Knowledge Management for Distributed Enterprises

Brian R. Gaines, Douglas H. Norrie, Andrew Z. Lapsley and Mildred L.G. Shaw
Knowledge Science Institute and Division of Manufacturing Engineering
University of Calgary
Alberta, Canada T2N 1N4
gaines@cpsc.ucalgary.ca, norrie@enme.ucalgary.ca,
drewl@insource.com, mildred@cpsc.ucalgary.ca

Abstract

The GNOSIS project in the Intelligent Manufacturing Systems international research program is concerned with the use of advanced information technology for knowledge systematization to support the complex intellectual and managerial processes involved in the manufacturing life cycle. It has developed technologies to coordinate distributed manufacturing enterprises, and these technologies have also proved useful in supporting the similar intellectual and managerial processes involved in distributed collaborative research. This article gives the background to the project, and illustrates its use of information technology to provide a corporate memory, its use of knowledge acquisition and modeling tools to model the project objectives and conceptual structures, and the architecture of the Mediator system to support knowledge processes in distributed enterprises.

1 Introduction

In October, 1989, an international Intelligent Manufacturing Systems collaborative program was proposed by Japan through the Ministry of International Trade and Industry (MITI). This proposal was aimed at maintaining and improving the vitality of manufacturing industry and at contributing to the sound development of the world economy, by undertaking joint international research and development among industrialized nations and orienting manufacturing systems towards the 21st Century. Three key research and development objectives were:

- to conduct research in basic and next generation technology;
- to organize and systematize the knowledge so it could be used in developing new technology and facilitating its transfer; and
- to standardize such technology and support the standardization efforts of international organizations.

A series of meetings took place to discuss possible ways and means of international collaboration in advanced manufacturing technology research. The final proposal was that a feasibility study should be undertaken in 1993 by the six Participants: Australia, Canada, the EC, EFTA, Japan, and the USA. The topics included:

- enterprise integration and global manufacturing;
- systematization of manufacturing knowledge;
- the control of distributed intelligent systems;
- techniques for rapid prototyping;
- “clean” manufacturing in the process industries.

Six test case projects were chosen to proceed with one year feasibility studies, amongst which was GNOSIS, concerned with the systematization of knowledge for design and manufacturing. This project was particularly interesting because it was concerned with developing technologies for knowledge modeling, concurrent engineering, virtual factories, fractal enterprises, and so on, but also used these technologies to manage the research processes of the international consortium involved.

This article gives an overview of GNOSIS and its use of various information technologies to support distributed enterprises, providing a corporate memory and modeling its own enterprise in an operational form.

2 The GNOSIS Enterprise

The GNOSIS test case consortium comprised 31 organizations in 14 countries (Figure 1), and involved over 100 researchers. The long term objective of the research program is to develop a new manufacturing paradigm which recognizes problems of the present manufacturing environment—growing scarcity of natural resources; problems of environmental destruction; and issues arising out of regional trade imbalances. The new *post mass production paradigm* is based on *systematization of design and manufacturing knowledge* to acquire and organize knowledge in a form that supports the design and manufacturing of *soft machinery*, i.e., products and factories which achieve reduced resource utilization and waste elimination throughout the whole life cycle from design to reuse or disposal. Soft products and factories are characterized by properties such as reusability, reconfigurability, and flexibility.

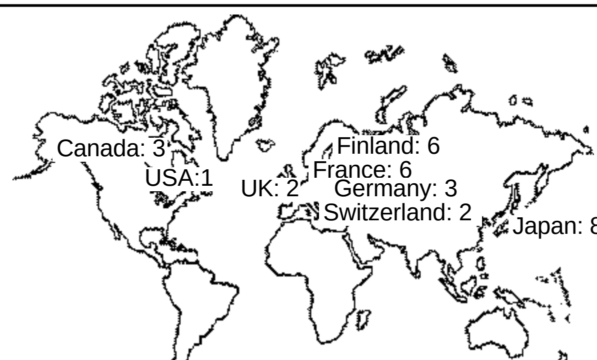

Region: Japan
Mitsubishi Electric Corp.
Nissan Motors Ltd
Fuji Xerox Ltd
Kajima Corp.
Shimizu Corp.
The University of Tokyo
Tokyo Institute of Technology
Kyushu Institute of Technology

Region: EFTA
ABB, Finland
Tehdasmallit Oy, Finland
Synergy Center Ltd, Finland
SIMSON, Finland
VTT, Finland
Tampere University, Finland
EPFL, Switzerland
ETHZ, Switzerland

Region: EC
ADEPA, France
Télémécanique, France
ARM Conseil, France
IBM, France
ITMI, France
LLP/CESALP, France
Fraunhofer Inst. IPA, Germany
MAN Roland, Germany
Brose, Germany
BICC plc, UK
Cambridge University, UK

Region: Canada
Alberta Research Council
DME, University of Calgary
KSI, University of Calgary

Region: USA
Deneb Robotics Inc

Figure 1 The GNOSIS Consortium

The short-term objectives of the GNOSIS test case have been to develop concepts, methodologies, technologies, and tools to implement the new paradigm, and provide a global infrastructure capable of supporting it. Major technologies investigated include *configuration management systems* supporting the reuse of engineering and manufacturing knowledge in routine design and *configurable production systems* achieving dynamic product-specific manufacturing in flexible production systems. The major characteristics, critical drivers, and obstacles to the post mass production paradigm were identified to illuminate a path to the new paradigm..

These objectives were addressed through 5 technical work streams related in content and membership as shown in Figure 2:

- **TW1: Systematization of knowledge and supporting information for design and manufacturing:** Knowledge systematization methods and tools were investigated and demonstrated to support the activities of the other streams.
- **TW2: Configuration management systems:** The focus of the work was the formalization of manufacturing problems and constraints in a distributed environment based on the sharing and reuse of product and manufacturing knowledge.
- **TW3: Configurable production systems:** Research investigated configurable production systems, production control, and factory design principles, with the longer-term goal of enabling soft factories and reuse of production process knowledge.
- **TW4: Soft machinery:** Research utilized the results of the other streams to develop conceptual prototypes of next-generation soft artifacts.
- **TW5: Post mass production paradigm:** Research concentrated on the characterization of the post mass production paradigm and its development strategies.

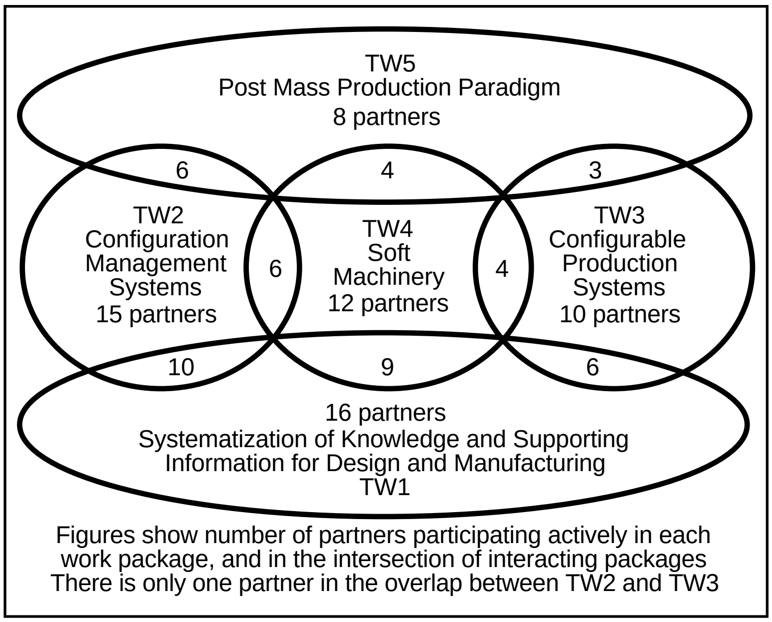

Figure 2 GNOSIS test case relations between work streams

3 Use of Information Technology to Coordinate the GNOSIS Enterprise

The high overheads of international collaboration were immediately apparent in the development of the GNOSIS program which involved in-depth negotiation and coordination across three continents. Because the participants were largely new to one another and came from differing cultures many face-to-face meetings were necessary both before and during the test case. As the overheads of such meetings might be a significant impediment to a successful long-term program, they were monitored carefully. As the GNOSIS research involved systematization of knowledge requiring major computing and communication facilities, it was also an appropriate test case in which to investigate such systematization and technologies to coordinate the project itself and to attempt to reduce the overheads of collaboration.

Each technical work stream involved researchers in many organizations in many countries. Coordination took place through fax, email and meetings. Figure 3 shows the major meetings together with their estimated costs. The total costs of the meetings were some $1.15M which is some 20% of the overall GNOSIS test case budget of $5M. In financial terms this overhead is substantial, but not unreasonably so for a research program based on international collaboration. However, what is rather more difficult to quantify is how effective were the meetings in coordinating the research. For example, to what extent did discussion at the meetings become concrete in terms of documented objectives, detailed schedules, contingency plans, and so on, and were these effective in actually managing the research?

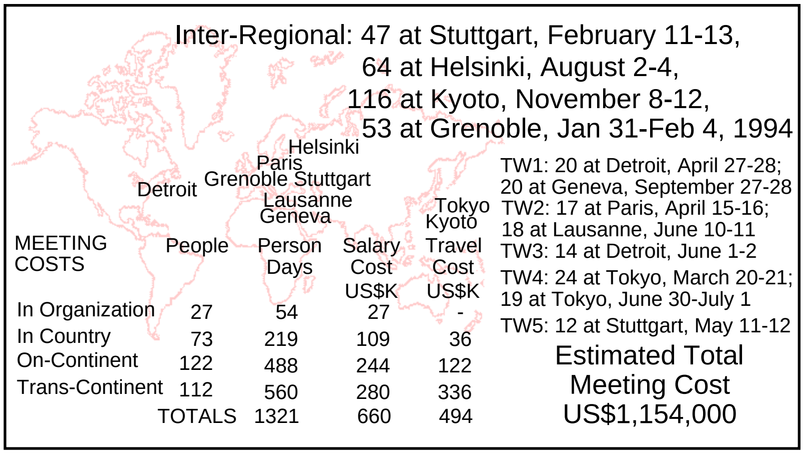

Figure 3 GNOSIS test case major meetings and costs

It was clear from the outset that discussion that did not lead to effective coordination had little value, and hence the emphasis in technical support was upon the capture and systematization of the outcome of the meetings. Maximizing the value derived from the meetings was at least as important as minimizing their cost. The interregional and technical work package meetings allowed partners to come to know one another, to exchange accounts of experience, research capabilities and needs, and to develop plans and agree responsibilities. Fax and telephone communication were used extensively to coordinate ongoing activities. Electronic mail was already in use by 23 of the 31 partner organizations, and mail list servers were set up at partner sites in Tokyo and Stuttgart that supported work package and administrative coordination. An electronic document archive with FTP and WWW access was established in Calgary.

To ensure knowledge capture and systematization at meetings, minutes were written immediately after the meeting, circulated by email and archived for access over the Internet. Experiments were also undertaken on the use of digital document technology to capture the technical contents of the meetings. For example, a camcorder was used to capture software demonstrations and presentations at a 3-day TW4 workshop in Tokyo in March 1993 involving some 30 participants. The slides presented were collected together with the handouts circulated, largely technical papers from conference proceedings and internal reports. On return to Calgary, the papers were optically recognized using Xerox *AccuText* with an accuracy of some 98%, remarkably high considering that some of the papers were late generation copies of double-column conference material in 9pt type. The slides were digitized and the movies converted to QuickTime and edited in Adobe *Premiere*.

The resultant material was put into a uniform style in Microsoft *Word*, resulting in a 300-page workshop proceedings (GNOSIS, 1993) containing 30 articles and 8 QuickTime movies. Figure 4 shows a page from the proceedings with the heading, comments, and movies of computer software demonstrations preceding a technical paper. The production time to OCR the source material, edit it into a uniform format, digitize the slides and movies, and issue the proceedings in print and CD-ROM was 1 week. The effort required was some 40 person-hours. Thus, a high-quality workshop proceedings was produced *after* the event without any additional effort on the part of the participants. The total cost of producing the proceedings was about $3,000, about the same as the air fare to travel to Tokyo.

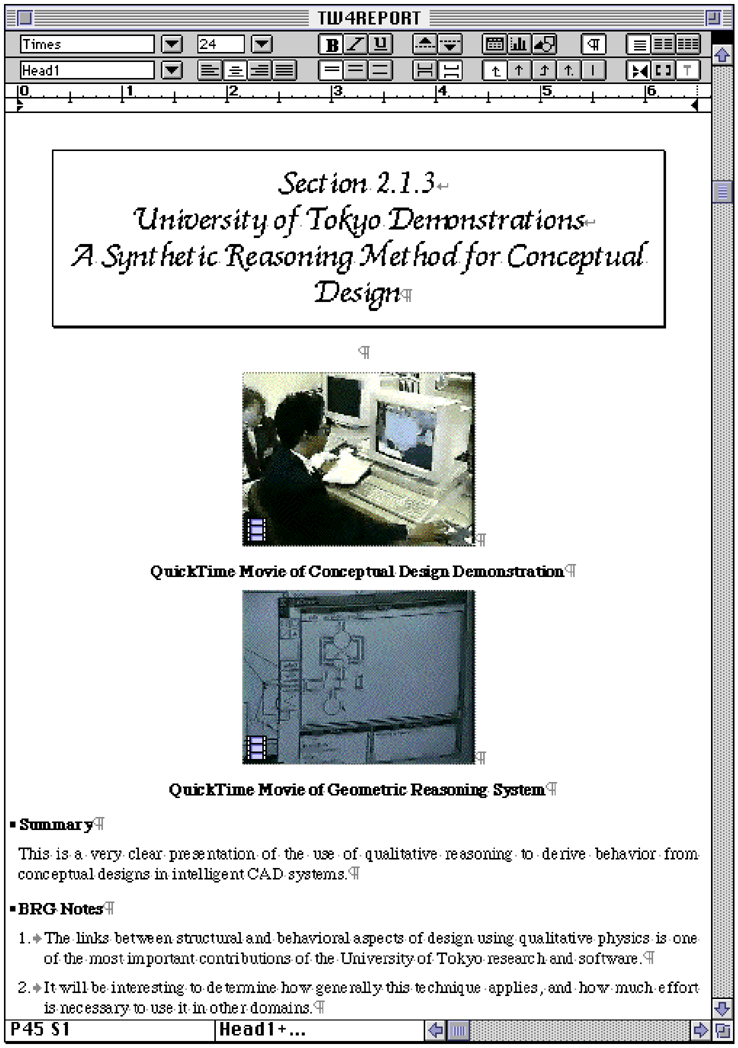

Figure 4 Workshop proceedings with embedded demonstrations

At the end of the GNOSIS test case, the project reports, slide sequences, photographs, movies, and so on, were issued on a CD-ROM. To enable widespread access to the reports, a hybrid format was used that could be read on Macintosh, Windows and UNIX platforms, and all the reports were issued in Microsoft *Word*, Farallon *Replica* and Adobe *PostScript* formats. The total volume of material on the CD-ROM was: 57 reports totaling 1590 pages, each in 3 formats; 11 movies totaling 70 minutes; plus software and maps of the material. 3 test masters were made of the CD-ROM using a Philips *CDD521* recorder and OMI *QuickTOPiX* software. The third one was sent to a CD pressing company that produced 500 copies in 10 days at a cost of $1500.

For the GNOSIS archives, it was appropriate to use as an indexing tool *Mediator* (Gaines and Norrie, 1994; Gaines, Norrie and Lapsley, 1995), a system that had been developed to support collaborative activities across the network as part of the GNOSIS research program. The *Mediator* implementation is based on groupware concept-mapping tools that were already in use for indexing multi-media materials (Gaines and Shaw, 1994). Figure 5 shows the GNOSIS project archives being accessed through layered concepts maps. The map in the window at the upper left is a top level “Server Agent” that manages a particular collection of material. In the example shown a local user is accessing material directly through this agent. Remote users connect to the server agent over the network using client agents that give them the same functionality through calls to the server.

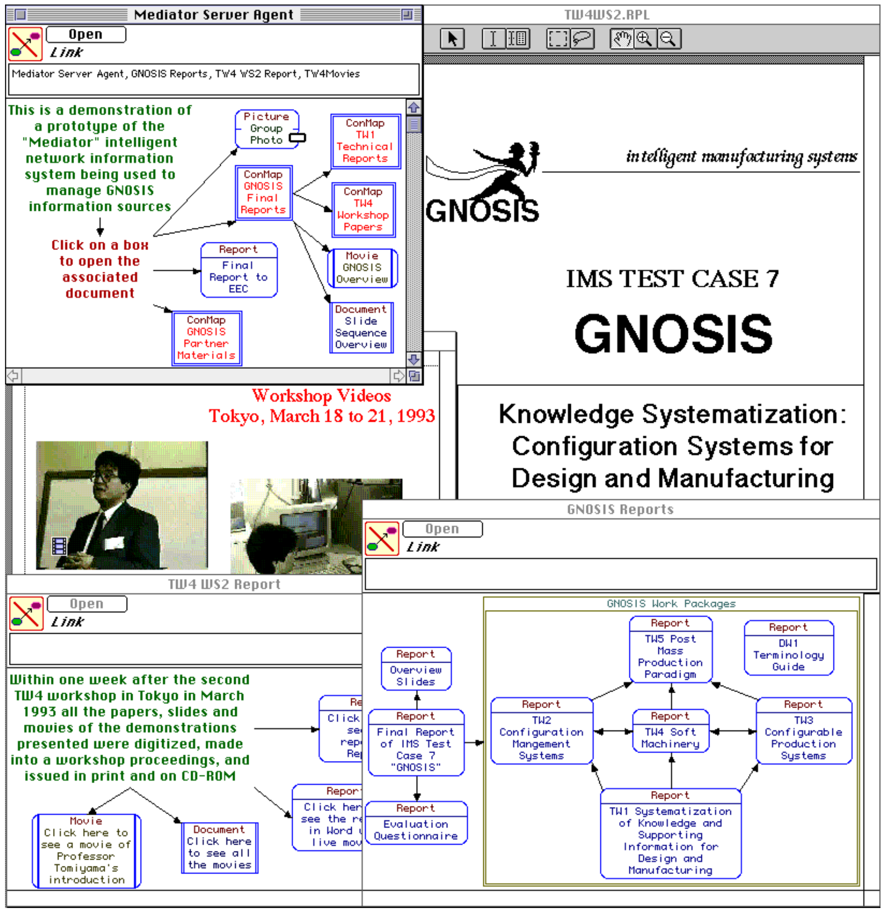

Figure 5 GNOSIS archives indexed hierarchical concept maps in Mediator

The concept map at the top left is currently write-disabled, and the cursor has changed to a button as the user mouses over the "Group Photo" node. Clicking at this point will display the photograph in a separate window. The user has already clicked on the node "GNOSIS Final Reports" to open the concept map shown at the lower right. This has a node for each report, and clicking on one will open the appropriate report, in this application using Farallon's *Replica*. The node at the top left gives access to a series of slides on the project displayed using *Replica*. A similar node in the original concept map at the top left gives access to a movie on the GNOSIS project that will be opened in Apple's *MoviePlayer*.

4 Using Knowledge Acquisition Tools to Model the GNOSIS Enterprise

GNOSIS provided a number of opportunities to use other knowledge acquisition and modeling tools to develop models of the research enterprise, its objectives and conceptual structure. Figure 6 shows a concept map developed automatically from a document using the Texan tool (Shaw and Gaines, 1987) in KSSn which analyses the co-occurrence of words in sentences, a technique commonly used in information retrieval systems (Callon, Law and Rip, 1986). The document analyzed is one on *The Technical Concept of IMS* (Tomiyama, 1992) that played a major role in the design of the research program. The document is treated as a set of entities which are sentences whose features are the words they contain. Rules are derived using empirical induction in which the premise is that if one word occurs in a sentence then the conclusion is that another will occur. The graph shows the links from premises to conclusions derived in this way. The tool is interactive and, as shown at the top center, provides access through a popup menu associated with each word to a list of occurrences of that word in context, and to the original document.

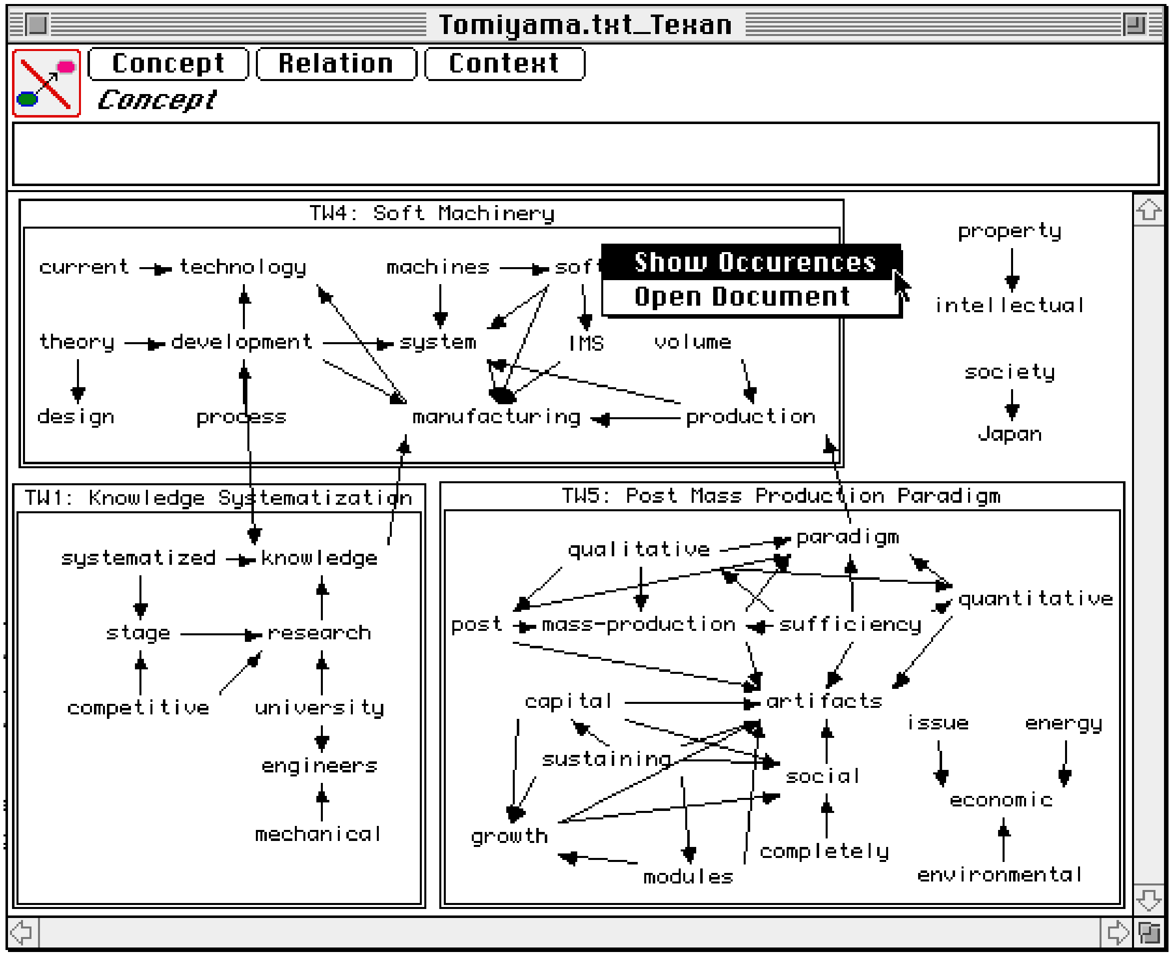

Figure 6 Analysis of paper describing objectives of IMS program

The initial output was a digraph consisting of one major connected component and some minor ones which correspond to significant topics such as intellectual property rights that did not directly relate to the socio-technical issues. The user noted that the major component itself consisted of 3 loosely connected sub-components, and added the context boxes shown to distinguish and name these parts. The significance of these parts is that they correspond to 3 of the 5 technical work packages (TW's) of the research program. What is particularly significant is the missing work packages, TW2 concerned with product configuration management systems, and TW3 concerned with configurable production systems. These were added to the GNOSIS research program during its formation through amalgamation of interests with other potential proponents of IMS test cases. These packages link technically to the knowledge systematization activities on the left of Figure 6 but are neutral to the major issues of the IMS program on the right.

Figure 7 shows the *mission statement* of the GNOSIS project as a concept map. This concise statement of the project objectives was taken from the introduction to the legal agreement signed by all participants, and much effort went into its formulation. The upper right part of the map is concerned with the *post mass production paradigm* studies (TW5) that show up on the lower right of Figure 6. The pivotal role of *knowledge systematization* studies (TW1) shows as linking TW5 to *configurable production systems* (TW3) and to *soft machines* studies (TW4). Studies of *configuration management systems* (TW2) are visible only in the nodes *product configuration* and *configurable products*. In practice, TW2 took off more slowly than the other work packages, and the analysis of the project documents tends to indicate that the objectives for this work package were not as clearly formulated or integrated into the overall project as for the other work packages. Examination of Figures 6 and 7 also suggests that there is a gap that needs filling between the very high level socio-economic goals of the IMS program and the very specific technical objectives of the work packages other than TW5. The popup menu on the right shows the user opening a document on environment issues through the access provided by the concept map environment.

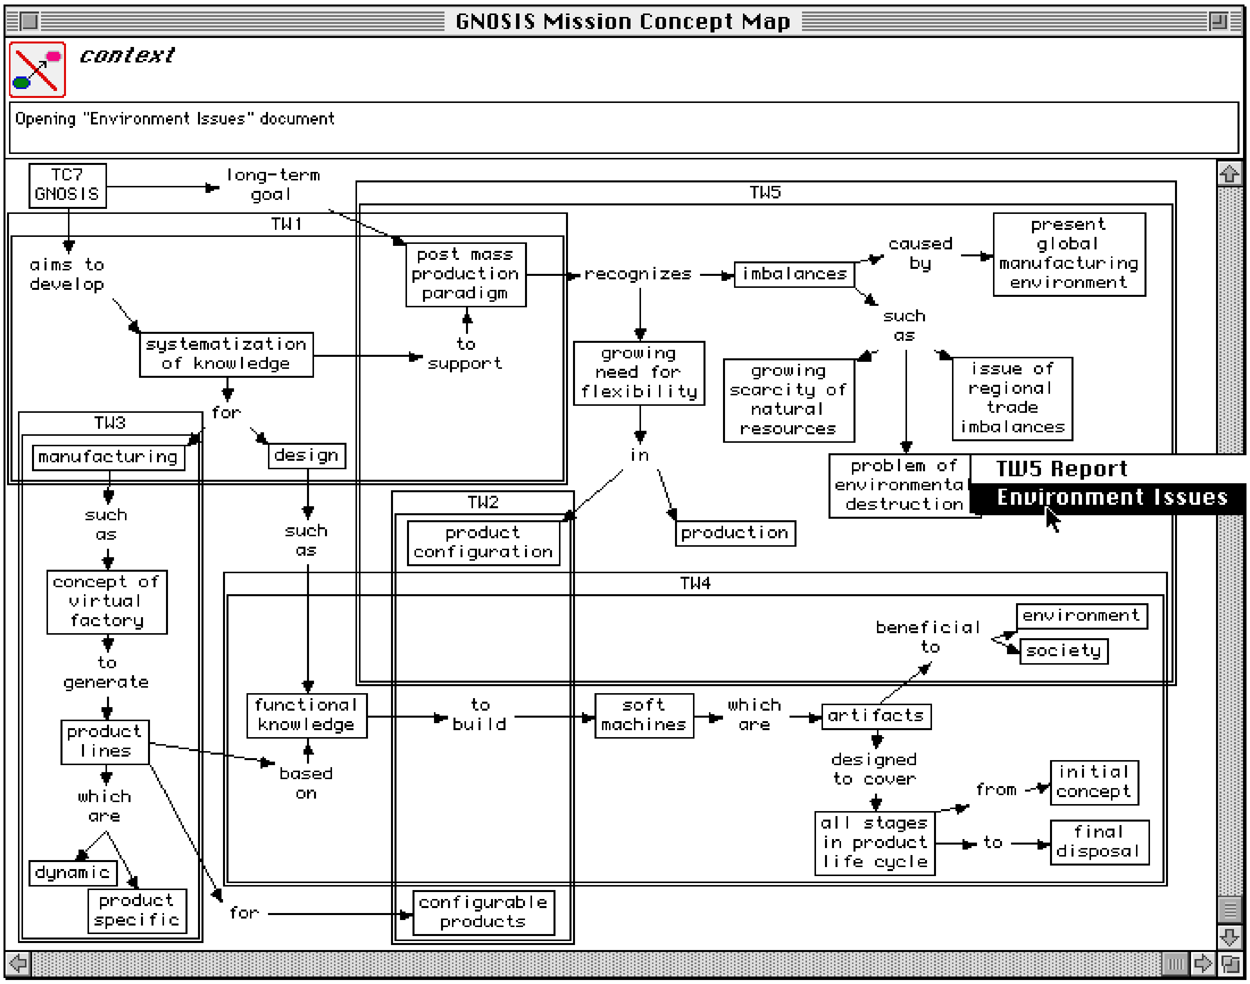

Figure 7 Concept map of GNOSIS mission statement

One problem for GNOSIS has been the presentation of the project objectives and activities to funding and reviews agencies concerned with the international and national programs. There was a common theme of *reconfigurable systems* in the three major technical work packages, but it became clear that this was inadequately projected in the project documents. In preparation for a major review meeting in June 1993 an analysis was made of the *soft machine* concept using a KA technique derived from personal construct psychology (Kelly, 1955), that of *repertory grid elicitation* (Gaines and Shaw, 1980; Boose, 1984; Gaines and Shaw, 1993).

Six major GNOSIS sub-projects were used as initial elements, and the ensuing repertory grid elicitation process resulted in the addition of another 10 elements, including human operators and organizational structures that provided contrasts to some aspects of the technological projects. Eleven distinctions were elicited that provided detailed insights into the complexity of the notion of reconfigurability, and these were presented on viewfoils to the review body to explain the roles of the GNOSIS projects and the relations between them relevant to issues of soft machinery.

Figure 8 shows the grid clustered through a principal components analysis to bring together similar distinctions and similar elements, and to show the relationships between them. It shows, for example, that some past distinctions in manufacturing are no longer as critical as they used to be—the *hardware—software* distinction did not characterize other distinctions—GNOSIS treats software manufacturing and hardware manufacturing alike. The main dimensions apparent are those typified by *system reconfigures itself—user reconfigures system* and that of *human intelligence—machine intelligence*. The IMS projects are on the machine intelligence side of the second dimension and cluster into two groups on the first, those concerned with self-reconfiguring systems and those concerned with user-reconfiguring systems.

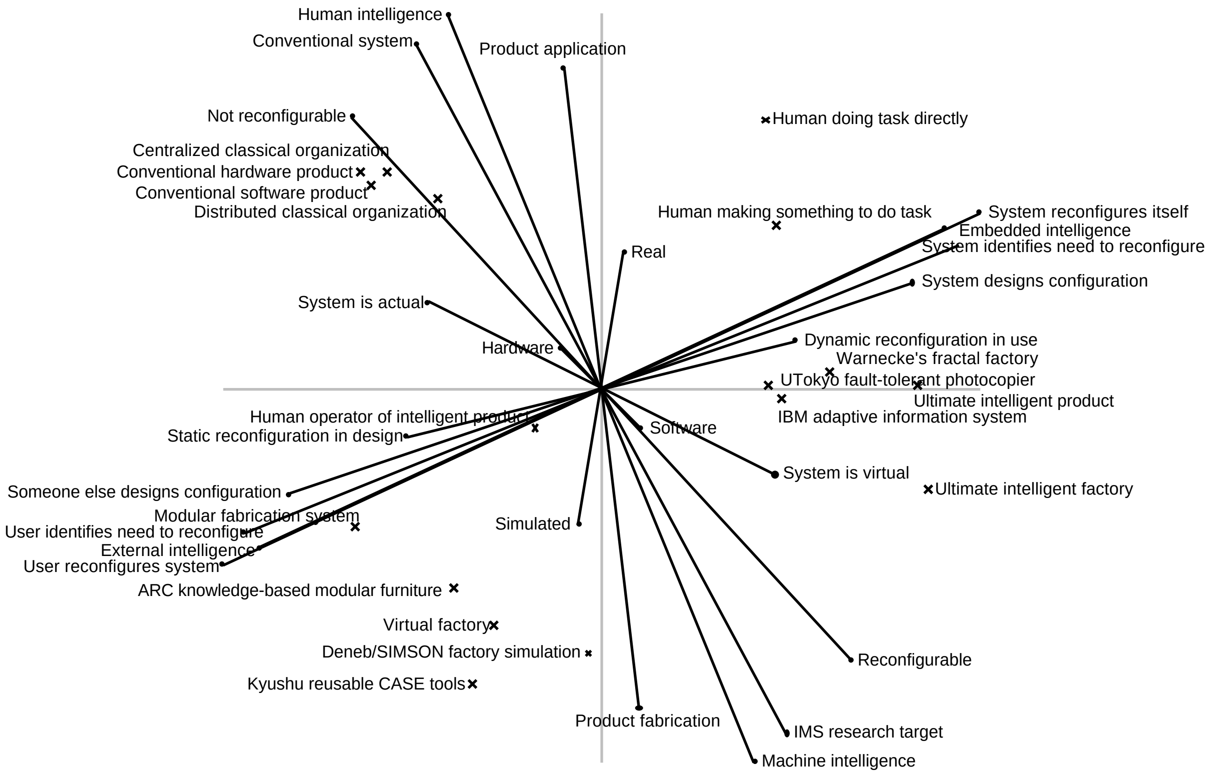

Figure 8 Principal components analysis of soft machine repertory grid

5 Mediator: Knowledge Management for a Distributed Enterprise

One of the research projects in GNOSIS is Mediator, a system designed to coordinate the overall manufacturing process over the complete product life cycle from requirements to recycling. In particular, it supports the integration of computer-based system design, production engineering and production sub-systems into a mutually collaborative framework for integrated design and manufacture. It provides facilities for recording and tracing decisions taken at each stage of the product life cycle, particularly the dependencies between knowledge, decisions, datasets, and so on.

The Mediator architecture is a distributed client, distributed server design, in which multiple users can collaborate synchronously or asynchronously through processes running anywhere on the network. The implementation supports a heterogeneous environment in which there are multiple protocols and multiple forms of user interface. Figure 9 shows the basic architecture. At the center is a collaborating and geographically dispersed user community. Beneath this is shown the computational infrastructure to support the collaboration.

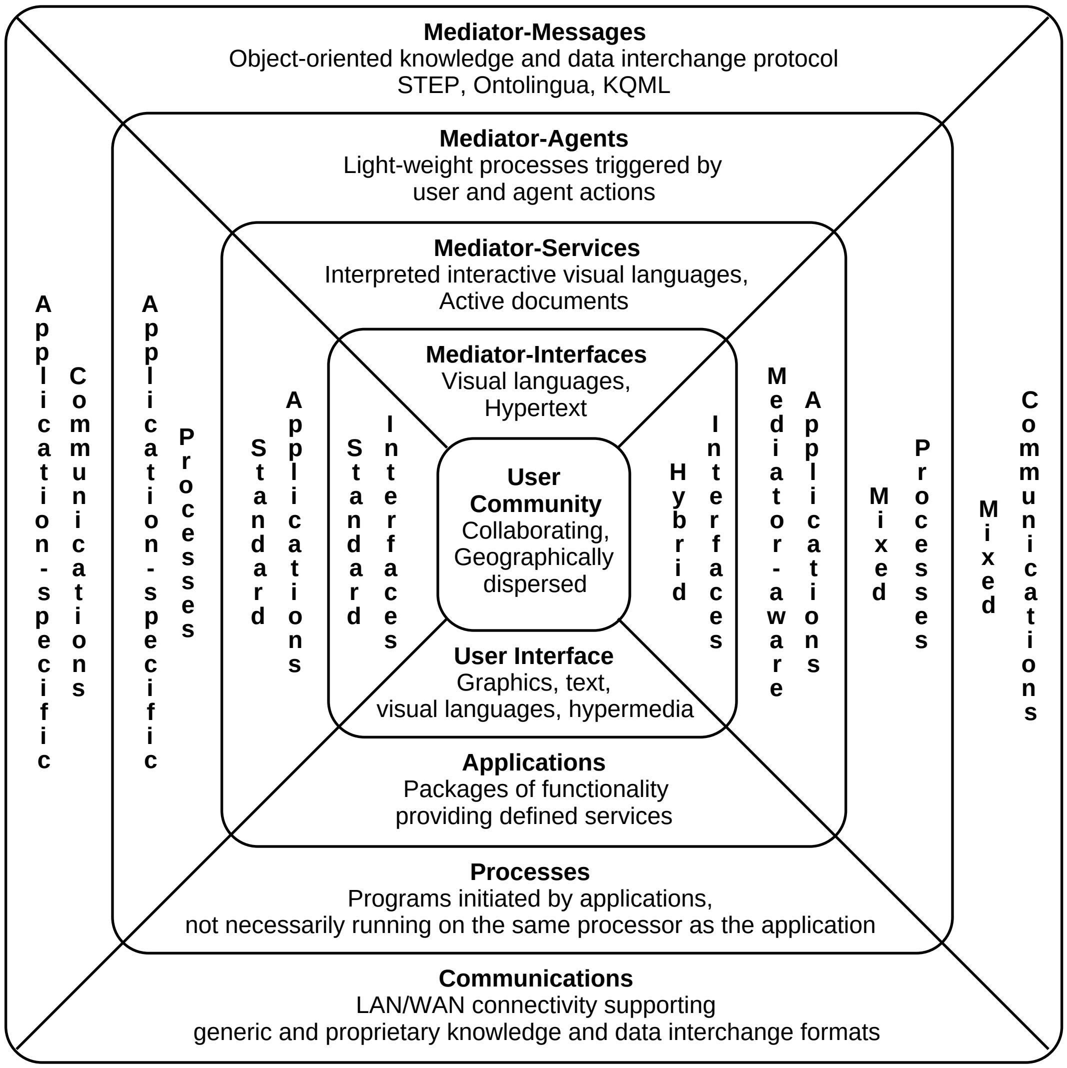

Figure 9 Mediator architecture

The community interacts with the system through a variety of forms of user interface, typically graphics, text, visual languages and hypermedia. They access a variety of ‘applications’ defined as packages of functionality providing defined services. The functionality is made operational by initiating processes which may run anywhere on the network; that is, remote procedure calls are expected to be common. Inter-process communications is provided through local and wide area networks supporting a range of generic and proprietary knowledge and data interchange formats.

On the left of Figure 9, these four layers are shown instantiated in terms of conventional software packages designed independently of mediator: as standard interfaces, standard applications, application-specific processes and application-specific communications. It is assumed that such standard applications can play a role in a Mediator-coordinated systems, minimally by the application and its datasets being registered in Mediator with its application windows open on the Mediator desktop, and maximally by Mediator controlling its inputs, outputs and operations by job-control scripts. Existing application software is assumed to play a major role in the operation of Mediator, and one can conceptualize this role by noting that the Mediator system will know a lot about such applications while they will know virtually nothing about Mediator.

At the top of Figure 9, the four layers are shown instantiated in terms of core Mediator-specific technology, what might be called a *Mediator-shell* since the majority of the software is not specific to any particular application of Mediator, but provides general collaboration and integration facilities. The primary user interface to Mediator is through visual languages allowing the general representation of semantic systems through graphical symbols. This is supplemented by hypertext and hypermedia as appropriate. The application layer uses particular instances of these visual languages with application-specific semantics and a visual appearance designed to be natural to use in the context of the specific application.

The process layer supports agents as light-weight processes triggered by user interaction with the visual languages, and also triggered by other agents (Kwok and Norrie, 1993). The communication layer supports object-oriented protocols for knowledge and data interchange such as STEP, Ontolingua and KQML (Finin, Weber, Wiederhold, Genesereth, Fritzson, McKay, McGuire, Shapiro and Beck, 1992). It also supports messages in arbitrary formats as appropriate to communication with other applications not designed with Mediator in mind. On the right of Figure 9, the four layers are shown instantiated in terms of separately designed applications that are 'Mediator-aware' to some extent, for example in using the Mediator interface technology, applications, agents or protocols as part of their normal operation. The Mediator shell technology is designed to be highly modular and readily integrated in whole or in part with existing applications. Thus, in a design domain, one might envision a Mediator graphic interface and data interchange protocols being integrated with functional design tools such as HUT's Δ (Gui and Mäntylä, 1993) or University of Tokyo's SYSFUND (Ishii, Tomiyama and Yoshikawa, 1993).

Much of the user interaction with Mediator is through an open architecture visual language system developed to support graphic interaction with computers on a wide range of topics including knowledge representation (Sowa, 1984; Gaines, 1991b), concept mapping (Gaines and Shaw, 1994; Kremer and Gaines, 1994), Petrinets (Reisig, 1985), bond graphs (Karnopp, Rosenberg and van Dixhorn, 1989), and so on. The system was developed as an alternative to textual interfaces to take advantage of modern graphic workstations through a simple and natural visual language of great generality. The generality is achieved because the underlying data structure is a sorted digraph which can represent typed binary relations and hence virtually any mathematical structure. A comprehensible user interface is provided through the capability to reflect the type structure in node decorations such as fonts, shading and color.

User interaction with Mediator takes place through the creation of statements in the visual language, and through interaction with these statements through popup menus whose content is specific to node type. Actions are context-sensitive: to the node selected for the popup, to nodes linked to it, and to other nodes preselected by clicking on them in the graph. This allows both simple and complex activities to be initiated by simple and comprehensible user actions. For example, Figure 10 shows a shop floor information access concept map being used to access bill of materials information for a particular product.

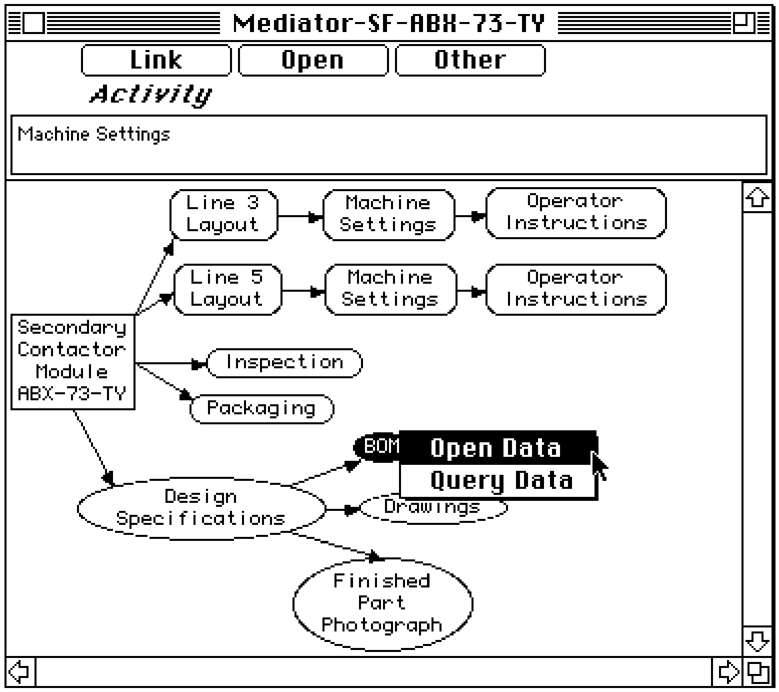

Figure 10 Interacting with a shop floor information map

The general concept map tool allows the style and functionality of such windows to be tailored to general classes of situations. Individual maps are then generated automatically or manually for particular situations, in this case a product with two possible line layouts. The user interface can be programmed using a wide range of possible dialogs. In this case a popup menu gives access to the bill of materials, through another concept map, or by initiating action with the appropriate database. The visual language provides a very general and easily customizable interface to the underlying knowledge and data structures, and processes operating on them such as specific applications and general agents, including communications with, and remote procedure calls to, other systems. The language may be used as a 'wrapper' to existing applications available only in binary form, and it may also be used as an embedded component of other applications available in source form. A computer-supported cooperative work approach has been adopted from the outset so that maps may be shared across local and wide area networks and used for distributed project coordination.

The open architecture, modular, networked, distributed client, distributed server technology underlying the Mediator implementation is well-suited to vertical integration with applications as shown on the right of Figure 9. The development of 'Mediator-aware' applications is attractive not only in increasing the functionality of Mediator, but also in allowing rapid prototyping of new applications. The effort required to develop good quality user interfaces is a major component of resource usage in most application development, and the use of the generic Mediator interfaces and protocols can increase the speed, and decrease the cost, of specialist system development. For example, a research group focusing on optimal scheduling and rapid shop-floor reconfiguration could develop algorithms that interface to Mediator for purposes of user interaction, and run on a high-performance server anywhere on the network.

Figure 11 shows a bond graph from HUT's Δ functional design system (Gui and Mäntylä, 1993) represented in the visual language, and Figure 12 shows a partial behavioral model from University of Tokyo's SYSFUND functional design system (Ishii et al., 1993) similarly represented. The generic Mediator data structures can be imported from, and exported to, the specific systems through simple converters.

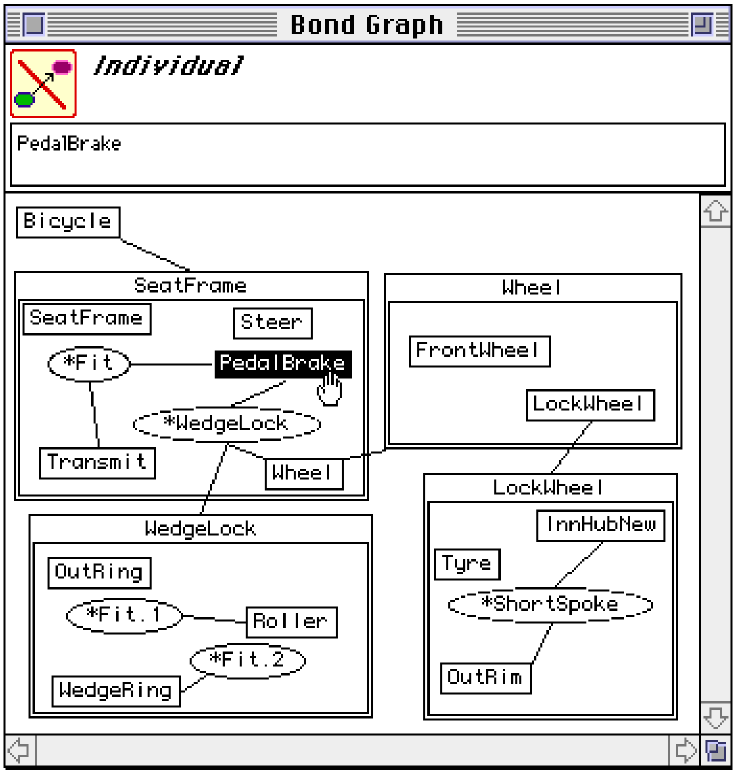

Figure 11 A Δ bond graph in Mediator

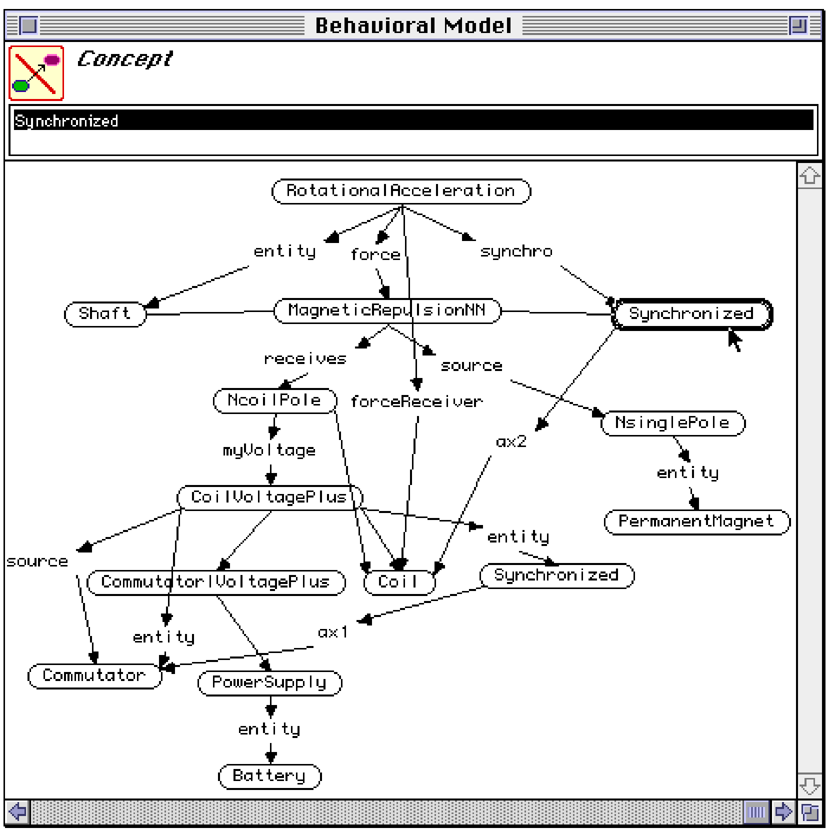

Figure 12 A SYSFUND behavioral model in Mediator

A semantic network visual language is already available for Mediator (Gaines, 1991b) that compiles into a KL-ONE-like knowledge representation system (Gaines, 1991a; Gaines, 1993a) used to represent formal knowledge structures (Gaines, 1993b). This has been used to operationalize a corporation's procedures manuals (Kremer, 1991) in a way that is applicable, for example, to the regulatory aspects of manufacturing.

6 Mediator Implementation on the World Wide Web

The initial Mediator prototype developed during the GNOSIS test case was implemented on the Apple Macintosh platform and operated over local area networks. It was issued on the CD-ROM that contained the final reports of the GNOSIS test case (GNOSIS, 1994), and it was used to index these reports and the accompanying demonstration software and digital videos through layered concept maps accessing the heterogeneous collection of files that comprised the final reports, demonstrations and data sets. Figure 5 has already illustrated Mediator in use to access material on the GNOSIS CD-ROM (Gaines and Norrie, 1995).

Since the completion of the GNOSIS test case in March 1994 we have had a systematic program of research designed to make Mediator technology widely available on a cross-platform basis operating through the Internet. The hypertext transport protocol (HTTP, Berners-Lee, 1993) of the World Wide Web (WWW, Berners-Lee, Cailliau, Luotonen, Nielsen and Secret, 1994) was chosen for the low-level communications layer since HTTP servers are widely available for a range of platforms, as were associated browsers such as NCSA Mosaic and Netscape. In particular, the browsers support not only the hypertext markup language (HTML, Barry, 1994) but also arbitrary data types through interface to other 'helper' applications such as our concept mapping tools. This makes it possible to implement the Mediator architecture of Figure 9 using standard protocols and application packages in large part. This is particularly significant because one objective of the research program in the past 2 years has been to move Mediator out of the research laboratory and make it widely available on the Internet to support the long-term IMS research program.

The main programming effort has been to port the concept mapping tools cross-platform so that they can act as client helpers to WWW browsers on any platform. Figure 13 shows Xcm operating under Motif and X-Windows under SunOS as a helper to Netscape.

At the top right of Figure 13 a Mediator concept map of a sequence of design processes for an engine mount has been opened. This was fetched by Netscape from an HTTP server on the Internet , and because the file extension was '.cm' the server and Netscape recognized it as a concept map data type and opened it in the Xcm application. A popup dialog associated with each node allows files to be specified that are associated with the nodes following the Mediator protocols. However, the file accesses are now specified through WWW uniform resource locators (URLs) which enables them to be fetched from any HTTP server on the Internet.

At the top left the user has fetched a FrameMaker document containing the requirements specification for the engine mount which Netscape has opened in FrameMaker. At the lower left he has fetched an annotated requirements specification in HTML which has opened in Netscape and itself contains hypertext links to other files.

At the top center the user has opened an HTML document containing the STEP specification of the mount, and at the lower right the AutoCAD drawing which Netscape has opened in AutoCAD. A similar screen would be generated on other platforms by the same sequence of actions.

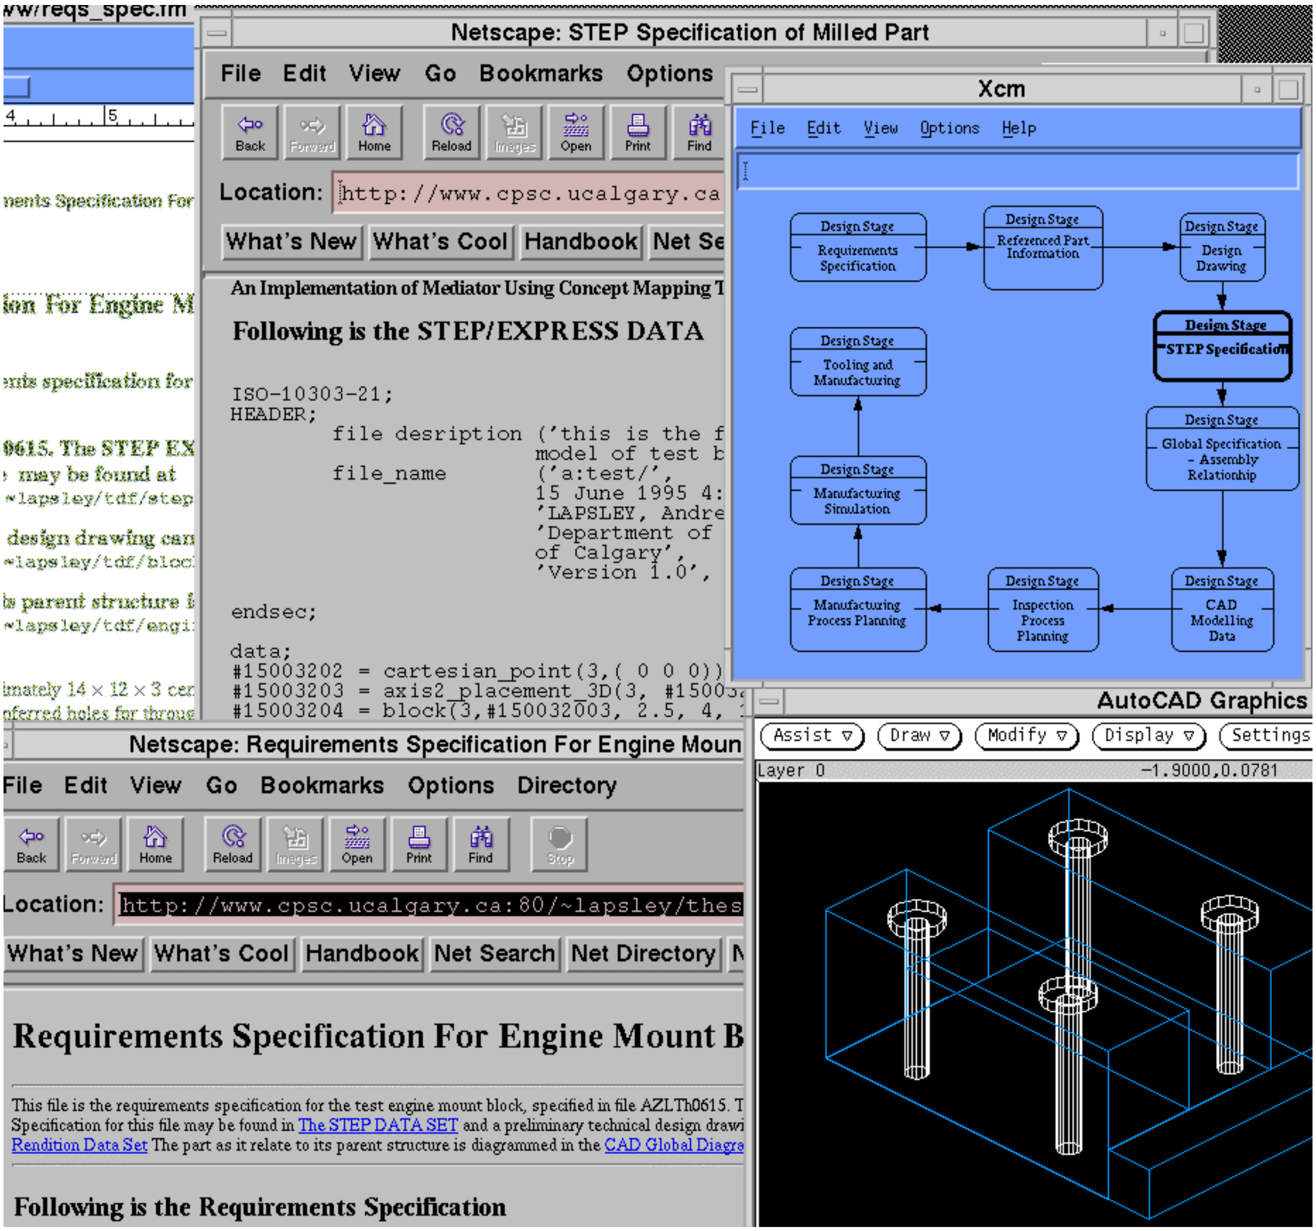

Figure 13 Screen dump of Mediator running on a UNIX workstation

One of the problems identified in using Mediator has been that of mutual awareness. If someone is working with partners at other sites, how do they know what they have done recently, what relevant new materials have been developed, what tasks have been accomplished? It is tedious and inefficient for them to search the others' sites for new material, particularly since the other partners are assumed relatively autonomous and may have undertaken unexpected activities. We developed CHRONO (Chen, 1995), a tool for automatically indexing a site in reverse chronological order in order to overcome this problem, and META-CHRONO for collating the CHRONO data from several sites in order to provide a unified awareness system of the activities of all relevant partners.

The problems of awareness in Internet communities are of interest in their own right and we have developed models of how knowledge is acquired through the net and web, and what patterns of discourse are involved which are described in an associated paper (Chen and Gaines, 1996).

7 Conclusions

The GNOSIS project is concerned with the use of advanced information technology for knowledge systematization to support the complex intellectual and managerial processes involved in the manufacturing life cycle. It has developed technologies such as Mediator to do so, and these technologies have proved useful in supporting the complex intellectual and managerial processes involved in distributed collaborative research. In the manufacturing environment we take it for granted that it is possible to develop a detailed technical model of the system to be operated. What is the equivalent for the research environment? Can we build a model of the system of researchers and resources, their processes and coordination, and make this overt to support the research itself?

These questions were not part of the mandate of the GNOSIS project, but they arise naturally out of the logic of the test cases, which was to investigate the cost-effectiveness of large-scale international collaboration in research. Kim's (1991) use of insights from manufacturing to develop a formal model of creative decision making, the AAAI92 workshop on the role of AI in communicating scientific and technical knowledge (Swaminathan, 1992), Rosenschein and Zlotkin's (1994) formalization of human negotiation conventions to apply to computational agents and other such studies that treat human and computer agents within a common framework, indicate that the questions posed are timely and appropriate to the current stage of information technology.

There are models of human society that treat it as a layered system of compound entities in which individuals, groups, and organizations are functional agents, each recursive sub-divisions of humanity itself conceptualized as an intelligent agent (Miller, 1978; Tracy, 1989; Gaines, 1994). From such a *collective stance*, the GNOSIS project may be seen as an intelligent compound entity consisting of distributed agents coordinated through the communication systems described in the previous section (Figure 17).

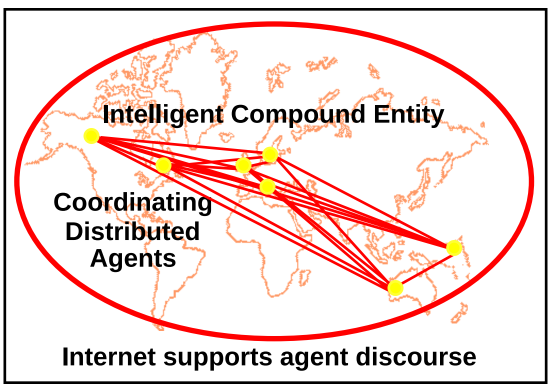

Figure 17 GNOSIS as an intelligent compound entity

The essence of all these models is to take a systemic view which abstracts both the human and technical components to form a society of agents. The problem is make the models operational and effective to enable the design of research support systems that facilitate the systematic acceleration of scientific research. In conclusion, we see it as a challenge for the systems community to develop operational models of knowledge creation that can be used to design support systems for international collaborative research. The basic computer and communication technology is already in place, the need to collaborate is recognized at the highest levels of government, and the will to succeed exists in individual researchers and their organizations. The formal foundations do not exist as yet, and without them we are undertaking major enterprises on a trial-and-error basis that will be counter-productive in the long run.

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada.

References

Barry, J. (1994). The hypertext markup language (HTML) and the World-Wide Web: Raising ASCII text to a new level of usability. **Public-Access Computer Systems Review 5**(5) 5-62.

Berners-Lee, T. (1993). Hypertext Transfer Protocol. CERN, Geneva. ftp://info.cern.ch/pub/www/doc/http-spec.ps.Z.

Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H.F. and Secret, A. (1994). The World-Wide Web. **Communications ACM 37**(8) 76-82.

Boose, J.H. (1984). Personal construct theory and the transfer of human expertise. **Proceedings AAAI-84**. pp.27-33. California, American Association for Artificial Intelligence.

Callon, M., Law, J. and Rip, A., Ed. (1986). **Mapping the Dynamics of Science and Technology**. Basingstoke, UK, MacMillan.

Chen, L.L.-J. (1995). CHRONO: A Chronological Awareness Tool. Knowledge Science Institute, University of Calgary. http://www.cpsc.ucalgary.ca:80/~lchen/cpsc.html#chrono.

Chen, L.L.-J. and Gaines, B.R. (1996). Knowledge Acquisition Processes in Internet Communities. Gaines, B.R. and Musen, M.A., Ed. **Proceedings of Tenth Knowledge Acquisition Workshop**.

Finin, T., Weber, J., Wiederhold, G., Genesereth, M., Fritzson, R., McKay, D., McGuire, J., Shapiro, S. and Beck, C. (1992). Specification of the KQML Agent-Communication Language. The DARPA Knowledge Sharing Initiative External Interfaces Working Group.

Gaines, B.R. (1991a). Empirical investigations of knowledge representation servers: Design issues and applications experience with KRS. **ACM SIGART Bulletin 2**(3) 45-56.

Gaines, B.R. (1991b). An interactive visual language for term subsumption visual languages. **IJCAI'91: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence**. pp.817-823. San Mateo, California, Morgan Kaufmann.

Gaines, B.R. (1993a). A class library implementation of a principled open architecture knowledge representation server with plug-in data types. **IJCAI'93: Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence**. pp.504-509. San Mateo, California, Morgan Kaufmann.

Gaines, B.R. (1993b). Experience with a class library for organizational modeling and problem solving. **Integrated Computer-Aided Engineering 1**(2) 93-107.

Gaines, B.R. (1994). The collective stance in modeling expertise in individuals and organizations. **International Journal of Expert Systems 7**(1) 21-51.

Gaines, B.R. and Norrie, D.H. (1994). Mediator: information and knowledge management for the virtual factory. **SIGMAN AAAI-94 Workshop: Reasoning about the Shop Floor**. pp.30-39. Menlo Park, California, AAAI.

Gaines, B.R. and Norrie, D.H. (1995). Knowledge systematization in the international IMS research program. **Proceedings of 1995 IEEE International Conference on Systems, Man and Cybernetics**. pp.958-963. New York, IEEE.

Gaines, B.R., Norrie, D.H. and Lapsley, A.Z. (1995). Mediator: an Intelligent Information System Supporting the Virtual Manufacturing Enterprise. **Proceedings of 1995 IEEE International Conference on Systems, Man and Cybernetics**. pp.964-969. New York, IEEE.

Gaines, B.R. and Shaw, M.L.G. (1980). New directions in the analysis and interactive elicitation of personal construct systems. **International Journal Man-Machine Studies 13** 81-116.

Gaines, B.R. and Shaw, M.L.G. (1993). Basing knowledge acquisition tools in personal construct psychology. **Knowledge Engineering Review 8**(1) 49-85.

Gaines, B.R. and Shaw, M.L.G. (1994). Concept maps indexing multimedia knowledge bases. **AAAI-94 Workshop: Indexing and Reuse in Multimedia Systems**. pp.36-45. Menlo Park, California, AAAI.

GNOSIS (1993). TW4 Soft Machinery Workshop Proceedings. Knowledge Science Institute, University of Calgary, Canada. ftp://ksi.cpsc.ucalgary.ca/IMS/GNOSIS/TC7/tw4ws2.

GNOSIS (1994). GNOSIS: Intelligent Manufacturing Systems: IMS Test Case 7: Hybrid-CD, Macintosh (native), PC, Unix (ISO 9660). Knowledge Science Institute and Division of Manufacturing Engineering, University of Calgary, Canada. ftp://ksi.cpsc.ucalgary.ca/IMS/GNOSIS/TC7.

Gui, J.-K. and Mäntylä, M. (1993). Assembly modeling on the basis of a mechanical design prototype. Laboratory for Information Processing Science, Helsinki University of Technology.

Ishii, M., Tomiyama, T. and Yoshikawa, H. (1993). A synthetic reasoning method for conceptual design. Department of Precision Mechanical Engineering, The University of Tokyo.

Karnopp, D., Rosenberg, R.C. and van Dixhorn, J.J. (1989). Bond Graph Techniques for Dynamic Systems in Engineering and Biology. **Journal Franklin Institute 308**(3)

Kelly, G.A. (1955). **The Psychology of Personal Constructs**. New York, Norton.

Kim, S.H. (1991). A formal model of creative decision making. **Robotics and Computer-Integrated Manufacturing 8**(1) 53-65.

Kremer, R. and Gaines, B.R. (1994). Groupware concept mapping techniques. **Proceedings SIGDOC'94: ACM 12th Annual International Conference on Systems Documentation**. pp.156-165. New York, ACM.

Kremer, R.C. (1991). Experience in applying KRS to an actual business problem. Boose, J.H. & Gaines, B.R., Ed. **Proceedings of the Sixth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.11-1-11-12. Calgary, Canada, University of Calgary.

Kwok, A.D. and Norrie, D.H. (1993). Integrating multiple reasoning in intelligent agent systems. **Integrated Computer-Aided Engineering 1**(2) 83-90.

Miller, J.G. (1978). **Living Systems**. New York, McGraw Hill.

Reisig, W. (1985). **Petri Nets: An Introduction**. Berlin, Springer.

Rosenschein, J.S. and Zlotkin, G. (1994). **Rules of Encounter: Designing Conventions for Automated Negotiation among Computers**. Cambridge, Massachusetts, MIT Press.

Shaw, M.L.G. and Gaines, B.R. (1987). KITTEN: Knowledge initiation and transfer tools for experts and novices. **International Journal of Man-Machine Studies 27**(3) 251-280.

Sowa, J.F. (1984). **Conceptual Structures: Information Processing in Mind and Machine**. Reading, Massachusetts, Adison-Wesley.

Swaminathan, K., Ed. (1992). **AAAI-92 Workshop on Communicating Scientific and Technical Knowledge**. Menlo Park, California, AAAI.

Tomiyama, T. (1992). The technical concept of IMS. RACE Discussion Paper, No. RA-DP2, Research into Artifacts, Center for Engineering, The University of Tokyo.

Tracy, L. (1989). **The Living Organization: Systems of Behavior**. New York, Praeger.