MEMOIR

ESPRIT Project Number 22153

Unichema International

Glaxo Wellcome Research and Development

IBEX Computing

Multicosm

PAC

University of Athens

D1.10 Final Report (Public version)

Version 1

PAC/MEMOIR/D1.10 Version 1

Chris Scott and Mark Stairmand, Parallel Applications Centre
Richard Miller and Steve Loades, Unichema International
Andrzej Glowinski, Glaxo Wellcome R&D
Eric Gerelle and Jacques Hale, IBEX Computing
Roger Rowe, Peter McManus and Gary Hill, Multicosm
Costas Halatsis, Mike Hatzopolous and Aggelos Pikrakis,
University of Athens
Wendy Hall, Dave DeRoure and Sigi Reich,
University of Southampton MMRG

28 September 1998

Executive Summary

This report is the final report of the MEMOIR Project (ESPRIT 22153). The project involves the following partners: Glaxo Wellcome R&D (UK), Unichema International (NL), IBEX Computing (F), Multicosm (UK), PAC (UK) and the University of Athens (GR).

The MEMOIR partners set out to demonstrate that advanced IT systems based on the concepts of hypermedia links, trails of captured user activity, software agents and open, modular systems architecture could deliver business benefits such as:

- increasing the productivity of researchers and information professionals within science and technology based companies;
- improving access to information which is relevant to particular needs at a particular time; and
- enabling relevant sources of knowledge available from colleagues to be identified. (This has been expressed by many large companies as "If only we knew what we knew…..").

A prototype software system with the potential to deliver such benefits has been successfully designed and developed and integrated into the IT infrastructures of the two end user companies, Glaxo Wellcome and Unichema.

User evaluation of the prototype MEMOIR system has not so far been as extensive as originally envisaged, and as a consequence conclusive validation of the benefits of MEMOIR or MEMOIR-like systems is still to be achieved as part of the follow-through from the project. However, the benefits gained by the project partners have been considerable, with much reusable knowhow acquired and many valuable lessons learned. In particular, MEMOIR has had a profound influence on the market positioning and technology development plans of both commercialisation partners in the consortium, IBEX and Multicosm.

Dissemination of MEMOIR results has been effective and widespread. Papers and demonstrations have been presented at several major IT conferences, briefings have been given to consultants and IT industry opinion formers, reviews and joint projects have been configured with major IT industry companies, and considerable use has been made of MEMOIR as a knowledge management infrastructure case study at Unichema's former and current parents companies, Unilever and ICI.

Contents

1 Introduction

1.1 General

This report is the final report of the MEMOIR Project (ESPRIT 22153). The project involves the following partners: Glaxo Wellcome R&D (UK), Unichema International (NL), IBEX Computing (F), Multicosm (UK), PAC (UK) and the University of Athens (GR).

1.2 Structure of the Report

This document conforms to the structure suggested in the EC's "FINAL REPORT OUTLINE" document (ref. Finrep.doc: final report outline – April 1994 -).

1.3 Abbreviations Used in the Report

The following abbreviations are used throughout the report:

SOP: Start of the project

M1 etc: Month 1…Month 24 of the project

Q1 etc: First quarter of the project …eighth quarter of the project

Deliverable reference numbers are taken from Table 5.2 in the Technical Annex.

2 Management report

2.1 Introduction

Section 2 reports on the main issues of the project and covers

- the objectives of the project and the background to it at the time the project commenced;
- the major changes in context which have occurred during the project lifetime;
- the originally envisaged market opportunities that would be enabled as a result of the project;
- the major achievements and results of the project;
- the significance of the achievements and results for the partners, the consortium, and for European industry in general;
- a summary of the commercial exploitation opportunities that have resulted from the project;
- the management of the project, and the major management issues that have arisen during the course of the project;
- the utilisation of resources within the project;
- a summary of the deliverables from the project; and
- a summary of the major dissemination activities and publications resulting from the project.

2.2 Objectives and background

The original objective of the MEMOIR Project was to demonstrate the applicability and integration of advanced multimedia systems in R&D-based enterprises (i.e. those which invest a significant amount in R&D towards the generation of new products) in order to meet the following business objectives:

- increase the productivity of researchers within R&D-based companies;
- improve access of sales teams to critical information right up to the moment of sale (the 'golden moment');
- improve access of technical specialist staff to critical information in their support of sales teams;
- enable teams and communities to arise within a research-based company which would not be expected based upon geographical distribution alone.
- For R&D-based companies, the provision of effective means for research staff, technical support specialists and sales staff to access and share all forms of technical information in ways which are relevant to their particular needs was, and still is, a particularly acute problem. Increasingly, the information that such staff generate and consume is multimedia in nature.

The ways in which staff within such companies use information differ considerably depending on their function. Researchers typically need to browse and search large numbers of documents prior to selecting those they wish to read, they need to access

data held in databases (both structured and unstructured), and they need to access information from experiments which can include alphanumeric data files, digitized images, video, and sound. Technical support staff and sales teams need to access the same information, but typically in a way which leads as rapidly as possible to particular information sources. Sales staff also need to be able to tailor the presentation of their products to the exact requirements of the customer. In addition to generating, sharing and using internally generated information, all three groups must acquire similar information from other parts of the company and from outside. Thus the information technology that supports this working environment *must* be multimedia-based, and must also enable different individual users and different categories of users to access the information in an appropriate manner. This tailored access must be provided as a combination of two features:

- customised information structure (e.g. different sets of hypertext links applied to the same information resource, depending upon the user);
- customised user interfaces (e.g. a researcher needs an interface which enables and encourages browsing and serendipity, whereas a member of a sales team needs an interface which provides rapid access to essential information as easily as possible.);

As the mountain of documents (whether paper- or computer-based) accumulates, the ability of an individual to hold all relevant indexes to the information in their head declines and vanishes. People therefore devise ways of organising the information on their computer systems to help them retrieve relevant information on demand. The typical situation in many large, geographically distributed, R&D-based companies is that 'islands' of information grow up which are physically accessible to others but, due to the personal and idiosyncratic nature of the 'organisation' of the information, are of little or no use to others. In addition there is the potential for massive duplication: small companies may lack the resources to carry out complex or expensive procedures once, but large distributed companies will unknowingly do things over and over again. Either way there is an endemic inefficiency caused by the inability to locate and share vital information.

Corporate Memory[1] relies almost entirely upon people knowing who in the organisation holds a piece of information (or knows where it is). For electronically stored documents, current practice relies heavily upon human recollection that the relevant report is "in such-and-such a directory", and is called "something like report.doc". For documents not yet stored electronically (the vast majority of historical information), current practice is to rely upon human recollection for approximately where, in an office full of filing cabinets, the document is to be found. This is followed by a search strategy based on pattern recognition using algorithms such as "it has a blue cover" (80% probability).

Within research-based companies such as Unichema and Glaxo Wellcome, four essential information management requirements were identified:

1) recording results from experiments;
2) running simulations;
3) providing information retrieval facilities for a variety of information needs;
4) identifying other people in the company who may be relevant to a particular line of work.

[1] The sum of the knowledge acquired and retained by individuals, past and present, within a company

On average, projects are repeated and products are re-invented every seven years within the chemical and pharmaceutical industries. There is also a constant risk that research work will ultimately be a waste of effort and resources, since it is not possible to know in advance whether a particular chemical or drug will be a) feasible, or b) profitable. There is therefore a need for a general mechanism to support the inventive process in order to reduce or eradicate these two problems, thereby increasing the productivity of researchers significantly. The requirement is not to be able to guarantee with absolute certainty the success of a project, but instead to avoid wasting time and resources on work which a part of the organisation "knows" will not succeed.

Other classes of workers who will benefit from having a similar support mechanism are the technical specialists and sales staff who need to be able to target their sales to the particular requirements of individual customers. It is a relatively easy task to manufacture many chemical intermediates—competitive advantage is gained by being able to inform the customer how to use these basic products in a particular manufacturing environment or production process (i.e. requirements 3 & 4 above).

A support system which will provide this competitive edge must therefore provide tailored access to the same basic information resources for different classes of users—the researcher/inventor, sales and technical staff all have different requirements of the same information. Researchers need to be able to browse, search for information and discover facts and relationships on the periphery of their specific area of interest; the sales and technical staff are not in a position to browse and discover information that may be relevant for a particular customer, instead they need to be directed to that information as efficiently as possible.

It is not enough to provide simple multimedia presentations, nor is it sufficient to leave the users to browse randomly through a potentially huge corpus of information. Mechanisms are required to support users in their quest for information, providing guidance based upon both manually defined relationships between items of information, and also on the activities of others who have carried out similar searches in the past. This guidance can also be used to bring together geographically isolated workers who are nonetheless involved in similar activities.

2.3 Changes of context

Since the start of MEMOIR there have been several significant changes in context for the project.

- Most significant has been the widespread take up of internet technologies and the World Wide Web. Glaxo Wellcome had embarked on an intranet implementation by the beginning of MEMOIR. However, Unichema's decision to implement an intranet occurred after the project start. The adoption of intranets by Glaxo Wellcome, Unichema and the majority of other organisations perceived to be likely users of MEMOIR-like systems resulted in an early decision by the project partners to change the technological basis for the project. It was decided at that point that the MEMOIR prototype would be implemented in an intranet environment.
- During the course of the project several products and research projects emerged that embodied some of the ideas and concepts being addressed in the MEMOIR project. Products include Alexa's *Alexa,* IBM's *Web Browser Intelligence (WBI),* Autonomy's *Agentware i3,* TDI's *NetAttaché Pro,* and Zuno's *Vrisko.* <<<text removed in public version>>> All of these developments have some overlap with, and some fundamental differences from, MEMOIR. The consensus of the project partners is that they add weight to the validation of the MEMOIR concept.

- Over the last 12 to 18 months, the concept of Knowledge Management has seen increasing levels of adoption by major companies worldwide. The interest and take-up of Knowledge Management practices has been particularly prevalent among the innovation-led companies typified by those in the chemicals and pharmaceuticals sectors, and identified as likely end users by the MEMOIR technical annex.
- During the project, Unichema was acquired from Unilever by ICI. While this presented opportunities for dissemination of MEMOIR to a wider corporate audience, the transition period had a negative impact on the level of Unichema's participation in the project.

2.4 Achievements and results

2.4.1 Functionality of final MEMOIR prototype

Agent-based services

Table 1 describes the MEMOIR user services that have been implemented. The set of functionalities in Table 1 are encapsulated in software agents that communicate with and respond to the MEMOIR message format.

Service name	Input	Output
Who else has seen this URL	One URL	An ordered set of users + a score for each user
Find trails like this	A set of URLs	A set of users' trails
Find the most common URLs in these trails	A set of users' trails	A set of the most common URLs ordered by frequency of appearance
Show trails	The user who makes the request (implicit)	The trails of the user who made the request
Find Similar people by trail	A set of URLs	An ordered set of similar users + a score for each user
Keyword extraction	One or more URLs	An ordered set of keywords
Find documents by keywords	A set of keywords	An ordered set of URLs + a score for each document
Find people by keywords	A set of keywords	An ordered set of users whose trails contain the keywords of the input
Suggested reading	A set of URLs	A set of URLs

Table 1: Agent-based services

Link creation

Facilities have been implemented to enable the generation of hypermedia links which are stored in linkbases independently of the underlying documents.

Dynamic hypertext mark-up

Multicosm's Webcosm product has been integrated into the prototype MEMOIR system to provide dynamic hypertext mark-up of browsed documents. Using this capability, predefined words or phrases occurring in any document browsed via the MEMOIR proxy server can be automatically provide links to related documents.

Persistent query support

Long-lived or persistent queries initiated in the MEMOIR environment may return results when the user who launched them is no longer logged on to the MEMOIR system. The prototype MEMOIR system provides a mechanism for trapping the results of long-lived or persistent queries for later delivery to the initiating user.

Customisation according to user profiles

A mechanism has been implemented to support a fixed set of user profiles, for example R&D and salespeople. In the prototype MEMOIR system, trails and links can be tagged with the type of user who creates them. This information can be used by all of the MEMOIR agents to provide exclusion filtering (by user type) of returned results.

Webcosm Linkbases™ can also be used to provide customised link presentation relevant to different user types.

Trail editing and annotation

Facilities have been implemented to enable the publisher of a trail to delete, edit and annotate trails.

Guided tours

Sequences of document identifiers can be useful in guiding a user through an information space with which he or she may not be familiar. Facilities have been implemented to enable a set of URLs to be published as a guided tour, which implies a sequential relation between the documents.

2.4.2 Tangible results

The implemented MEMOIR system consists of the components below. A brief outline of the function each component performs is given below. <<<text removed in public version>>>

- message router
  forms a lightweight binding between components;
- user interface
  enables access to all MEMOIR functions via a WWW browser;
- persistent storage for links and trails (Linkbase and trailbase)
  most functionality involves the sets of trails and links stored in the ITASCA Object Database;
- software agents
  a set of dedicated and sometimes interacting software modules performing MEMOIR functions;
- authentication facilities
  to ensure that users' concerns about published data are addressed, and their privacy of this data respected;
- an HTTP proxy
  enables both tracking of user behaviour and on-the-fly document markup.

The system design showing the relationship between these components is shown in Figure 1.

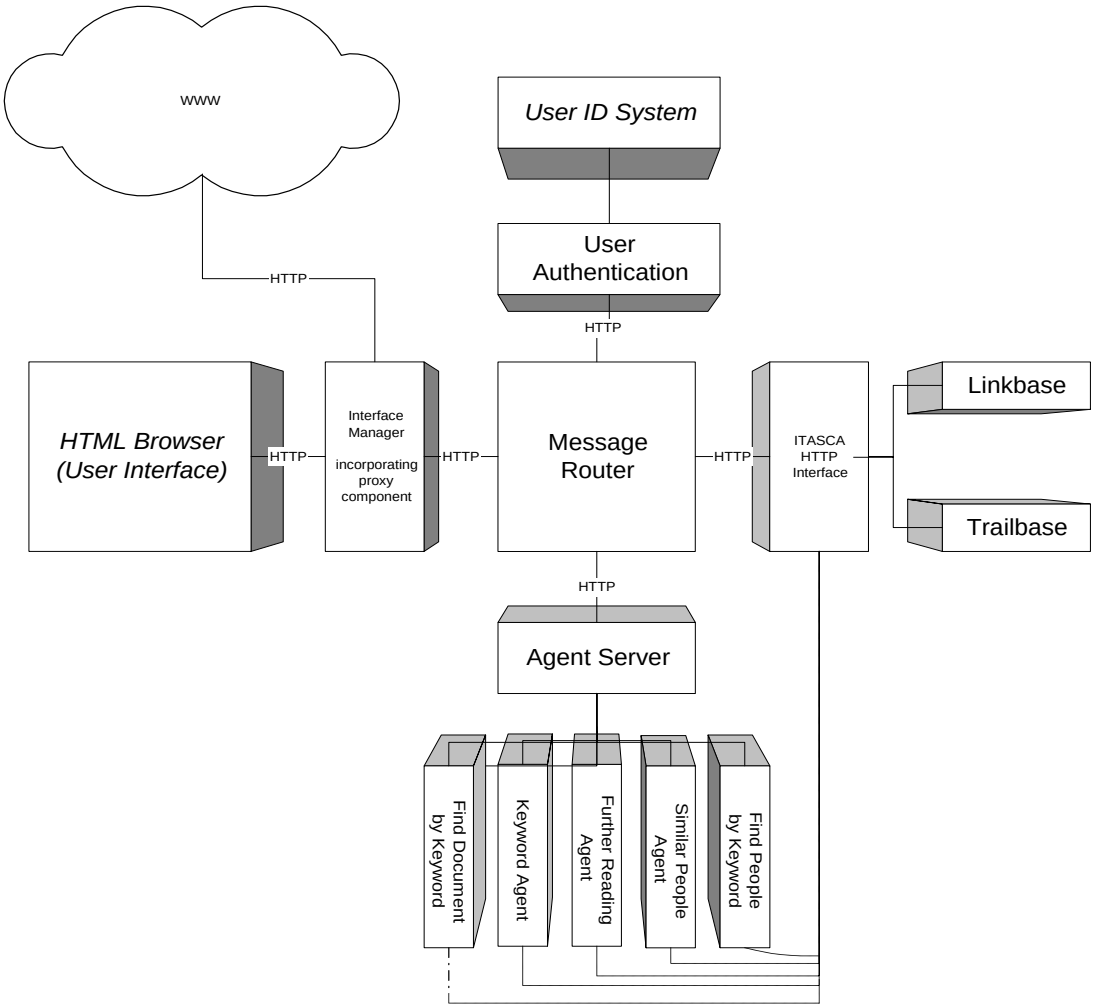

Figure 1: MEMOIR System Architecture

2.4.3 Intangible results

The intangible results achieved by the project partners through their participation in MEMOIR are likely to have a greater impact on the businesses concerned than the tangible implemented system.

MEMOIR has enabled the partners to acquire valuable knowledge which can be categorised as follows.

- Vision and principles
- User-centred design
- Technical development process
- User evaluation
- System integration
- Management
- Promotion and market positioning

The knowledge acquired in these categories is summarised below.

Vision and principles

The original vision and principles underlying MEMOIR, namely:

- that it should account for the people and the roles that these people fulfil in an organisation;
- that it should be open, especially to the incorporation of new technologies;
- that it should be distributed; and
- that it should address the reuse of information for different purposes

have been shown to be implementable in the context of what is now seen as a system capable of aiding the management and exploitation of knowledge within an organisation.

The partners have acquired a detailed understanding of the practical implications of the vision and principles outlined above in the context of their own business.

User-centred design

The importance of investing in a user-centred design approach to a "people-centric" system such as MEMOIR has been strongly reinforced by the project experience. In the project there has, to some extent, been a misalignment between the largely technology driven development process and the user-centred approach to evaluation. Given the technological immaturity of the MEMOIR concept at the outset, a technology driven stream for early exploration of possibilities and the implementation of prototypes was necessary. However, in retrospect the timing of the plan for user evaluation was over-ambitious and probably distracted the partners from effectively increasing the level of user participation in the design and development process.

The relatively immature design of the MEMOIR user interface has been a major constraint on an effective evaluation process. User feedback, particularly from Glaxo Wellcome, has indicated that the user interface has been perceived by many as a barrier to effective use of the MEMOIR services.

An important lesson for developers of user interfaces for systems like MEMOIR is that the use of terminology can be critically important in determining the usability of the interface. In hindsight, the current MEMOIR interface uses terminology that is too abstract for most users. There is a trade-off between "abstract and universal" and "precise and user group specific". In addition to the need for the right terminology to be employed, it is important for user acceptability that user expectations are met for a look and feel that is contemporaneous with their existing desktop environment.

Technical development process – key lessons

As described above, there was insufficient user participation in the design and development process and an over-ambitious emphasis on providing polished, robust systems to users for evaluation in the workplace.

Key technology design principles (component-based development and asynchronicity of component interaction) were established early in the project. However, these well-conceived principles turned out to be more difficult to enforce (due to expediency at various points in the project) than envisaged. There still remain several interdependencies between components that make configuration management a more complex task than it should be. The MEMOIR prototype is now largely asynchronous in

the way that the components interact. However, this had to be retro-fitted rather than being developed-in from the outset.

The issue of configuration management and version control for co-operative software development projects is one that deserves special attention. The system within MEMOIR now works well, but the effort required to implement the system and to make it a natural part of the way in which the developers (at seven different geographical locations) worked together was considerable. The time and investment required to make such a system work should not be under-estimated; the results obtained were a remarkable testament to what can now be accomplished via sophisticated use of Internet facilities.

The importance to an effective development process of getting the configuration and version management processes right is fundamental. The MEMOIR project has resulted in an undeniably complex software system that is able to be continually developed in a controlled manner. This achievement was critically dependent on the configuration and version management system underpinning the development process.

Vital attributes of the MEMOIR configuration and version management system were the central role of the PAC in the process, and the use of web sites at the PAC and developer companies to maintain consistency and instant access to latest component releases.

The overhead of developing a user interface capable of operating in a consistent way on a large number of operating system/browser combinations was significantly underestimated. In practice, the experience of developing a widely deployable Java-based user interface has been a long way from SunSoft's marketing proposition of universal portability.

User evaluation

It is acknowledged that the plans for user evaluation on increasingly wide scales within the user companies were over-ambitious. For users outside a typically small, geographically localised group of pioneers, user evaluation is only practically possible with very mature, product quality software that has very carefully designed and trialled user interfaces. In retrospect, the perceived need to achieve modestly widescale user evaluation led to the introduction of insufficiently trialled software.

The expectation that the developers would essentially have only one opportunity to gain any significant commitment from users to the evaluation process was strongly borne out by experience. It is very easy to demotivate users and nearly impossible to re-recruit them without effectively relaunching, under a new identity, the software for evaluation.

In seeking to have users evaluate software like MEMOIR in the course of their normal work activities it is vital that there is a well-understood, and explicit, alignment of work activities and the facilities/services provided. In other words, their motivation to participate must be extremely carefully thought through. There should be a recognised distinction between the early users and the eventual broad-based user population.

It is important (and acknowledged as difficult) to gain an early understanding of the critical mass of users and captured information necessary to create a perception of value in the use of a system like MEMOIR. This is often referred to in the literature of recommender systems as the 'incentive problem' or 'cold start problem'. A key response from some users may be paraphrased as "The concept is exciting, but I hardly ever learn of similar people and I am only rarely provided with any suggested reading. Why should I invest in publishing to MEMOIR if I get little of value in return?"

Very early in the project, the risk was identified that MEMOIR might be regarded as a "Big Brother" system spying on the activities of staff. The publishing mechanisms were designed to take account of this.

System integration

The integration of a system like MEMOIR with the existing infrastructure of a large user organisation is a task requiring considerable planning and co-ordination across all relevant parties as well as technical implementation effort. This takes a significant time and if the experience of MEMOIR is typical, results in questions where existing policies do not yield ready answers. Many of the issues and problems relating to the system integration task were attributable to the requirement to account for the security policies at end-user IT environments. The resolution of these issues for MEMOIR involved unplanned development work in some cases.

Concerning the technical aspects of the integration problem, the openness and flexibility of the MEMOIR architecture made it relatively straightforward to interface MEMOIR services to existing authentication and directory services.

Management

With hindsight, the scope of MEMOIR was too great for a single two year project. The project plan implicitly assumed that a technology driven development project implementing radically new and untried ideas could successfully transform itself into a user-centred evaluation of the resulting systems. A better structure would have been to have two separate projects: the first focused on IT development, and the other focused on user evaluation of (near) product quality systems.

Some valuable management tools have been employed in MEMOIR. Of particular note are the following.

- Task forces. At various critical points during the development process, very successful use was made of developer task forces assembled in a single location to address a specific, well defined problem. The task forces were constituted of appropriate representatives from the consortium, and lasted from two days to two weeks.
- Management by embarrassment. During the later stages of development and the transition to the project's evaluation phase it became apparent that some problems were remaining nearly, but not fully, resolved for several weeks. This issue was discussed at a PST meeting. A problem management process involving the publication of (approximately) weekly bulletins identifying issues together with their status and *executive level* ownership was agreed and successfully implemented.

Promotion and market positioning

The major contributions of the two highly motivated University participants in MEMOIR has resulted in wide reaching, high profile dissemination of the MEMOIR experience and results in a research community context.

The potential importance of internal dissemination of a project like MEMOIR within a large end user company should not be under-estimated. The Unichema team, in particular, has been highly successful in their use of MEMOIR as an exemplar for knowledge management system implementation projects both within Unichema and more widely in ICI.

The MEMOIR project has had to evolve its technical focus to take account of the growth of the Web, and its marketplace positioning to take account of the emergence of knowledge management as a high profile trend among the potential customer base for MEMOIR-like systems. In the rapidly evolving context of the IT industry, R&D projects such as MEMOIR must maintain a constant watch for relevant changes and be ready and able to review the positioning and promotion of the project in response. As is the case with MEMOIR, project partners must also be open to the possibility of profound and fundamental changes to their positioning partly driven by project experiences.

2.5 Significance of achievements and results

2.5.1 MEMOIR partners

The MEMOIR project has had a strong effect on the market positioning of both IBEX and Multicosm. During the lifetime of the project, participation in MEMOIR and the results emerging from the project have been significant contributing factors to both IBEX and Multicosms' positioning of themselves as suppliers of technology to support knowledge management processes within large enterprises.

MEMOIR has provided an important learning vehicle for both Glaxo Wellcome and Unichema.

For Glaxo Wellcome, the most significant impact is on information systems strategic planning, in particular with respect to high potential projects.

The key areas of MEMOIR influence are on the following.

- The evaluation and appraisal of novel, high potential, high risk IT components for Glaxo Wellcome's IT infrastructure.
- Glaxo Wellcome are developing processes for systematically reviewing where the enterprise should seek to be in terms of IT and what IT can enable.
- The development of improved sharing and access to the company's knowledge assets.

<<<text removed in public version>>>

For Unichema, participation in MEMOIR, and the issues that the project has raised, has helped to shape the corporate approach to knowledge management. In particular, the MEMOIR experience has helped Unichema to understand approaches to the introduction of IT systems to support knowledge management. Following Unichema's acquisition by ICI, Richard Miller, Unichema's R&D Director, has been able to draw on the Unichema experience in MEMOIR to help shape ICI's evolving strategic adoption of knowledge management. <<<text removed in public version>>>

The PAC is in the process of developing a more commercial position than it has previously operated concerning the exploitation of knowhow and experience required in R&D projects. The PAC now has a deep understanding of the issues associated with developing and implementing people-centred knowledge management tools. MEMOIR has provided the PAC with experience of co-ordinating a multi-site development process for a complex software system and of implementing complex systems at user sites. This experience will form a basis for the PAC to develop new service propositions.

Participation in MEMOIR has significantly strengthened Athens' and Southampton MMRG's research profiles in the agents and distributed people-centric systems fields.

2.5.2 European industry

MEMOIR has had significance for European industry outside the MEMOIR partners. The experiences of the partners have been widely disseminated, and the project partners believe that MEMOIR provides a valuable exemplar for a new class of people-centred system. MEMOIR has been promoted as an exemplar of how IT systems capable of supporting knowledge management processes can be constructed.

MEMOIR has had particularly strong promotion within the previous and current parent organisations of Unichema (Unilever and ICI respectively). Several other subsidiaries of Unilever and ICI have also been the recipients of presentations on MEMOIR. Both Unilever and ICI are among the early adopters of knowledge management practices.

2.6 Commercial exploitation

Both IBEX and Multicosm are strongly committed to commercially exploiting the results of MEMOIR by producing products aimed at the field of Knowledge Management The commercial exploitation plans have been strongly influenced by information gleaned in the process of disseminating the MEMOIR project. The exploitation plans are described in Section 3 of this document. <<<text removed in public version>>>

2.7 Management

2.7.1 Introduction

Unichema were responsible for the overall co-ordination of the MEMOIR project. The PAC under a management services contract with Unichema undertook day-to-day management of the project.

The following Sections describe the major management issues addressed in the project and some of the key lessons learned.

2.7.2 Major decisions

The process of preparing the technical annex provided the focus for most of the major decisions underlying the project. Subsequent to the preparation of the technical annex there were two further decisions which had a major impact on the approach taken in the project. These were:

- the decision to make the MEMOIR implementation based on http; and
- the decision to make the MEMOIR client as close as possible to requiring no configuring whatsoever.

The technical annex did not originally envisage an http-based implementation. During the project design phase a variety of alternatives for message passing (including CORBA/IIOP and proprietary systems such as IBM's MQ Series) were considered. http was selected as a widely supported standard and a lightweight protocol. The decision to adopt http was also compatible with widespread adoption of intranets by large enterprises worldwide.

The early vision for MEMOIR included client software that would be installed on the desktops of users. User needs for low maintenance costs in large user communities resulted in the decision to make the MEMOIR user interface a Java applet, with the client browser being configured by a simple proxy setting. While it met the needs of users, this decision exposed the limited portability of Java user interface code in practical applications using a multitude of browser / operating system combinations. Significant

development effort was needed to implement and maintain a version of the user interface which would reliably and consistently execute on even a restricted set of client browser / operating system platforms.

2.7.3 Major milestones

The first major milestone for the project was the completion of a pre-prototype demonstrator system in M4. The process of implementing a pre-prototype demonstrator proved to be highly valuable in providing a tangible focus that the partners were able to work jointly towards. This formed a strong basis for the formation of an effective multi-partner development team.

The second major milestone was the successful implementation of an initial MEMOIR system suitable for evaluation. In the technical annex this was scheduled for completion by the end of M9. In practice, the challenges associated with translating fuzzy visions of what MEMOIR should do into real software meant that the first MEMOIR installation was made during M15 at Glaxo Wellcome. This installation was suitable for evaluation by only a small number of IT professionals, not by business end-users. The slippage of the second major milestone reflects the over optimistic scope of the project discussed in Section 2.4.3 *Management*.

2.7.4 Management challenges

An important challenge for the project management team was to maintain the coherency of the project partners in working for common objectives despite rapidly changing business contexts:

- technological, e.g. the widespread adoption of intranets and the Web, and emerging MEMOIR-like tools;
- market positioning, e.g. the emergence of a knowledge management tools marketplace and Multicosm and IBEX's positioning in it; and
- Unichema's change of ownership.

The co-ordination of highly distributed software development was a significant challenge for the consortium. This was successfully achieved through the implementation of appropriate configuration management practices and supporting infrastructure, and through the co-ordinating role played by the PAC.

While the need to re-orient the project from being R&D-focused to being user evaluation-focused was identified and acknowledged by the partners, the challenge of actually accomplishing this re-orientation within the resources and time available to the project was not fully achieved.

2.7.5 Problem management

In the course of attempting to re-orient the project to focus on user evaluation, the need for an effective problem management process was identified. This was achieved by gaining executive level ownership of the process at each project partner and sign-up to a principle of "management by embarrassment". A weekly bulletin was issued by the PAC identifying problem issues, their executive level owners, due date, and status.

2.7.6 Management lessons

Many management lessons have been learned by the project partners in the course of MEMOIR. Some of the lessons worthy of note are as follows.

- Project scope. As discussed in Section 2.4.3 *Management*, it is now believed that the scope of MEMOIR was too great for a single two year project. With hindsight it may have been better to have embarked on a technology R&D project with very little planned exposure of resultant technology to end users for evaluation.
- Contingencies for unproven technologies. The importance of allowing for contingencies in plans to allow for the inevitable problems associated with immature technologies (e.g. portability problems of Java user interface software, or software features of the HP-UX operating system resulting in insurmountable difficulties in enabling multi-threaded MEMOIR components to be ported to HP) cannot be overestimated.
- Design and development co-ordination. Early experience of the development team working on a clearly defined task (i.e. the pre-prototype) was a vital instrument in creating mutual understanding of the aspirations and visions of the developer partners, and therefore provided a basis for opening up the channels of communication for the design and development process. The willingness and capability of a single organisation (here the PAC) to take on the responsibility of co-ordinating the design and development process has been a key element in the effective development process.
- Demarcation of ownership. A vital task for the co-ordinator of the development work is to ensure that development areas and the interfaces between them are clearly defined, and that the developer partners agree ownership of the responsibility for execution of required development work.
- Task forces. The use of task forces made up of personnel drawn from all relevant organisations and brought together at a single location to address a specific problem or issue has proved highly successful in the MEMOIR development process.
- Version control. The implementation and adherence to a version control system for the individual components of a complex software system such as MEMOIR has been a key element in enabling the efficient incremental development of new capabilities.
- Configuration management. Related to version control, it has been vital for system problem fixing to provide an infrastructure capable of specifying and recreating the precise system configuration tested at a particular time or installed at an end user location at a particular time.
- User evaluation. The experience of MEMOIR has shown that the resources, planning and organisational commitment needed for large-scale user evaluation are considerable (and were underestimated in MEMOIR). A potentially important lesson for future projects is that the organisational structures and culture of an R&D project that has an initial focus on technological development are ill-suited to supporting evaluation. If the transition between these different focuses is to be made successfully, significant planning and commitment is required.
- Problem management. While it is not seen as an instrument which should be deployed from the outset, a problem management system such as the MEMOIR "Rocket" has been demonstrated to be a valuable option in ensuring rapid attention is given to potentially lingering problem issues. The Rocket took the form of a weekly bulletin identifying issues, executive owners, due date, and status. A key element of its success was sign-up, at the executive level, to the concept of "management by embarrassment".

- End-user participation. While Unichema has been able to successfully justify following up the MEMOIR project with further evaluation of the prototype system after the end of the project, Glaxo Wellcome has chosen not to do so. Glaxo Wellcome has identified that its mechanisms for evaluating and exploiting some types of novel IT, especially generic information management applications, are not as fully developed as those in areas that are closely linked with specific business activities.

2.8 Resource utilisation

<<<text removed in public version>>>

2.9 Deliverables

<<<text removed in public version>>>

2.10 Publications and dissemination

There has been sustained academic and commercial dissemination activity throughout the project. MEMOIR papers have been published in the proceedings of the following conferences:

Hypertext 97: The Eighth ACM Conference on Hypermedia and Hypertext

PAAM 98: The Third International Conference and Exhibition on the Practical Application of Intelligent Agents and Multi Agents

WWW-7: The Seventh International World Wide Web Conference

WETCE '98: IEEE Seventh International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises

<<<text removed in public version>>>

The MEMOIR prototype system has been demonstrated at the PAAM 98 conference, the Hypertext '98 conference,the MEMOIR seminar, and it will be demonstrated at the MULTIMEDIA '98 conference in September.

Dissemination in a commercial context has been achieved through approaching relevant IT consultancies and opinion formers such as Meta Group, Patricia Seybold Group, Gartner and Delphi, exhibiting MEMOIR at commercial events, organising a seminar to disseminate the project in the context of Knowledge Management, and through raising the project's profile within the partner corporations and their parent organisations.

In the lifetime of the project there has been significant commercial interest in Knowledge Management (KM), a field which embodies many of the principles underlying MEMOIR; for the purposes of exploitation, it is therefore appropriate to set MEMOIR in the context of KM. This was achieved through a seminar, co-hosted by Ibex and Sun Microsystems in London, which aimed to overview the MEMOIR project, position the work in the field of KM, and explore the uses and applications in R & D environments.

Unichema's relation with their former parent organisation Unilever, and their new parent organisation ICI, has provided an opportunity to disseminate the project in both these organisations.

MEMOIR has been exhibited by the project partners at a wide range of commercial events, including Business Intelligence conferences, a Government Computing

exhibition, Computer 97 and Computer 98 exhibitions and at seminars held by the British Computer Society and by Unicom.

The project has made a submission for inclusion on the EC's PROSOMA web site.

3 Technology implementation plan

The technology implementation plans for the exploitation of results from the MEMOIR project by each of the commercial technology suppliers in the project are overviewed in Sections 3.1 and 3.2 and for each of the end users in Sections 3.3 and 3.4. <<<text removed in public version>>>

3.1 Ibex Implementation Plan

The Ibex technology implementation plan is based on integrating the MEMOIR project results into a range of Ibex Knowledge Management Systems (IKMS). Ibex is developing these products by integrating the experience from a number of R & D projects into its products and applying them to current business. Through consultation with leading IT consultants, Ibex has identified that a Knowledge Management System should:

- provide tools to support the integration of both information and experience to build knowledge; and
- provide tools to organise, structure, store, index and retrieve that knowledge; and
- exploit existing computerised information sources such as electronic documents and relational databases; and
- be able to handle and execute complex structures in real time; and
- exploit industry standards, taxonomic conventions and high level end-user oriented interface specifications.

The foundation of IKMS is a Base Product providing core management services including persistence, knowledge representation, reasoning and workflow. (see Figure 3) The Base Product is integrated with domain modelling and knowledge management tools many of which are provided by Ibex partners. The IKMS Base Product is integrated with production environments through Knowledge Management Applications that are also provided by IBEX and its partners. The Base Product is likely to be based on the existing MEMOIR architecture, although the message router, agents and ITASCA interface would be integrated into a single component. A user application will also be developed to include the functionality of the current interface, whilst being more user friendly and supporting more functionality than the current Web browser.

The IKMS services will be available over the Intranet / Internet, through corporate Web applications servers, providing an interface of the IKMS to other applications. The IKMS will manage "Knowledge Objects" on behalf of these applications; these include the Link, Trail and Profile objects from the MEMOIR project along with other objects such as queries, decision trees and business models.

IBEX will be applying the MEMOIR results to develop custom Knowledge Management Solutions using the IKMS Base Product as the starting platform. Initially the custom developments will be focussed on the application areas of Chemical Synthesis and High Throughput Screening, and Risk Management for futures and Options Trading. An IKMS custom solution is built in four layers. The IKMS Base Product provides a basic knowledge acquisition mechanism at the first layer. Providing a standard database adds domain-specific content to create the Knowledge Base. The top two layers add customer-

specific rules and procedures to integrate the Knowledge Base with the customer specific production environment. Tools to support database searching, rulebase specification and decision support are added at each level to build a complete IKMS solution.

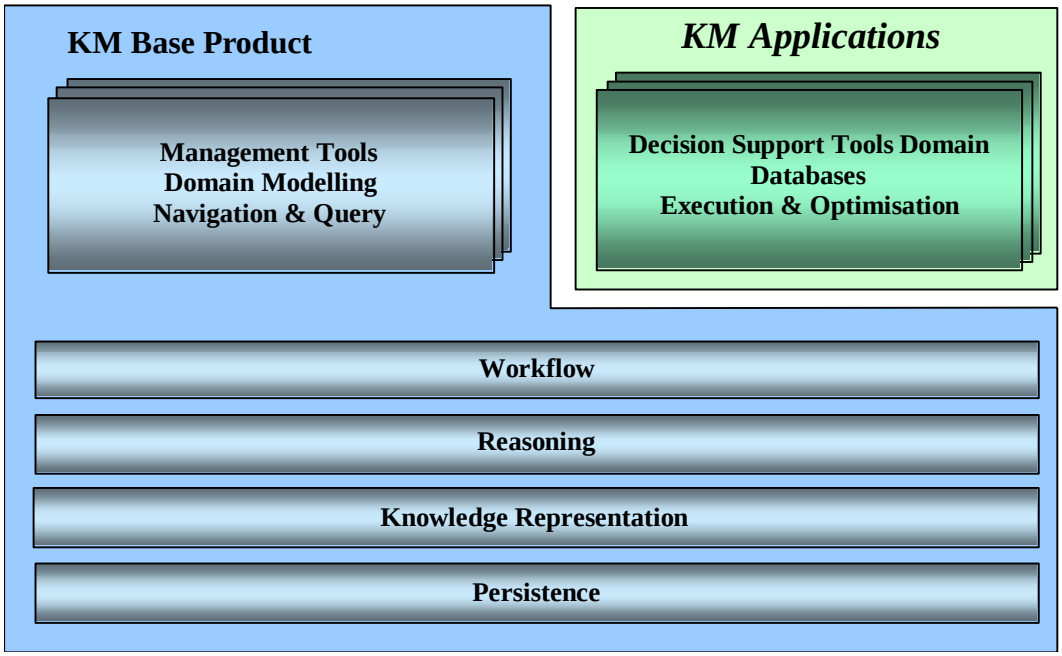

Figure 3: *Ibex Knowledge Management System*

3.2 Multicosm Implementation Plan

When discriminating between emerging 'Knowledge Management' offerings it is possible to make a broad distinction between systems aiming to classify the knowledge assets of an enterprise, and those designed to deal automatically with unstructured information. In these terms, the MEMOIR project falls into the latter category, and given the relative lack of existing established competition, rapid exploitation of successful project results has a high chance of success.

In parallel with participation in MEMOIR, Multicosm has been investing in complementary technologies, including linguistic technology, and can now offer a total solution for automatically managing the relationships between dynamically selected documents. The company's Open KnowledgeWare product set is based around a user requirement of asking two questions when looking at a document or a set of documents:

- What else is relevant?
- Who else knows about this subject?

These questions also underlie the functionality available in the final MEMOIR prototype system. With the same rationale that led to the use of a Web-based interface to MEMOIR, Multicosm note that key technologies must operate in conjunction with the Web browser, normally in an Intranet environment.

The Multicosm set is based on the following elements:

- Dynamic link insertion (Webcosm)
- User Document History Processing (TrailBase Agents, etc.)
- Automatic Document Linguistic Analysis (Themes)

The design objective for Open KnowledgeWare products is to intrude as little as possible into user's consciousness. This is compatible with feedback from the MEMOIR evaluation process, where some users reported that the interface to MEMOIR was overwhelming and a barrier to using the functionality offered. The Multicosm products will provide a simple set of "buttons" providing access to key knowledge functions. Ideally, system objects should be transparent but MEMOIR experience suggests a conscious user action to "publish" a trailbase is necessary to preserve a sense of privacy.

In order to further enhance the product set, Multicosm has started to partner with leading suppliers of sophisticated search engines and document management systems. Of particular interest are search engines which represent user queries, and documents, semantically and can therefore provide better quality information in response to user searches; the same principle is also applicable to link insertion, and Multicosm is pursuing this to improve the quality and relevance of the link insertion process.

Overall, Multicosm is confident (based on feedback from industry experts) that its knowledge management technologies are state-of-the-art and have benefited greatly from the experience gained during the MEMOIR project.

3.3 Unichema exploitation plan

Unichema's participation in MEMOIR, and the issues that the project has raised, has helped to shape the corporate approach to knowledge management. In particular, the experience of participating in MEMOIR has helped Unichema to understand approaches to the introduction of IT systems to support knowledge management. Following Unichema's acquisition by ICI, Richard Miller, Unichema's R&D Director, has been able to draw on the Unichema experience in MEMOIR to help shape ICI's evolving strategic adoption of knowledge management. Richard has recently been appointed as Director of Knowledge Management and Sustainable Development within a new company formed by a merger of Unichema and ICI's Surfactants business.

Unichema will continue to evaluate the MEMOIR prototype system after the end of the project. No end date has been set for the completion of this evaluation.

3.4 Glaxo Wellcome exploitation plan

For Glaxo Wellcome, the most significant impact is on information systems strategic planning, in particular with respect to high potential projects.

The key areas of MEMOIR influence are on the following.

- The evaluation and appraisal of novel, high potential, high risk IT components for Glaxo Wellcome's IT infrastructure.
- Glaxo Wellcome are developing processes for systematically reviewing where the enterprise should seek to be in terms of IT and what IT can enable.
- The development of improved sharing and access to the company's knowledge assets.

4 Technical report (for publication)

The material presented in this section is based on the content of a paper published in the proceedings of the WET ICE '98 conference (IEEE Seventh International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises. This material has

been updated to reflect the status of the technical developments, and the end-user evaluation feedback at the end of the project.

4.1 Introduction

Effective access to documents and effective collaboration between researchers can be crucial to the competitiveness of large companies. MEMOIR, which stands for Managing Enterprise-scale Multimedia using an Open Framework for Information Re-use, addresses this problem by providing the user with extra assistance in these tasks in return for minimal extra effort by the user. Unlike other recommender systems, MEMOIR recommends people as well as documents. The application areas that we focus on pertain to large distributed organisations that have to manage diverse sources of information and that rely heavily on their research and development function. The system described in this paper is currently being trialled by users in two such organisations.

The MEMOIR architecture is an evolution of the Distributed Link Service (DLS, (Carr et al., 1995)). While the DLS, like other open hypermedia systems, treats hypermedia links as first class objects, MEMOIR additionally promotes another kind of object: the trail. A user's trail is the set of actions on documents that they have visited (such as opening the document) in pursuing a certain task. By matching trails, we match users. MEMOIR lets the user ask questions such as `who else has read this document?' and `what else should I read?'.

MEMOIR can be thought of as a navigation assistant, and it borrows from DLS the idea of implementing this agent as an HTTP proxy. With the proxy on a separate machine, this approach has minimal impact on the user's desktop (perhaps requiring a routine configuration change) and it steals no cycles. Trails, like links, are persistent, shared objects, hence multiple proxies communicate with a logically central (but perhaps physically distributed) database to store and process them.

The notion of supporting collaborative work by hypertext technology dates back to the very first ideas of hypertext, as expressed by Bush's `Memex' (Bush, 1945). He envisaged trails (also called paths (Furuta et al., 1997; Zellweger, 1989) or footsteps (Nicol et al., 1995)) in order to provide a mechanism for finding a user's personal information but also in order to allow these trails to be available to other users. Reading (i.e. opening and viewing) a document might be the most common type of activity in a trail, though there are many other actions users perform; for example, bookmarking, printing or mailing a document could indicate a special interest. In the system described in this paper, we focus on `open document' actions and support collaboration by providing agents that perform intelligent analysis on the documents being viewed by users.

The architecture is presented in Section 4.2. Section 4.3 describes experience we have had with the framework, both as a research vehicle and at the end user sites. We report on related work in Section 4.4, followed by concluding remarks.

4.2 The Open Architecture of MEMOIR

In open hypermedia architectures such as the Distributed Link Service, information about the links between multimedia documents is stored and managed separately from the documents themselves, which remain in their native formats; hence links are objects in their own right, stored in link databases (or linkbases). The criteria for open hypermedia have been extensively discussed at Open Hypermedia Systems Workshops, the last of which was held at the 1997 ACM Hypertext Conference (Wiil, 1997); this group is also developing a protocol for open hypermedia systems.

The MEMOIR project (Hill et al., 1997) extends DLS to support collaboration by storing and processing trails. It introduces new system components: a distributed object-oriented database, a framework for agents, and a communications message router. Figure 4 gives an overview of the different components and their interaction. Components displayed in dashed lines are third party components; interactions using other than HTTP as a protocol are drawn using dashed arrows.

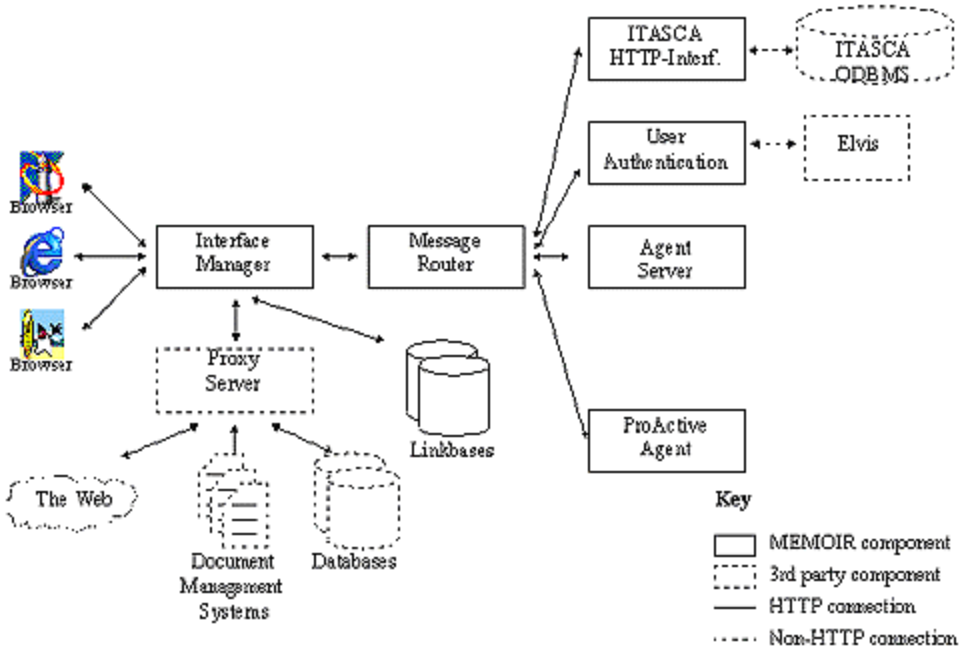

Figure 4: The MEMOIR system components

The key component is the message router which acts as a hub with which other system services register. A system may employ more than one message router. Its design is based on Microcosm's filter model (Hill et al., 1993). Any component makes its services available by registering details of its services with the message router, which in turn can route particular requests of other components to the newly registered one. Messages for querying all currently available services are part of the set of message router services. The model of servicing a component allows the system to be dynamically tailored to specific needs. New services can be easily added, and multiple providers for the same service are possible. Multiple linkbases, for example, could be managed by one component providing services for link manipulation; they could equally be managed by different components providing the same services.

Distribution is addressed at various levels. Firstly, all components can be run on different hosts. Secondly, the use of HyperText Transfer Protocol (HTTP) allows components on heterogeneous platforms to communicate. Thirdly, the underlying object-oriented database system ITASCA (see below) can be distributed and by that provide the means for sharing data.

Standard Java enabled Web browsers provide the user interface to the MEMOIR system in order to minimise installation and maintenance overhead. The base panel shown in Figure 5 is a Java applet that provides the user interface. The only configuration necessary is to set the interface manager as a proxy server and to connect once to it in order to log in. As browsing activity progresses the trail display in the user interface is

updated, and users may select documents from the trail as input to MEMOIR functionality. For instance they might chose to select some trail marks, press the button named `suggested reading' and the results would be returned in the lower list box of the base panel applet.

The applet serves as an interface to the following functionality:

- link management, including creation of links and link navigation;
- trail management, including `finding similar users', querying for suggestions of other documents to be read' and trail creation;
- keyword services, such as extracting keywords from a document as well as finding users or documents based on keywords.

The interface manager behaves as a proxy server and by that is able to listen to a user's requests and commands. It can itself be configured to use another proxy (e.g. cache or firewall). It is via the interface manager that users view the Web, document management systems and also linkbases. Thus, existing data can be re-used and is not stored within the system (though meta data such as a document's keywords are extracted and stored in the data repository, ITASCA).

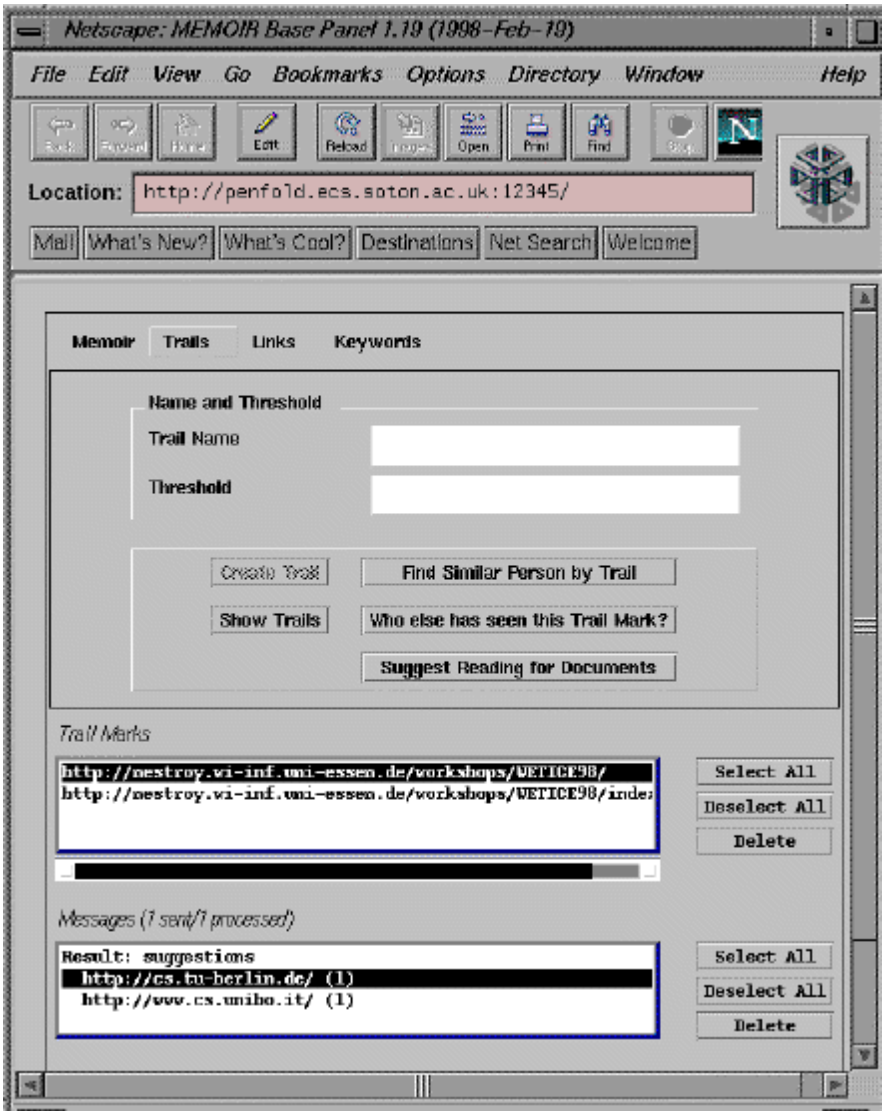

Figure 5: Basic control panel

An authentication component is responsible for managing user authentication. As far as possible this component relies on existing user management systems. A dedicated agent server manages a set of software agents and provides different agent services that are essentially responsible for data mining and resource discovery tasks such as searching the Web (Pikrakis et al., 1998).

The ITASCA distributed object oriented database system (IBEX, 1995) is used as an underlying data repository. Additionally it provides the framework for a number of agents that are implemented directly in LISP, the data manipulation language of ITASCA. As HTTP is used as communication protocol an interface component to ITASCA is provided allowing the communication with the database via HTTP (see Figure 1).

The Hypertext Transfer Protocol is used as a communication layer and provides the necessary infrastructure for the components to inter-operate. Thus all components `talk' HTTP, albeit with MEMOIR-specific message semantics. We distinguish between two basic forms of communication, namely notifications and requests. These are expressed using POST or GET methods respectively, a convenient mapping which sits comfortably in existing Web infrastructure:

- Where a process simply wishes to provide notification that an event has occurred, such as the opening of a document, the POST request can be used to deliver a message of a particular type to the message router. The message router can in turn report the event to any processes registered to receive messages of this type using further notifications.
- For queries or requests where a reply is required, such as following a link, the GET request is used. As with POST requests, the message is passed on to the appropriate components, and all results can be packaged together to be returned to the calling component.

It has proven very useful to support an asynchronous model in addition to the synchronous client-server model imposed by the Web architecture. The notifications enable the system to exhibit proactive behaviour and avoid the need to maintain processes that poll for state changes of other components or remote objects. Some operations do not need to occur in real-time. For example, a trail is submitted to the system when a user logs out, and this does not need to be processed instantly and a result returned to the user.

The MEMOIR message format is based on the open, extensible structure used in an earlier open hypermedia project (Microcosm (Davis et al., 1992)). This format uses simple tag/value pairs and although there is a required set of tags for normal operation of the system, additional tags can easily be added to extend the information carried to suit new functions.

To allow asynchronous interaction with the user, we have extended the MEMOIR set of messages with additional tags. These include a query ID, a request ID as well as a flag expressing asynchronous behaviour; a query ID is composed of a component ID and a simple string or number that is managed by each component in order to keep a record of sent and received messages. This permits, for example, the user interface to group a user's requests and answers together, identifying those that refer to the same ID. If a result is returned while the user is not logged on to the system, it is stored in a message box within the ITASCA database.

4.3 MEMOIR Functionality

Although MEMOIR is primarily an extensible framework for supporting collaboration, basic functionality is required to carry out the tasks of helping a user to determine, when working with a set of documents:

- Who else is relevant?
- What else is relevant?

The MEMOIR system now includes a rich set of functions, some implemented as software agents, to support these and other user actions.

4.3.1 Trail-related functionality

As a user access documents, this activity is recorded in the MEMOIR user interface and a trail representing the recent history is cumulatively built up. The user may invoke several operations based on documents in this trail.

- Publish trail. This function will add the trail to the database, associated with the current user, associating that user with the documents in the trail.
- Who else has seen this document: This function will tell the user which other MEMOIR users have published trails containing this document, thereby registering their interest in the document.
- Find similar people by trail: By selecting a set of documents which are of interest the user can ask which other MEMOIR users have related interests, based on the selected documents.
- Suggest reading for documents. This function will bring the user's attention to documents which may also be of interest based on those selected. The suggestions are based on published user trails which contain the selected document, and builds up a set of documents based on other documents in these trails.

Trail editing facilities are provided in MEMOIR such that users can review, edit, remove, and annotate trails they have created.

4.3.2 Link-related functionality

A key component of MEMOIR is the ability to record associations between information, independently of the documents themselves. This enables different views, or contexts, to be provided for the same set of documents. Central to this notion is a linkbase, which is used to store associations between documents and terms.

MEMOIR users may:

Publish a link based on a specified anchor. Having identified a document which is of interest, a user may publish an association between a representative term (or link anchor) and that document.. When publishing a link users specify a context in which the association is valid.

Select a context in which to view all information to be presented. Before documents are presented to the user, hypertext links will then be applied to them dynamically according to the current context.

4.3.3 Keyword-related functionality

Keywords may be used as indexes to both documents and people. When a trail is published by a user, the system automatically extracts keywords from the documents referenced in the trail and stores these as metadata with the trail. MEMOIR users may:

- Request to see the keywords characterising a document, or set of documents.
- Request to see which other MEMOIR users are associated with a specified keyword or set of keywords
- Request to see which other documents associated with other MEMOIR users contain a specified keyword or set of keywords.

4.3.4 Guided Tours

Guided tours help a user to navigate through an unfamiliar domain in a logical way, exploiting the experience of others who have previously navigated the same domain. MEMOIR users may create guided tours, a sequence of logically related steps which dictate the order in which documents are viewed. Other users of the system may then choose to follow a tour corresponding to the domain they wish to navigate.

This functionality is based on work carried out at the Multimedia Research group of Southampton University, and demonstrates the extensibility and flexibility of the MEMOIR system., as well as representing another way to explout the trail information held in the system.

4.4 End User Evaluation

The MEMOIR system has been evaluated at Glaxo Wellcome and Unichema, and the process has generated valuable feedback for all the project partners; this will influence the product offerings from the commercial partners. The consortium acknowledged in the latter stages of the project that the initial plans for very large-scale evaluation had been over-ambitious. In the following sections the feedback relating to various aspects of the MEMOIR system is summarised.

4.4.1 Interface and Usability-related issues

The evaluation process has revealed that the user interface to the MEMOIR is, for some users, a barrier to the exploitation of the system's functionality. Many users have reported that the MEMOIR interface seems out of date when compared to contemporary desktop applications, particularly users who have had exposure to the latest Web technology. The unsophisticated appearance of the MEMOIR interface is due to the requirement to maintain consistency amongst a very wide range of client machines. Some users also maintain that creating a trail is a lot of work, involving retrospective selection of items in the trail history; this is very useful feedback, since the success of MEMOIR relies on users' publishing trails. Some users have expressed a dislike for having two browser windows open to use MEMOIR because of the amount of 'real estate' this requires on their desktop. The development team has carried out experimental work to minimise the interface's presence on the desktop, although the results have not been subject to end-user evaluation.

There is no clear consensus amongst evaluators on the naming of functions in the interface. Whilst some users have expressed strongly that the naming conventions are obscure, others have expressed that invoking functions is intuitive, as the buttons do what they say. The service entitled 'Suggest Reading for Documents' was noted as

misleading, because the name implies the service involves actively making suggestions to other users, rather than requesting additional relevant material.

It was generally felt that there was insufficient notification to users regarding the status of the system, particularly when problems occur. For example, users receive no notification when the system is restarting, and there is no notification when dropped connections have been re-established.

4.4.2 Privacy-related issues

The consortium recognised the potential 'big brother' reaction to the MEMOIR system from an early stage; this is why a user must actively create a trail rather than the system automatically recording document access activity. It is clear that a balance must be struck between ensuring that the user recognises that nothing is recorded without his or her consent, and minimising the effort required to store information in the database.

Some users have expressed privacy-related concerns, however. Users have questioned whether it is fair to want to know which documents their boss, or other members of their team, have accessed. Similarly, the 'Who-has-seen-this-document' functionality has been perceived in a negative way by some users, who feel that others may be able to check up on whether that person is doing their work.

4.4.3 MEMOIR functionality

The evaluation process aimed to determine whether MEMOIR functions operate as expected, and whether these function are perceived to add value. A number of specific points relating to this issue were made in the evaluation. Users were disconcerted that a 'create trail' request, which seems like a simple operation, took a long time to process. It was noted that a function to actively inform specific inform specific individuals in a team about a trail would be useful, rather than relying on potentially interested parties coming across the trail by generating a relevant request A similar sentiment for notifying users of interesting documents was expressed.

From the point of view of link-related functionality, some users noted that they had to be careful with the selection of context to view information; an inappropriate choice can result in a multitude of added links in a single paragraph which do not add any value to the document. On the positive side, there was a favourable reaction to the citation-style of presenting links. Users would have liked the facility to make their private linkbases public once sufficient useful information had been incorporated. The concept of links was well received although some people thought this was a duplication of functionality already available to them in Lotus Notes.

There was some confusion over the effect of adding further keywords to an existing query. Specifically, some feedback indicated that adding keywords appeared to widen the scope of the original query rather than narrowing the scope, as would be expected. The keyword-related functionality was deemed useful by many evaluators. An example is a report that one evaluator was able to find interesting sites once he had established what his boss was interested in by exploiting the keyword-based functions.

There was some feedback relating to MEMOIR's referencing of other system users. In some cases another user's identification was returned to a user in response to a query, but there was insufficient information to identify and contact the individual. In direct response to this, a modification was incorporated so that all information pertaining to an individual, collected at the time of registration, was readily available through the MEMOIR user interface. People were favourable to the concept, and noted that it would

be useful to telephone contacts suggested, find out what they were trying to do, and establish whether they could be of mutual assistance.

Some users reported surprise that MEMOIR did not suggest relevant people for topics with which the user knew other users in his group were associated. This situation arises because in the early stages of an evaluation, with a restricted group of subjects, very little information has been published and consequently the MEMOIR trailbases and linkbases are sparsely populated. A decision had previously been taken not to 'bootstrap' the system based on system log files, although the appropriate technology for doing this was available. Even if users were made anonymous, this process would constitute inappropriate use of log files since users were not aware that the fact *someone* had accessed documents in a particular domain would be made public.

4.4.4 System stability

Problems relating to the robustness of the system from the user's viewpoint were encountered in the evaluation, in some cases attributable to problems relating to the Java Virtual Machine implementation in desktop Web browsers.

Users noted that sometimes the applet would report that it had stopped. In some cases it would restart, but it was not clear what had caused this occurrence. The 'show trails' functionality in particular was often attributed to causing problems with a MEMOIR session, often terminal. Users of the system were particularly frustrated when they had built up a interesting trail which they intended to publish, and then encountered a crash of the session. This prompted some users to suggest and 'auto recovery' facility, similar to those available in many desktop packages.

5 Technical report (for project documentation/follow-up)

<<<text removed in public version>>>

6 Conclusions

The MEMOIR project has been successful in many respects. It has taken some visionary ideas for a new class of IT tools supporting knowledge management processes within large globally distributed companies, and demonstrated that they can be implemented successfully. It has shown that these tools can be implemented in a way which integrates with, and is complementary to, the complex IT infrastructures of such companies.

The participants in MEMOIR have learned a great deal about the issues associated with developing and implementing IT systems to support knowledge management. Participation in the MEMOIR project has had a profound influence on the market re-positioning of the two commercialisation partners, IBEX and Multicosm, and it has been a key element in enabling Multicosm to secure a significant injection of venture funding.

The value and usability of the specific implementations made in MEMOIR have not yet, as had been hoped, been fully proven through evaluation by large communities of end-users. However, the project partners have successfully disseminated and demonstrated the results of MEMOIR to a large, worldwide audience of researchers, industry consultants, opinion formers, and prospective corporate users of MEMOIR-like systems.

MEMOIR project partners are committed to continuing evaluation and refinement of the approaches developed in MEMOIR after the formal end of the project. This represents a strong endorsement of the value of what has been achieved so far.

References

Autonomy. Technical report, Autonomy Inc., 301 University Avenue, Suite 200 Palo Alto, CA 94301., 1997. Available as http://www.agentware.com/.

R. Barrett, P. P. Maglio, and D. C. Kellem. How to personalize the Web. In Conference on Human Factors in Computing Systems (CHI97), Atlanta, Georgia, USA, Mar. 1997.

V. Bush. As we may think. The Atlantic Monthly, 176(1):101-108, July 1945.

L. Carr, D. DeRoure, W. Hall, and G. Hill. The Distributed Link Service: A tool for publishers, authors and readers. In Proceedings of the Fourth International World Wide Web Conference: The Web Revolution, pages 647-656, Boston, Massachusetts, USA, Dec. 1995.

H. Davis, W. Hall, I. Heath, G. Hill, and R. Wilkins. Towards an integrated information environment with open hypermedia systems. In European Conference on Hypertext Technology (ECHT92), pages 181-190, Milano, Italy, 1992. ACM Press.

C. Dodge, B. Marx, and H. Pfeiffenberger. Web cataloguing through cache exploitation and steps toward consistency maintenance. In Third International World Wide Web Conference, Darmstadt, Germany, Apr. 1995.

R. Furuta, F. M. Shipman, C. C. Marshall, D. Brenner, and H. Hao-wei Hsieh. Hypertext paths and the World-Wide Web: Experiences with Walden's paths.In Proceedings of Hypertext '97, pages 167-176, Southampton, UK, 1997.

G. Hill, G. Hutchings, R. James, S. Loades, J. Halé, and M. Hatzopulous. Exploiting serendipity amongst users to provide support for hypertext navigation. In Proceedings of Hypertext '97, Southampton, UK, pages 212-213, 1997.

G. Hill, R. Wilkins, and W. Hall. Open and reconfigurable hypermedia systems: A filter based model. Hypermedia, 5(2):103-18, 1993.

IBEX Corporation SA, F-74160 Archamps, France. ITASCA Distributed Object Database Management System Release 2.3.5, Technical Summary, 1995.

T. Joachims, D. Freitag, and T. Mitchell.WebWatcher: A tour guide for the World Wide Web. In Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97), Japan, Aug. 1997.

B. Kahle and B. Gilliat. Alexa - navigate the Web smarter, faster, easier. Technical report, Alexa Internet, Presidio of San Francisco, CA., 1998. Available as http://www.alexa.com/.

H. Lieberman. Letizia: An agent that assists web browsing. In International Joint Conference on Artificial Intelligence (IJCAI-95), Montreal, Canada, Aug. 1995.

D. Maltz and K. Ehrlich. Pointing the way: Active collaborative filtering. In Proceedings of the Conference on Human Factors in Computing Systems (CHI95), Denver, Colorado, USA, May 1995. ACM.

D. Nicol, C. Smeaton, and A. F. Slater. Footsteps: Trail-blazing the Web. In Third International World Wide Web Conference, Darmstadt, Germany, Apr. 1995.

A. Pikrakis, T. Bitsikas, S. Sfakianakis, M. Hatzopoulos, D. DeRoure, W. Hall, S. Reich, G. Hill, and M. Stairmand. Memoir - software agents for finding similar users by trails. In PAAM98. The Third International Conference and Exhibition on The Practical Application of Intelligent Agents and Multi-Agents. March 23-25, London, UK, pages 453-466, Mar. 1998.

P. Resnick and H. R. Varian. Recommender systems. Commun. ACM, 40(3):56-58, Mar. 1997.

M. Röscheisen, C. Mogensen, and T. Winograd. Beyond browsing: Shared comments, SOAPs, trails and on-line communities. In Proceedings of the Third International World Wide Web Conference, pages 739-749, Darmstadt, Germany, Apr. 1995.

A. Schickler, M. Mazer, and C. Brooks. Pan-browser support for annotations and other meta-information on the World Wide Web. In Fifth International World Wide Web Conference, Paris, France, May 1996.

Tympani products. Technical report, Tympani Development Inc., 1307 South Mary Ave., Suite 102 Sunnyvale, CA 94087., 1997. Available as http://www.tympani.com/.

U. K. Wiil, editor. Proceedings of the 3rd Workshop on Open Hypermedia Systems, Hypertext '97, Southampton, England, April 6-11, 1997.

P. T. Zellweger. Scripted documents: A hypermedia path mechanism. In Proceedings of Hypertext'89, Navigation in Context, pages 1-14, 1989.