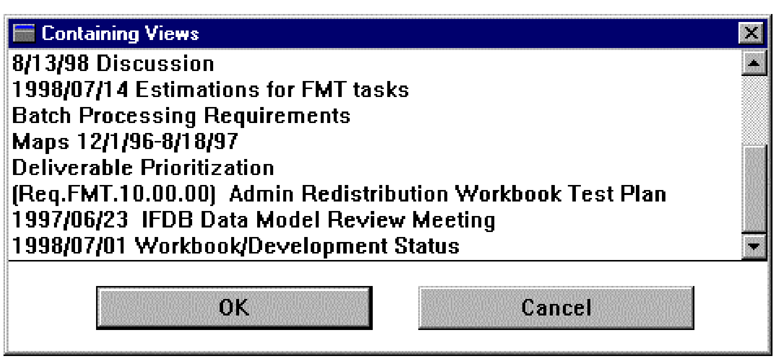

Discussion and Future Work

The examples above provide a small sample of the ways in which granular reuse of knowledge elements in an OMS can be supported. With the proper tools and methods, a deeply interlinked OM resource can offer a great deal of help to individuals and teams engaged in constructive events. As the examples above show, some of the ways an OMS can help are:

- by providing starting points for new work that bring to bear knowledge elements and relationships developed in an earlier context, by different people

- by providing ways to share similar knowledge elements among multiple groups working concurrently

- by providing objects created for one purpose that can be manipulated for other purposes, while still containing all their original content and relationships to other objects

To date, use of Project Compendium tools and methods at Bell Atlantic have yielded a sizable amount of such granular reuse. It is too early to generalize from our experience to make many wider claims about the sustainability, viability, and usability of the OM resources we have developed. In future work, we will look more deeply at some of the databases that have been (and are continuing to be) built in an attempt to understand the interplay of situated project team use of the OMS, the amount, type, and purpose of the reuse that has occurred, and the tools and methods employed. Future research will provide an analysis of how often such reuse has happened and a further breakdown of types of reuse (node level, structure level, etc.), and look more deeply into the question of OMS "health" mentioned above. We are also interested in providing a deeper analysis of the success factors for the types of granular reuse covered in this paper (why was reuse successful? What purposes did it serve?).

One observation we can make now, though, is that a critical factor in the success of the approach so far has been the role of social relationships in making people aware that these resources exist and helping other teams to make use of OMS material in a new context. More research is needed on such social network issues, especially as they relate to skill development, resource awareness, implementation and promotion strategies, and tool/method design.

Finally, as the examples in this paper indicate, Project Compendium work to date has been done with using GDSS, Inc.'s QuestMap™ product as the hypertext authoring tool and data repository, along with a number of small helper applications written at Bell Atlantic. However, driven in part by our desire to analyze and understand the types of reuse and other data relationships in our OMS databases in ways that QuestMap does not currently support, we are now building our own Java-based toolset to serve as the collaborative hypermedia environment. We will report on that work in a future paper.

Acknowledgements

The author wishes to thank Maarten Sierhuis (NASA Ames Research Center), Jeff Conklin (GDSS Inc.), and Simon Buckingham Shum (Open University) for the many conversations and collaborative efforts that gave rise to the ideas in this paper.

Citations

1. A. Selvin and M. Sierhuis, "Experiences with Conversational Modeling: Building Models Through Collaboration," *NYNEX Science & Technology Technical Memorandum TM96-0045*, White Plains, NY, 1996.

2. A. Selvin, "Supporting Collaborative Analysis and Design with Hypertext Functionality," in *Journal of Digital Information,* Special Issue on Hypertext Functionality, forthcoming.

3. A. Selvin, "Leveraging Existing Hypertext Functionality to Create a Customized Environment for Team Analysis," in *Proceedings of the Second International Workshop on Incorporating Hypertext Functionality Into Software Systems*, March 1996. Also at http://space.njit.edu:5080/HTFII/Selvin.html.

4. L. Bannon, K. Kuuti, "Shifting Perspectives on Organizational Memory: From Storage to Active Remembering," in *Proceedings IEEE HICSS '96*, Hawaii, Jan. 1996.

5. S. Buckingham Shum, "Negotiating the Construction and Reconstruction of Organisational Memories," *Journal of Universal Computer Science* 3-8, 1997. [http://www2.kmi.open.ac.uk/tr/kmi-tr-56.pdf]

6. Maarten Sierhuis, personal correspondence.

7. Bannon and Kuuti, op. cit.

8. J. Grudin, comments as discussant during a session at the Human Computer Interaction Consortium 1998 Workshop, Winter Park, Colorado, March 1998.

9. J. Conklin, C. Ellis, L. Offermann, S. Poltrock, A. Selvin, J. Grudin, "Towards an Ecological Theory of Sustainable Knowledge Networks," Submission for National Science Foundation *Knowledge & Distributed Intelligence Proposal Solicitation* NSF9855, 1998.

10. Cite *Computerworld* Knowledge Management article on this point – citation forthcoming!

11. M. Sierhuis, personal correspondence.

12. A. Selvin, *Human Interaction Issues with Group Support Technology in a Software R&D Organization*, NYNEX Science & Technology Technical Memorandum TM 94-0059, 1994.

13. A. Selvin, *Conversational Modeling: A Software-Supported Technique for Analysis by Teams,* NYNEX Science & Technology Technical Memorandum TM96-0013 (1996).

14. B. Zimmermann, A. Selvin, "A Framework for Assessing Group Memory Approaches for Software Design Projects," in *Proceedings of Designing Interactive Systems '97*, Amsterdam, 1997.

15. F. Shipman, C. Marshall, "Formality Considered Harmful: Experiences, Emerging Themes, and Directions," *Technical Report CU-CS-648-93*, Dept. of Computer Science, University of Colorado, Boulder, CO (1992).