



Technical Report
RAL-TR-96-016

Heterogeneous Database Access and Multimedia Information Presentation: The Final Report of the MIPS Project

**A Bruffaerts J Donald J Grimson D Gritsis K Hansen A Martinez
H Williams and M Wilson**

March 1996

© Council for the Central Laboratory of the Research Councils 1995

Enquiries about copyright, reproduction and requests for additional copies of this report should be addressed to:

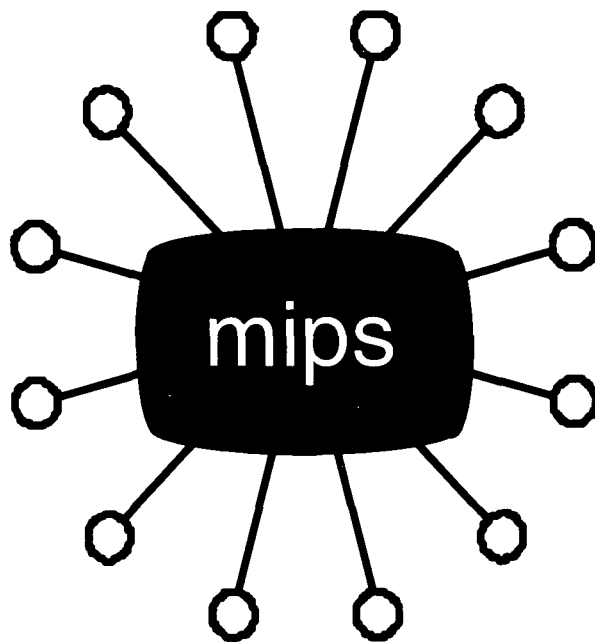
The Central Laboratory of the Research Councils
Library and Information Services
Rutherford Appleton Laboratory
Chilton
Didcot
Oxfordshire
OX11 0QX
Tel: 01235 445384 Fax: 01235 446403
E-mail library@rl.ac.uk

ISSN 1358-6254

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

Esprit III Project 6542

Heterogeneous Database Access and Multimedia Information Presentation : The Final Report of the MIPS Project.



A Bruffaerts,	Sema Group Benelux SA
J. Donald,	Cartermill International Ltd
J. Grimson,	Trinity College Dublin
D. Gritsis,	Epsilon Software SA
K. Hansen,	Danish Technological Institute
A. Martinez,	Sistemas y Tratamiento de Informacion SA
H. Williams,	Heriot-Watt University
M. Wilson,	Rutherford Appleton Laboratory

13th February 1996
Project Document Ref. : MIPS/LC/125 V2.0

Acknowledgements

The MIPS consortium is grateful to the MIPS Project Officer and Reviewers for all their helpful and constructive comments which have greatly contributed to the success of the project.

We are especially indebted to the Corfu Initiative for all their support, help, and many kindnesses to the MIPS team during the project review held on their beautiful island.

There is a special thanks to all those project team members who spent long hours foregoing the pleasures of foreign places for the sake of the project.

Finally, we are highly appreciative of the foresight of the European Commission and the partners commercial and business managers in sponsoring the project.

TABLE OF CONTENTS

Acknowledgements	3
1. Introduction	6
1.1 Background	6
1.2 Motivation	7
1.3 Objectives	9
2. Approach and Method of Working	12
3. Requirements	14
4. MIPS Functionality	19
4.1 Overall functionality	19
4.2 User	22
4.3 Application Scenarios	23
4.4 Conclusions	25
5. MIPS Architecture	26
6. MIPS Modules	30
6.1 Presentation & User Interface	30
6.1.1 The General Query Tool	31
6.1.2 The Knowledge Base System	32
6.1.3 The MIPS Presentation Tool	33
6.1.4 The Presentation Tools Interface	33
6.1.5 The Presentation Manager	33
6.1.6 The HyTime Engine	33
6.2 Selection and Delivery	34
6.2.1 Data Model and Queries Description	34
6.2.2 Selection and Retrieval Tool (SRT)	35
6.2.3 Web Builder	36
7. Demonstrator	37
8. Contribution to the state-of-the-art	40
8.1 Access to heterogeneous data sources	40
8.2 Hypermedia application description and multimedia data	43
8.3 User Interfaces and HCI issues	44

9.	Results	46
9.1	MIPS Toolkit	46
9.2	MIPS Modules	49
9.3	MIPS Demonstrator	54
10.	Conclusions	56

Appendix 1 - Information Available

Appendix 2 - References

Appendix 3 - MIPS Conference Paper References

Appendix 4 - MIPS Consortium

1. Introduction

1.1 Background

The Multimedia Information Presentation System (MIPS) project was conceived in the early '90s against a background of major advances in communications offering the prospects of access to large amounts of information and the advent of multimedia capabilities in computing offering the prospects of presenting information in a form which could be readily assimilable and usable by a range of users, including those with little knowledge of computing.

The spread of networks in the time leading to the project had resulted in a general acceptance of the concept of large-scale interconnection of computer systems and the benefits that this may bring with it. The development of large distributed database systems was expected to provide a similar revolution for information storage to that which networks did for communication. The potential to link together a large number of different though related databases to provide a massive distributed information store already existed and the spread of such systems was seen as simply a matter of time.

The explosion of information electronically available through the interconnection of a large number of different information sources was expected to bring with it problems, both in terms of finding the information which a user requires and of presenting it in a suitable format. Users cannot generally be expected to know exactly where information is stored and in what format it is held. Presentation of returned information is partly a question of taking account of user preferences in the way in which data are presented, and partly a question of reducing a large quantity of data to a form which is meaningful to the user.

The projected growth in the multimedia market was seen to have major impacts for all aspects of home and business systems because the user could both hear information and see it in a graphical representation. Current research had shown that we recall 20% of what we see, 20% of what we hear and 50% of what we see and hear. Multimedia is particularly important to those members of the population who prefer graphic images to text and have not, until fairly recently, had a technically feasible option.

Also in the time leading up to the project the benefits of creating and maintaining documents in the Standard Generalized Markup Language (SGML) (ISO 8879) were being realised and a development of SGML in terms of the HyTime standard (ISO 10744) to deal with multimedia information was under discussion. Creating HyTime compliant hyperdocuments for future presentation was seen as offering considerable advantages to publishers and information providers.

The project was, therefore, clearly focused on developing a toolkit to meet the requirements of users to gain more effective access to multimedia information sources. The increasing availability of information and the lack of technologies in the market place that allowed the user control over the way in which the information is presented, were exacerbating the

problems of information overload. These problems were manifest as sub-optimal productivity, lack of competitiveness, reduced quality of life and inhibited interchange of information.

Considerable activity at the time was being devoted to the development of multimedia tourist information systems. However, these were being developed on different platforms with different user interfaces making it difficult to offer access to these services to travel agents as the agents would require different hardware and/or software platforms.

Although most travel agents use a centralised airline reservation system to book flights to these destinations, there is no way for agents to reference any of these tourist information systems to book accommodation or other tourist attractions. There is an obvious requirement for a service offering access to travel agents to this information. A distributed database capable of storing and displaying multimedia elements could provide such a service to travel agents.

The tourism area was therefore considered to offer good potential for a demonstrator of the technology to be developed.

1.2 Motivation

The consortium brought together to undertake the project was a vertical integration of users and technology developers complemented by specific research expertise in Universities.

All of the partners had identified how the results of the project could generally contribute to the development of their organisations and plans for full exploitation and commercialisation were to be finalised in the light of technology and market developments during the course of the project.

In the first year of the project the Spanish end user/integrator withdrew and their contribution was undertaken by a Greek system integrator with strong links to the travel industry.

1.2.1 Cartermill International Ltd

Cartermill International Ltd (formerly Longman Cartermill) have been successfully involved in the creation of electronic information products and services since 1985. The company addresses the information needs of a wide variety of customers in Europe and USA. Cartermill was keenly interested in the publishing opportunities offered by the project to address the new markets for high-quality interactive publications that multimedia technology offers. It was anticipated that MIPS technology would offer competitive advantage in developing new media publications through incorporating multimedia information from external databases to complement the current database information held by the company and in delivering information from multiple sources in an effective way for non-expert end-users of medical, financial, current affairs and business information.

1.2.2 Rutherford Appleton Laboratories

RAL is a large research institute within the Council of the Central Laboratory for Research Councils (CCLRC) and was formerly part of the Science and Engineering Research Council (SERC). One of their major objectives is the support of UK HEIs (Higher Education Institutions) in research and development. One outcome from RAL work is dissemination of results and know-how to the UK academic community, and direct technology transfer to UK industry. RAL saw a need for the MIPS technology in further projects for industry customers to support a range of applications with multiple, multimedia information sources and their delivery to the end-user. Areas of application included catalogues - presented as multimedia information, the provision of workstations that can handle up-to-date, on-line multiple information sources for training and education in HEIs and consultancy services based on this technology for agencies such as ESA and CERN, and more locally, for UK industry, and government administration.

1.2.3 Sistemas y Tratamiento de Informacion SA

STI has been very active in the development of customised databases and had identified the need to develop and maintain, over a set of different vendor machines, a generic and ambitious database product. Their involvement in a number of Framework projects specially RACE Multimed and ESPRIT EDS, had allowed the development of skills and expertise on databases and user interfaces.

The challenge seen by STI was to use advanced techniques to retrieve information from alien heterogeneous databases and design an appropriate, expressively rich and operationally powerful, language to allow the interworking of different data structures and the extraction of the information sought. STI envisaged that the methods, technologies, and tools developed by the project would be applicable to the design of commercial products that would greatly enhance the usefulness of many Value Added Services, in particular those related to on-line databases, by allowing the extraction of information of heterogeneous databases to take place transparently - an impossible task at the start of the project.

1.2.4 Sema Group

Since 1987 Sema Group has been actively marketing SGML-based tools round the world, including the setting up of a network of agents to spread European expertise into America and Asia. The aims of MIPS fitted directly into the medium/long term development plans of the Sema Group in the areas of sales of the proposed HyTime engine as an extension to the existing set of SGML-based tools available from Sema Group, integrating the complete toolset into existing CASE tools to provide multimedia facilities for software developers, and utilising the toolset to provide multimedia information sets and training packages for customers. A wide range of possible products was foreseen for the proposed toolset.

1.2.5 Epsilon Software SA

Epsilon Software has a very strong relationship with a number of tourist-oriented enterprises, such as tour operators, hotels etc. Epsilon has provided the software for the operation of one of the biggest and certainly the most prestigious hotel in Athens, the "Grand Bretagne". It saw a need in the home market for multimedia presentation stations related to the tourism, such as for local travel agents in Crete and in Corfu and for providing visual and audible information on tours, along with the latest prices and new package deals, which the customer could use in order to choose a tour which fits his entertainment and financial needs.

The demonstrator system was envisaged to be adaptable to kiosk operation and as a stand alone or networked workstation which would open up a substantial market for the MIPS end product and would further enhance the company image as a high-technology solution provider.

1.2.6 Danish Technological Institute

The Danish Technological Institute (DTI), with its 1200 employees of which 700 are consultants, is Denmark's largest supplier of technological services including industrial management. The Institute's activities include Industrial R&D, Consultation, Testing, Training and Information services.

DTI foresaw MIPS as complementing their expertise and business activities as consultants and in R&D projects through sales of applications and application shells developed by DTI with the MIPS toolkit, sales of parts of the demonstrator to charter tourism enterprises and travel agencies and sales of the toolkit and parts thereof to suppliers of value-added services and software developers of industrial information solutions.

DTI also intended to exploit the tools for internal use and for use by existing clients by creating intelligent presentation system solutions for its own departments, using the developed toolkit in connection with the work on market research and development for clients, and through development of training and educational materials which make use of its position as one of the largest suppliers in the Danish market of courses, seminars and conferences.

1.3 Objectives

This project sought to address the problem of information overload by filtering and prioritising multimedia information before it is presented to the user in an appropriate form. The aim was to create a tool set based around the ISO 8879 Standard Generalized Markup Language (SGML) standard that could be used to create products that allow the user to obtain and present distributed heterogeneous multimedia information from a variety of sources in a controlled, timely and cost-effective manner. The tool set proposed consisted of two main modules:

- The presentation and user interface module which provides the user with intelligent support for the presentation of retrieved information in the most appropriate manner. The application of mediators and agents allows the data to be filtered and prioritised before delivery.
- The selection and delivery module which provides access to the various sources of information and mechanisms for navigating between and within these sources.

Both modules utilise SGML data and link models in conjunction with a hypertext presentation mechanism and knowledge bases to perform the required interactions. Interaction with the information sources use international standards such as SQL and SFQL to implement transport mechanisms.

The main objectives of the project were:

- To develop tools which can be assembled to produce a system to deliver information from heterogeneous, distributed multimedia sources to the user in a uniform format. This system was to be built on existing multimedia tools and standards.
- To develop a capability for the introduction of personalised user profiles and to apply the concepts of mediators and agents in order to filter and prioritise the information presented in a personalised information space.
- To ensure that the interface developed for the personalised information space meets the needs of non-professional users in terms of consistency and usability. This is of major importance to the exploitability of the MIPS system.
- To define the market requirements, industry benefits and exploitation potential for its developments and to show how the developments lead to improved IT-based solutions for use by industry.
- To design a demonstrator addressing the information handling problems of the tourism industry sector which is based on the user requirements obtained from:
 - the participating user industry representatives
 - the target end users of these industries
- To realise this demonstrator using the tool set developed and to evaluate it. This multimedia application was to clearly show the benefit to the particular end-user group and the exploitation potential of the development.
- To prepare for the exploitation of the results of the project.

As a subsidiary objective the project was to evaluate the advantages offered by the new ISO HyTime standard (ISO 10744) for linking together heterogeneous data sources by applying this emerging standard to the development of the demonstrator.

The chosen area for development of the demonstrator was the field of tourism. This area had clear potential for multimedia developments and it was envisaged that the presence of other potential user group representatives in the consortium would facilitate the development of a complete presentation system which could be applied to other multimedia application areas.

2. Approach and Method of Working

In order to attain a leading edge in a rapidly expanding and evolving marketplace the consortium sought to complete the project within a 36 month period with the main development effort occurring between months 6 and 30 of the project and evidence of the successful development of the systems proposed in the project provided by means of a major demonstrator in the tourism area. Work on defining the format of this demonstrator - including collection of appropriate materials - was scheduled to start at month 6 and to be completed by month 27. Integration and evaluation of the demonstrator was due to start at month 18 with completion due at the end of the project - month 36.

Following a week long workshop in month 2 of the project during which the perceived functionality of the system was assigned to specific modules within an overall system architecture, responsibility was assigned to partners on a module by module level in accordance with the technical annex. Work on the definition of the architecture of the system and the functional specification for each module was then undertaken as part of the production of the project's first technical deliverables [1] and [2]. During this time several meetings were held between partners in order to define interfaces between modules and ensure that any misconceptions or misunderstandings could be identified and resolved as early as possible in the project.

The original intention in the workplan had been that the requirements workpackage would be progressing in parallel with the technical discussions taking place in order to ensure that the requirements identified would have a major impact on the system design. In fact, due to the withdrawal of the Spanish user, delays in starting the requirement capture activities meant that the architecture of the system was well developed before any input was available from the requirements workpackage and therefore the user interface was restricted to functionality within specific MIPS modules rather than to the system architecture itself.

Epsilon's integration into the project meant that some of the lost time was made up in terms of the user representation on the definition of the demonstrator itself. However work on the technical development of the MIPS toolkit was still progressing as planned and therefore by the mid-project review it was felt that it was possible for the project to get back onto schedule. However the development work which was intended to be completed by month 27 was in fact more adversely affected by the lack of user involvement than expected and resulted in the consortium requesting extensions to the project duration.

Following the initial workshop the overall approach to progress meetings was that the Technical Team would meet every two months under the direction of the Technical Project Manager who was then responsible for reporting the outcome of the Technical Team's meeting to the Project Steering Committee which would meet on the day immediately after the Technical Team meeting. This Project Steering Committee operated as the main decision-making body within the consortium and although a Project Management Board (comprising of representatives of full partners only) was established as part of the project management

structure it met only three times during the project life time.

Individual workpackages met as required and any issues or decisions having an impact outside an individual workpackage were discussed and reported at either the Technical Team meetings (for those workpackages directly related to technical developments with the project) or the Project Steering Committee (for exploitation issues).

3. Requirements

As a pre-competitive research and development project, MIPS is positioned on the technology/product life-cycle between the solution of basic research issues, and the development of product prototypes to assess the market positioning of products (Figure 1). The main issues to be addressed were those of technology integration and systems engineering, including the early stages of product definition. Therefore, the requirements on the system are a combination of both market pull and technology push.

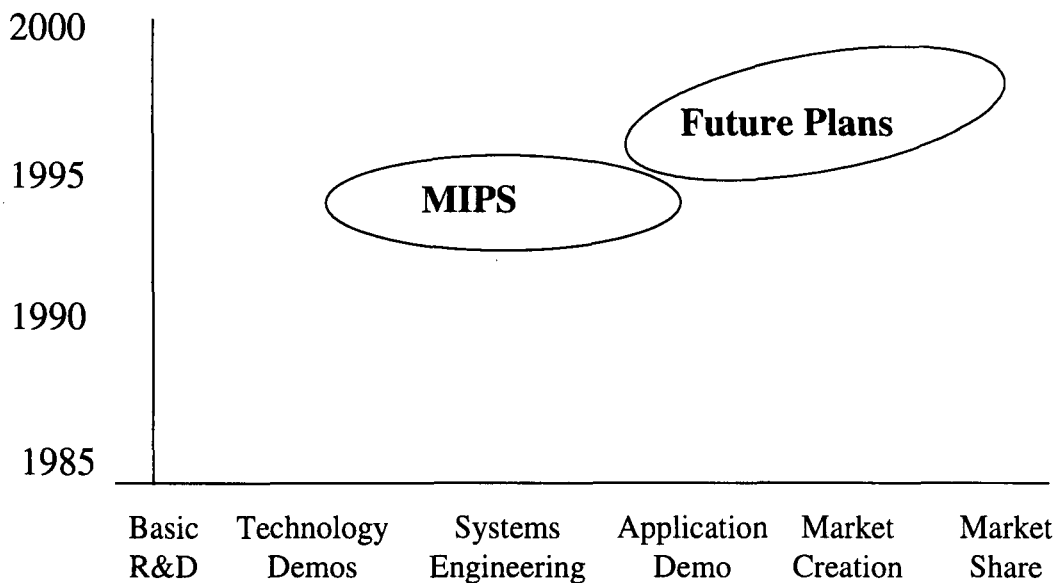


Figure 1 - MIPS in the RTD to Market Share Spectrum

The initial vision of MIPS usage before user requirements were detailed was:

A MIPS system will be used to present information to end users. The information will exist in the form of an application document. This will be read into the MIPS system which will allow the user to browse it.

An example would be tour operators producing their brochures as MIPS applications which would be read by potential customers on a MIPS system running in a travel agency. Such brochures could contain all the required information in stand alone form;

they could contain pre-written queries which call up some information from local or remote databases when it is required, or they could allow the user to create “free” queries to these local or remote databases to obtain information. Three scenarios further describing these alternative uses of a MIPS system have been presented in the MIPS architecture deliverable [1].

Behind this vision there were three main areas of technology push to include functionalities:

- 1) To construct a HyTime engine which could be used to present hypermedia applications.
- 2) To allow access to multiple heterogeneous distributed databases from single queries in order to retrieve information for display.
- 3) To apply knowledge engineering techniques to dynamically tailor the overall system to the user and task.

Each of these areas was one where a partner had already been involved in basic research, and the desire to further existing results through the technology life-cycle was independent of market pull, although clearly within the initial vision of MIPS usage.

Having acknowledged that the requirements for MIPS would be derived from both technology push and market pull, a set of user requirements was developed starting with those provided by the ISO generic user requirements for multimedia systems (reported in [4] Appendix 6). The domain of intended demonstration was derived from the initial vision, and not from a clear understanding of market needs, or the multimedia market economy. The area of tourism clearly employed information in both static and continuous media, and had been an early application area for shallow multimedia systems in kiosks. However, the requirements for MIPS needed to be more generic than those of the single tourism domain, and requirements needed to be collected for the tool envisaged in the vision of MIPS usage, incorporating the three elements of technology push, more widely than purely multimedia authoring or presentation.

The first stage in developing requirements was to extend the initial vision of MIPS, and 17 scenarios were developed of how a MIPS system would be used in different vertical market sectors (these are reported in [4] Appendices 1 & 2).

To develop these scenarios (reported in [4] Appendix 3) the MIPS team either used domains about which they had personal experience and knowledge, or involved working with others who had expertise in those domains. These produced many requirements, but their importance was hard to determine since business benefits and costs were not clearly calculable from these general scenarios.

Following the scenario development four areas were selected for detailed case studies

including the specification of Enterprise, Task and User Models: the Manufacturing Industry, the Publishing Industry, Art Festival Management, and a Multimedia Information Kiosk. Overall Enterprise Models included a Strengths, Weakness, Opportunities and Threats (SWOT) analysis of the enterprise; showed the role of the task supported by the tool in the enterprise, and both the financial importance of the task, and the notion of quality which could be applied to it, resulting at the lowest level in descriptions of the Critical Success Factors (CSF) for the enterprise.

To complement this analysis, the user models suggested the domain and system experience of potential users, allowing an estimation of the importance of the task to them, the likelihood of training being made available to them, and the likelihood of them develop experience with the system. These factors influence user interface design, documentation, training and the supporting methodology around the use of the tool described in section 5. These studies are reported in Appendix 4 of [4].

Having successively detailed generic requirements from 17 scenarios to four detailed enterprise studies, the final user requirements capture was to study the demonstration domain of tourism in greater detail. Interviews and observational studies were undertaken at two tour operators and five travel agents in Spain. The tour operators develop holidays and publicise them, the travel agents have contact with the public and sell the holidays as part of their business. The structure of these seven analysis follows that used for the four case studies.

From the demonstrator domain study, the four case studies and the 17 scenarios, 14 classes of detailed user requirements were derived:

1. Display of hypermedia information
2. User and task dependent presentation
3. Information previewing & sampling
4. Query refinement
5. Transparent access to remote and local heterogeneous information sources
6. Integrated access to existing applications
7. Distributed application (i.e. support for cooperative working)
8. Local report generating facility
9. Authoring of hypermedia information
10. User and task dependent tailoring
11. Ability to add new information sources
12. Acceptable Quality/Price/Performance/Usability
13. Security and Integrity
14. Installation requirements

These requirements were divided into three strengths. The strongest requirements are those which are essential requirements for the user and business needs (e.g. 1.1, 1.5 below). The second strength of requirement are those which it was considered that the system may meet. Mostly these were placed in this category since it was not clear how technically complex

meeting them would be. In addition, requirements were categorised in terms of the type of user who was more likely to benefit if the requirement was to be met. These were:

- N - Naive end user
- S - Skilled end user
- A - Application Builder
- C - Application Configurer

For example, requirement 1.14 states:

Should include control of video and audio sequences. Among other facilities to freeze, rewind go forward etc. (N,S)

The weakest requirements are those which are deemed to be outside the scope of the MIPS system. It is clearly understood that these will not be met by the MIPS system although they could be met at some point in the future. Mainly these would require considerable implementation or support effort for little perceived gain. For example, requirement 1.28 states:

Users should be able to see all monetary values in a stated value, using today's rate. At the same time the monetary units which payment is required to be made in must be clearly indicated. (N,S)

This would require an on-line currency rate server which was constantly maintained at the correct exchange rate. Although it may be possible to establish such a resource (maybe on the WWW by a currency business) and contact it when required, it is not essential to the system, and the costs involved would clearly be prohibitive.

Of the essential requirements, many are obvious although necessary:

Should include provision for the support of multimedia object notations including text, recorded sound, moving video, still raster image and record-oriented data. (N,S,A)

Other requirements resolve potential design trade-offs and require interpretation. This view is summarised in requirement 1.5:

Should include provision of a user interface that presents a single system image to the user. (N,S)

[3] includes all the user requirements for the system which were then translated into a set of technical requirements for the functional and non-functional aspects of MIPS, as well as to produce constraints to be applied on the later design stages as to which standards to conform to, and which hardware and software platforms should be supported. These technical requirements are reported in [5].

The last piece of work done under the general heading of requirements in the MIPS project was knowledge acquisition for the KBS inside the MIPS architecture. The two major areas of KBS usage are in remote database selection and information presentation design. Both of these are supported by a generic ontology of terms which can be used in queries for data to the data sources that MIPS has access to. In order to allow access to many different data sources from a single query, it is necessary to use terms in the query which are not derived from each particular data source, but which are defined centrally. Otherwise the different usages of the same term in different target data sources would result in conflicting interpretations of a query. In MIPS a general ontology of terms has been used which can be elaborated for individual domains such as tourism. The general ontology was collected from previous knowledge based systems, and from ontology studies. However the tourism terms used in the demonstrator system were gathered from potential users of the MIPS system through interviews. The knowledge acquisition techniques, and the knowledge representation used are both described in [18] along with a copy of the ontology itself printed as an ISA hierarchy, and as a set of object, attribute, value range tables.

4. MIPS Functionality

4.1 Overall functionality

It is currently possible with commercial products on a PC to retrieve information from SQL servers, databases, and documents across a network whether that information be text, relational tables, still images, sound or video. That information can be presented through tools in a commercial windowing system to users for them to read, or cut and paste into multimedia documents. However, the range of different data sources which can be used is limited; queries must be specified separately for each of the data sources and not as a single query to retrieve the information from any or all of them; and the tools to present the information will each occupy a different window on the screen and use proprietary presentation styles which differ from each other depending on the source format of the information. Another currently available form of presenting multimedia information is by authoring it in a proprietary tool into a discrete document or hypermedia network which the user can then browse. However, the user has no access to documents outside the hypermedia network and is tied to a proprietary representation format.

MIPS has sought to combine the best of these two approaches and thereby to overcome the problems of each and can viewed as providing, in a single integrated system, the essential functionality of currently separate systems for multimedia authoring, database access, report generation and information browsers (Figure 2).

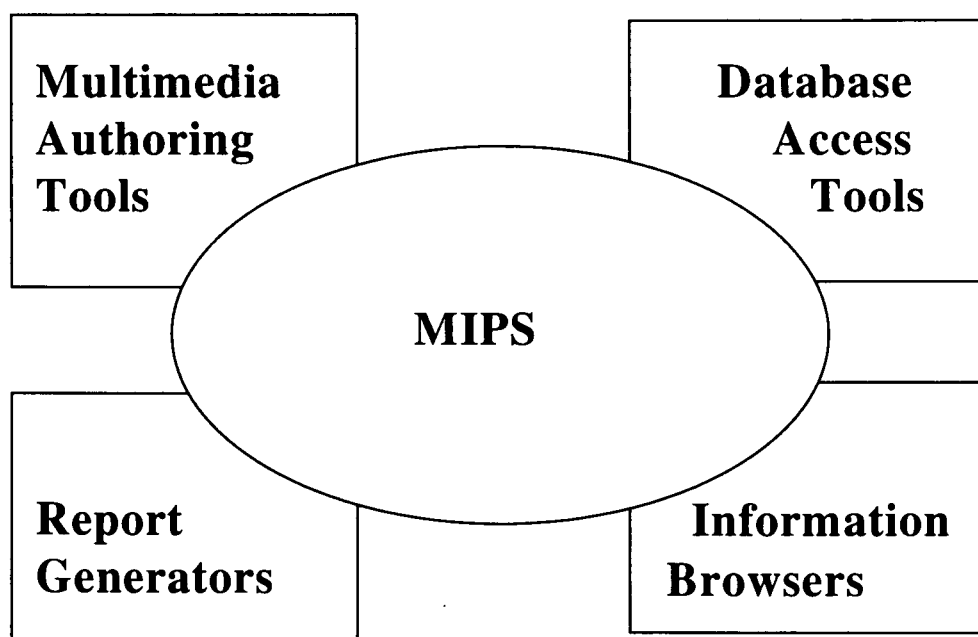


Figure 2 - MIPS in the Application Spectrum

MIPS allows an end-user to retrieve and browse heterogeneous multimedia information structured and presented in a user-friendly and consistent manner. In its most powerful configuration, it will allow the end-user to formulate queries which are interpreted, analysed, and despatched by the system to heterogeneous distributed external data sources, and to view a coherent and customized presentation of the data retrieved as answer. Pre-specified or dynamically retrieved assets are stored in, or referenced from, a set of hyperdocuments conforming to the ISO standards HyTime (ISO 10744) and SGML (ISO 8879). The hyperdocuments contain a navigation web of information nodes and hyperlinks. Each information node specifies presentable-data elements (which may be instantiated with data) and the means by which they may be presented as components of a graphical user interface (GUI) for the node. The data assets, web, and application-specific knowledge in a knowledge base system (KBS) module, together constitute the application description. This is authored by an application builder for a particular domain - e.g. tourism, aircraft maintenance, medical anatomy and surgery, etc.

The system uses the KBS:

- (a) to mediate the interpretation and analysis of queries to data sources,
- (b) to act in lieu of an application builder in dynamically specifying fragments of hyperdocument, especially for the incorporation and presentation of retrieved data, and
- (c) to provide graceful adaptation to the needs and preferences of end-users and the constraints of the system configuration in use.

The growing need for information which incorporates updatable data which is often held in widely distributed sources in heterogeneous formats and accessible remotely is met in MIPS by the separation of data and control allowing the reusability of both data and control mechanisms. MIPS separates 'control' structures - navigation structure, node structure, specification of presentation mechanisms, specification of presentation tools - and data. (A presentation mechanism is a set of window instances, plus any audio output, and a script of actions associated with windows or subwindows). In MIPS, data may be incorporated in the web, stored in a local filestore, or accessed in remote data sources. Locally or remotely stored data and SGML are used for the specification of the 'control' structures. This is a significant advance, even for the simplest scenario of MIPS usage - a fixed presentation of fixed assets which is equivalent to conventional pre-authored multimedia applications. In addition, the specification of presentation mechanisms is separate from the specification of presentation tools, easing portability to varying environments. Further still, node structures are separable from the specification of presentation mechanisms, thus facilitating variation in the modes of presentation of the data. Nodes are also separable from navigation structure, which facilitates reuse of the contents of information nodes for different tasks within an application or for different applications. This contrasts with systems in which links are embedded in node structure (e.g. the hypertext system Hyperties). The use of generic presentable-data elements

in node structures eases reuse of those structures for other applications. The web structure indicating how data and control structures are separately definable and incorporated into a presentation are shown in Figure 3.

The conceptual structure of the web is created by the application builder to meet user requirements for a particular application. In so doing, the application builder should be aware of the ontology of the application domain, the subset of object classes about which the user is likely to require information for any given activity within that domain, and the ways in which these object classes are related to one another. The types of nodes required, and the patterns of hyperlinks among them, can then be designed so as to facilitate browsing of the web by the user. The conceptual structure of the web will be related to the domain ontology constructed for the KBS module, and it may be possible to create different conceptual types of hyperlink to reflect different types of relation between objects in the ontology.

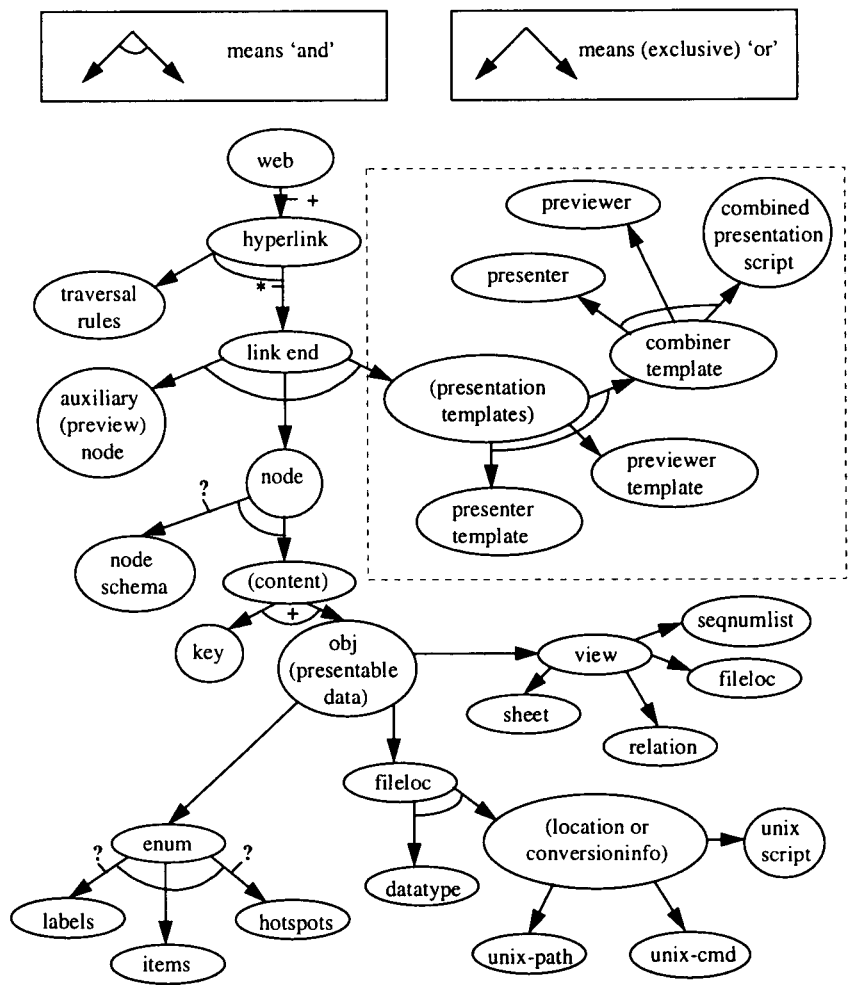


Figure 3 - The major components of the HyTime Web available from the MIPS DTD

A simple example of conceptual web structure, for the domain of tourism, is one which nodes having the application-specific types 'country', 'region', 'city', and 'hotel' are connected in a tree where the links have the conceptual meaning of 'part-of'. In the domain of aircraft maintenance, a large and complex web of nodes representing components might be organised so as to be browsable by structure of the aircraft (probably as a tree where the links have the meaning 'part-of'), or by 'walk-through' sequence of connectivity (where the links have the meaning 'gives-access-to'), or by functional system or maintenance cycle (possibly as a general graph where the links represent the functional roles of components or the sequence in which they should be inspected at a given maintenance period). The information contained in the nodes would thus be reusable for different activities within the domain, with each activity mapping to a certain conceptual type of hyperlink.

4.2 User classes

Arising from the user requirements, three different classes of user are envisaged at different stages in the development and use of an application.

- First, an *application builder* is required to author the initial HyTime hyperdocument containing the link and node structure, provide data assets to be presented or identify sources from which they can be obtained, and enter domain knowledge into the MIPS KBS. Data assets of text, static images, or video could be produced by application builders in their preferred tools, or used from existing sources provided by third parties. General KBS domain knowledge will already exist in MIPS for the application builder to expand in the application domain to enter details of data sources from which presentable information can be retrieved. If the application domain is one in which previous applications have already been produced (e.g. tourism) then this would only require very minor development.
- The second user of the system will be a *site configuration manager* who will install the MIPS application and further update the KBS with details of the presentation tools available at the site and the exact class of end-users who will use the system.
- The third class of user comprises the *end-users* who will be presented with the assets represented in the hyperdocuments. They will obviously vary in their experience of the domain and in using a MIPS application. To support this variety of *end-users*, the KBS has facilities for including detailed user models which are used to select appropriate presentation mechanisms and tools.

Each of these users has a role to play in the progressive tailoring of a MIPS application. The application will be tailored to a domain, site, and finally to the actual end-user at that site, in order to provide the most effective and efficient presentation of information. Details of application building are given in [24].

4.3 Application Scenarios

MIPS provides six levels of functionality to meet the six scenarios of usage which are outlined below. The functionality for each scenario includes that for preceding scenarios.

Scenario 1: Fixed presentation of fixed assets. All assets are already incorporated in the application description. The navigation structure is fixed. The presentation mechanism for each node is pre-specified, as are the presentation tools to be used. The application is immutable, standalone entity, and it might typically be stored on CD-ROM. This scenario is suitable for a simple hypermedia presentation.

In this scenario the application builder has to author the HyTime Application Description (HAD) in which the information content of each node is defined, the navigation structure of the presentation is determined along with the presentation tools to be used. In addition the KBS ontology needs to be defined and this involves specifying the objects applicable to the domain, the parent/child hierarchy, the object attributes and domain specific rules.

Scenario 2: Tailorable presentation of fixed assets: Again all assets are already incorporated in the application description and the navigation structure is fixed. The presentation mechanism for each node is template based allowing the assets to be presented with separately defined presentation tools. The assets are portable between presentation tools and the scenario is suitable for small hypermedia applications.

For applications in this scenario the KBS ontology needs to be defined as for scenario 1 and the HAD authored in the same way except that the HAD has to contain the description of the template-based presentation which allows the presentation to be used with alternative presentation tools.

Scenario 3: Tailorable presentation of explicit assets: The assets are held explicitly as raw datasets in local files. The navigation structure is template based. The presentation mechanism for each node is template based as in scenario 2. The assets are fully portable between applications and the scenario is suitable for large hypermedia applications.

This scenario needs to be set up as for scenario 2 with the addition of template based navigation structure which includes the schemata of the explicit datasets to be presented.

Scenario 4: Tailorable presentation of assets from databases: Assets can be dynamically retrieved from external data sources in answer to pre-specified queries. The application description contains queries with a fixed answer schema which have been pre-specified by the application builder, each of which is despatched if the end-user accesses the node containing it.

The application description is expanded to incorporate or refer to the data retrieved in answer to the query. The fragment of web which is to accommodate the retrieved assets may be

pre-specified by the application builder.

The KBS is expanded to include descriptions of the external data sources mapping the queries to known or potential sources of information with the specified attributes.

This scenario is suited to static hypermedia report generation.

Scenario 5: Adaptive presentation of assets from databases: In this scenario the system can cope with unexpected answer schemata. The KBS is further extended to allow the dynamic specification of the fragment of the web accommodating the retrieved assets and provides the functionality for adaptive hypermedia report generation applications.

Scenario 6: Assets can be dynamically retrieved from external data sources in answer to user-generated queries: The user can, at any juncture in navigation, formulate a query to external data sources. The query may be formulated by completing a query template or by construction of an open query using a constrained vocabulary. An appropriate web fragment and presentation mechanism for the answer are specified dynamically by the KBS.

In this scenario the application builder has to create the query templates and add the user modelling functionality to the KBS.

The user modelling facility is a trade-off between the expressiveness and coverage of a natural language interface and the run-time efficiency of a tunable defaults file approach. The goal of the KBS user modelling system is to have the required information available while the user is interacting with the system for the interface options to be set in the most cooperative way for users to achieve their task objectives. The development process includes: identifying the variations in the user population; identifying the variations in the interface/application functionality; identifying correspondences between classes of user (stereotypes) and the required interface functionality; identifying trigger or cue events which determine that a user has changed a user class (macro rules); representing this correlation as efficiently as possible in a user model; embedding the user model in a user modelling component which eavesdrops on the user's actions. At run time, the user's actions are monitored for cue events or combinations of cue events which trigger macro rules and change the user stereotype which in turn change the customisable user interface or application options.

As far as practical implementation of the MIPS design is concerned, the MIPS modules have been fully tested and demonstrated to meet the functionalities of Scenario 4 above, in the Tourism demonstrator. Unfortunately it was impossible, due to time and budget constraints, to achieve the objective of the full functionalities of the Scenario 6. The sound general design and the intermediate results obtained make us confident that the critical components of the proposed approach are perfectly suitable for the planned purposes, within the limits of course of a R&D prototype.

4.4 Conclusions

The design of MIPS provides a comprehensive range of functions for creating multimedia presentations not only for the development of hypermedia titles by authors and publishers, but opens up those documents in two significant ways: first, they are standardised to the HyTime ISO standard to ease portability and market growth; second, they may be expanded by users to accommodate data from online data sources, to provide up-to-date information tailored to the end-user. To do this, MIPS provides facilities (1) for application builder to incorporate queries to local and remote data sources in their application hyperdocuments, and (2) for end-users to issue new queries, the data retrieved in answer to which is used to expand the hyperdocuments. This expansion mechanism requires knowledge of screen, graphic, and other forms of design, currently only being developed in the hypermedia community. The query mechanism itself is driven by queries which are constrained to use the terms which are available in local or remote databases to index data.

5. MIPS Architecture

The generic MIPS system is made up of two subsystems: a Presentation Subsystem and a Retrieval Subsystem. The subsystems themselves are several UNIX processes and one MS-Windows task. The processes and tasks are communicating together via a cross-platform synchronous RPC mechanism, implemented by the Software Bus, which is a result of the Eureka ESF project. Figure 4 represents the whole MIPS system, its processes and task, and their interactions.

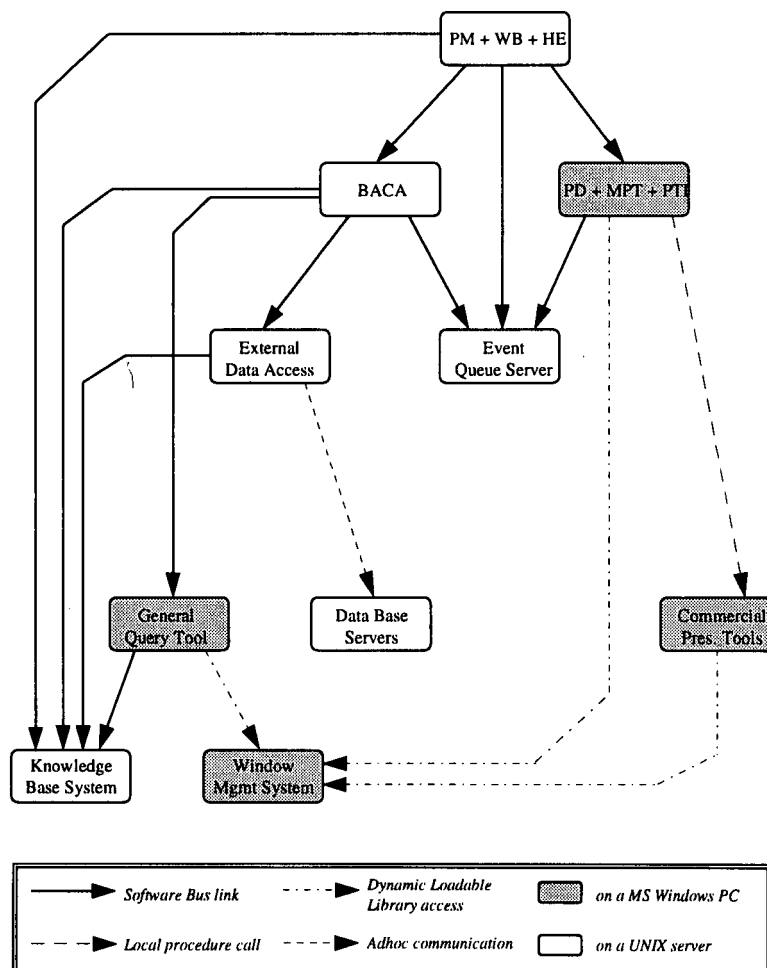


Figure 4 - The client/server control structure of MIPS processes

Integration technology

The Software Bus was selected as an integration technology because it provides a high-level environment for RPC with structured parameters and independence from the underlying interprocess communication primitives across several platforms.

The Software Bus also provides a so-called Service Abstract Description Language (SADL) which forces providers and clients of a service to share the same service definition. The software plugs, ensuring the communication among the processes and the marshalling and un-marshalling of the data structures passed as parameters, are generated automatically by the Software Bus SADL compiler in C or C++, for UNIX processes or MS-Windows tasks, for server or client components, on demand.

The adoption of the Software Bus allowed the easy physical re-distribution of the MIPS processes among several servers and the replication of subsystems of intercommunicating processes independent of their physical location.

Core architecture

The synchronous aspect of the RPC mechanism of the Software Bus motivates the diamond architecture of the generic system core.

At the top of the diamond, the Presentation Manager sends requests either to the Presentation Dispatcher, for executing presentation primitives, or to the BACA process, for evaluating queries. At the bottom of the diamond, the Event Queue Server stores the feedbacks it receives, either from the BACA process, when a query answer is available, or from the Presentation Dispatcher, when the end-user has expressed his own feedback by interacting with the graphical user interface.

Whenever the Presentation Manager is idle, it polls the event queue to determine what to do next: either integrate some data retrieved from external databases, or execute the script associated to the event generated by the user feedback.

MIPS Presentation Subsystem

The Presentation Subsystem is a self-contained playback system for MIPS Hypermedia Application Descriptions (HAD), represented as a set of HyTime-compliant SGML documents. The Presentation Subsystem only needs the Retrieval Subsystem if queries are embedded within the description; it is made up of a Presentation Manager and a Presentation Dispatcher, both using the service of an Event Queue Server to handle the asynchronicity of the user feedbacks.

The Presentation Manager (PM), which is a UNIX process, uses the API of 3 static libraries:

the SGML Engine (SE),
the HyTime Engine (HE)
and the Web Builder (WB),

which are implemented in C++ above the OODBMS ObjectStore.

The SGML Engine implements a set of classes to handle SGML documents as trees of typed nodes with lists of attributes. The HyTime Engine implements a set of classes to support a HyTime view of a set of HyTime-compliant SGML documents. The Web Builder implements the instantiation process underlying the concept of HAD templates introduced to shorten the HADs and to integrate raw data within navigation and presentation structures.

The role of the PM is to playback the HAD loaded into the system.

The Presentation Dispatcher (PD), which is a MS-Windows task, uses the API of 2 libraries:

the MIPS Presentation Tools (MPT)
and the Presentation Tools Interface (PTI).

They are both implemented in Borland C++.

The MPT implements the set of interaction mechanisms available in MIPS windows and their associated primitives. The PTI gives access to third-party applications controlling their own windows.

The role of the PD is to translate the abstract presentation primitives generated by the PM into the calls to the procedures implemented by the MPT and the PTI. The PM also translates the events generated by the end-user into abstract user feedbacks and stores them in the Event Queue Server.

The only other process in the Presentation subsystem is the General Query Tool (GQT). This process is launched by the BACA process (see hereafter) whenever a user-defined query is required. This should only happen in advanced uses of the MIPS system. The GQT uses the MPT to interact with the end user.

MIPS Retrieval Subsystem

The Retrieval Subsystem is a query evaluation system which is able to retrieve data from heterogeneous data sources. The Retrieval Subsystem is independent of the HyTime/SGML aspects of the Presentation Subsystem: it receives queries encoded in IRL(QD) and yields query answers encoded in IRL(RDD).

It is made up of 3 UNIX processes:

the BACA process,
the EDA process
and the KBS process.

The BACA and EDA processes are implemented in C/C++. The KBS process is written in Prolog.

The BACA process is responsible for the clarification and breakdown, with the help of the KBS process, of the IRL(QD) queries into possibly several subqueries targeted to specific data sources. According to the paradigm obeyed by the data sources (relational DBMS or document DBMS), the subqueries are encoded in different sublanguages and passed to the EDA process.

When the subquery answers are available, they are collected by the BACA process and assembled together in a consolidated answer corresponding to the original IRL(QD) query. The query answers are then stored in the Event Queue Server from where they will be picked up by the PM.

The EDA process is responsible for translating the subqueries into the query dialects specific to the targeted data sources, retrieving the data from the specific data sources and converting them into the format accepted by the BACA process.

The KBS process is responsible for providing the necessary information on the available data sources, their accessibility and cost and their schematic description. It also provides the ontology which defines the terms to be used in the IRL(QD) queries and their mapping onto the schematic description of the external data sources. It contains the rules to be used for breaking down queries in subqueries.

Hardware and software platform

The current version of the MIPS system runs on a platform made up of a Sun server under SunOS 4.1.3 and an Intel-based PC under MS-Windows 3.1 connected together by Ethernet.

The use of the Software Bus allows the distribution of the UNIX processes on several Sun servers without any recompilation. Porting to other UNIX systems will be possible if they support C and C++ compilers, a compatible Prolog implementation for the KBS, and an implementation of ObjectStore for the PM.

Porting to MS-Windows NT should be relatively easy since it provides real multi-processing and an implementation of sockets.

6. MIPS Modules

The overall architecture of MIPS showing the interconnectivity modules is shown in Figure 5.

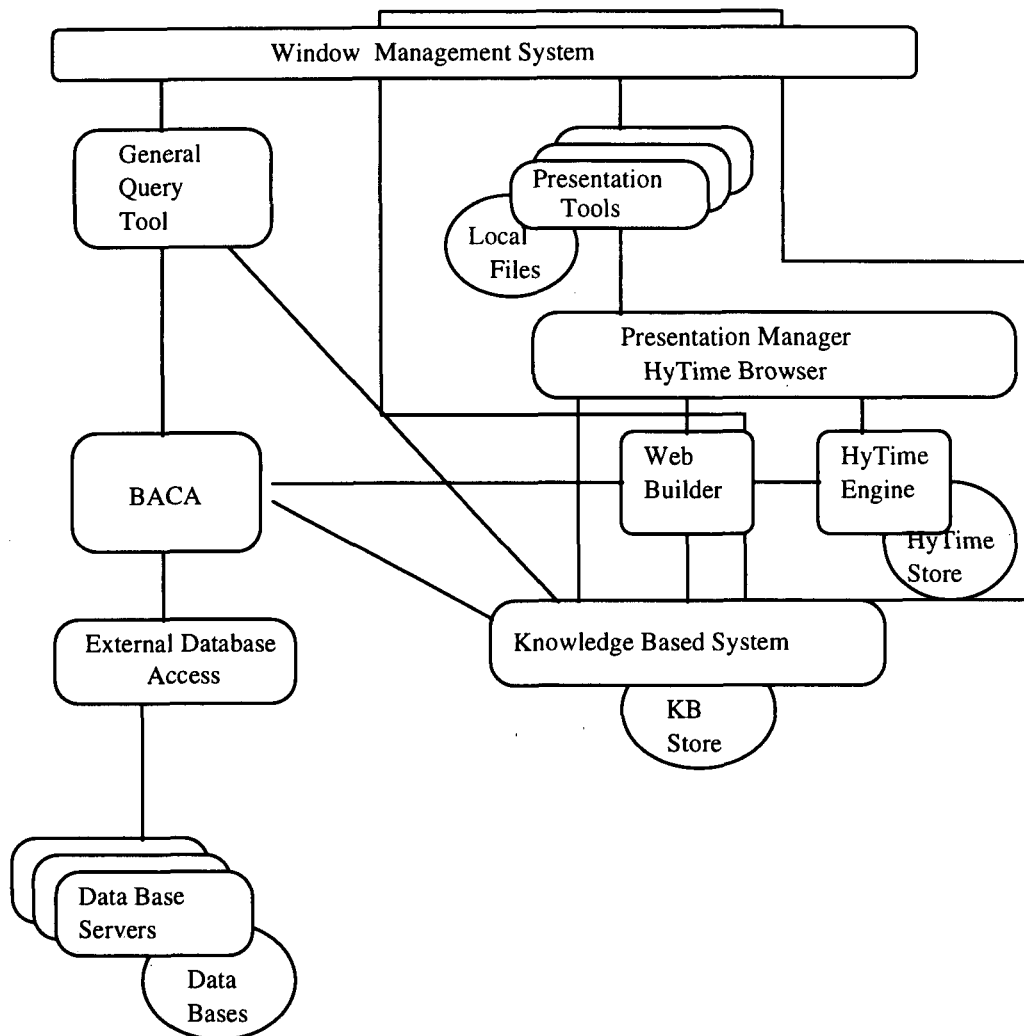


Figure 5 - The MIPS Modules

6.1 Presentation & User Interface

The modules concerned with presentation and user-interface issues are: the General Query Tool (GQT), the Knowledge Base System (KBS), the MIPS Presentation Tool (MPT), the Presentation Tools Interface (PTI), the Presentation Manager (PM), and the HyTime Engine (HE).

The functional specification for these modules were provided in [10], [11] and [12], with the

designs of the GQT, MPT and PTI described in [16], and the design of the PM, HW and KBS presented in [20]. The final technical and user documentation for all these modules for all these modules was then presented in [23] and [24].

6.1.1 The General Query Tool

The purpose of the GQT is to allow end-users of a MIPS system to create queries dynamically during an interactive session, which are then used to retrieve data from data sources accessible to the system. A graphical user interface allows this to be achieved in a usable and simple manner. The queries are output to the Selection and Retrieval Tool (SRT) in the Query Dialect of the Internal Representation Language, IRL(QD).

The GQT is loaded with query templates created by an application builder for a particular application of the system, e.g. tourism. Each query template represents a particular user task: e.g. flight booking, car hire, or accommodation reservation. When the GQT receives a request from the SRT to create a query, it activates its initial screen which displays an icon for each template. An end-user can select an icon, or several in turn, to specify a query. Once a template has been selected, the end-user will then specify constraints on the elements of the query, to tailor the data retrieved to the user's information requirements. For example, in the case of accommodation reservation, an end-user might:

- specify the location, e.g. Paris, or several locations, e.g. Paris or Rome;
- specify a range for price, e.g. < 200 FF;
- specify media types to be retrieved, e.g. a still graphic or video of the accommodation;
- specify whether a constraint is mandatory, e.g. show hotels costing more than 200 FF if no cheaper accommodation is available.

Only the most experienced users will have access to all these facilities, and a default value can be suggested for each entry. Default values can be fixed and the option to change them disabled for certain users.

A natural-language description of the query so far is provided to the end-user, updated after specification of each element of the query. When all the elements that constitute a 'minimum query' for a given template have been specified, the end-user can add them to the overall query so far. When this happens, the natural-language description is frozen and the template icon column is redisplayed. When the end-user has used all the templates they require for their query, they select the appropriate button to enable retrieval of data. This results in an intermediate description of the query, containing a combination of information from the various templates and the end-user's input, being passed to the Query Translator (QT) for translation into IRL(QD) and despatch to the SRT.

6.1.2 The Knowledge Base System

The KBS's role in relation to presentation and user-interface issues is to act as an 'automatic author' of application descriptions - specifying modes of presentation, GUI design details, and navigational structure on the basis of default heuristics and any expressed user preferences.

The MIPS approach to the construction of application descriptions, both by the application builder and dynamically by the KBS, involves the use of templates. Templates specify fragments of information web in a range of granularity from navigation structure to the basic components of information nodes. They may refer to other templates, so that high-level templates may be constructed from lower-level 'building blocks'. The rationale for the use of templates is twofold: to minimize the size of the application description by allowing reuse of structures specifying uniform presentation formats for an arbitrary number of instances of the same type, and to ease the task of the application builder (or the KBS) in creating the application description.

The KBS, acting as an 'automatic author', is called upon by the PM to provide 'presentation' web templates to specify the presentation format for the contents of information nodes, and by the Web Builder (WB) to provide 'navigation and node structure' web templates (plus node schemas if appropriate) specifying those aspects of web fragments to accommodate retrieved data. The KBS attempts to dynamically match or construct templates on the basis, not of data itself, but of schematic representations of the structure and content of the data as provided by the PM/WB. There are three major kinds of schema: answer schema, web template argument schema (WTAS), and node schema. They specify the structural interrelationship of raw or presentable data components and characterize those components in terms of types in the KBS ontology. The syntax of each kind of schema is similar to that of the others and mirrors that specified by the MIPS HyTime DTD. An answer schema is an optional part of an IRL(QD) query; it can specify how the various components of the query are grouped together for the purposes of navigation and presentation - on the basis of both semantic (conceptual) and HyTime syntactical criteria - and parameters of their presentation. Where a query contains no answer schema, the KBS is capable of constructing one on the basis of the formula of the query and various grouping and design heuristics. The WTAS and node schema are similar structures which describe web templates and information node structures in the application description.

The full power of MIPS is manifest when the end-user retrieves data from heterogeneous distributed external data sources and views a customized presentation of that data. To accommodate retrieved data, whether the query was (a) prespecified by the application builder and stored in the web or (b) generated through end-user interaction with the GQT, requires that the web be dynamically expanded during the end user's session. This is done by the WB module according to a specification supplied by the KBS in the form of a web template whose instantiation with the retrieved data will result in an appropriate fragment being incorporated in the web. This web template is selected by the KBS according to the best match of the query answer schema to a library of WTASs indexing web templates in the

application description. Failing to match to a WTAS/template 'off the peg', the KBS attempts to construct a template 'custom tailored' on the basis of the answer schema. A similar process is followed when the PM passes a node schema to the KBS and requests a presentation template in return.

6.1.3 The MIPS Presentation Tool

The MPT has been developed to facilitate seamless multimedia applications and, in particular, the demonstrator system. It comprises a set of interface elements and an API for manipulating them. It provides a GUI with a coherent look-and-feel that is consistent with that provided by the GQT.

The MPT can render:

- plain text and RTF produced by MS-Word;
- .BMP, .GIF, and .PCX graphics formats;
- .WAV sound files;
- .AVI and .MOV video files.

6.1.4 The Presentation Tools Interface

The PTI interfaces with the Presentation Dispatcher to undertake the control of proprietary presentation tools (PTs) for the rendering of information that cannot be rendered by the MPT. Each PT is a MIPS-external, commercially available Windows tool.

6.1.5 The Presentation Manager

The role of the PM is to interpret, i.e. to play back, a HyTime Application Description (HAD). In MIPS, the application is described, in part, by a set of interrelated SGML documents compliant to the HyTime standard and stored in the HyTime Store module. (Other information resides in the KBS and in the local asset base.) These documents conform to the MIPS Document Type Definition (DTD). The MIPS DTD allows the representation of all the important aspects of a hypermedia application in an abstract and platform-independent way: data and references to multimedia assets, navigational structure, presentation structure and layout, and dynamic behaviour. The PM is in charge of implementing the MIPS-specific semantics of the DTD.

6.1.6 The HyTime Engine

The purpose of the HyTime Engine is to manipulate SGML documents and HyTime hyperdocuments to:

- import linear representations from external files and load them in the HyTime Store,

-
- export linear representations to external files,
 - retrieve information inside documents,
 - navigate through them and among them, update them and create some of them.

These functions are defined to support the ISO standards ISO 8879:1986 for SGML and ISO/IEC 10744:1992 for HyTime.

The HyTime Engine is split into two layers: an SGML layer and a HyTime layer. It is built on top of an object-oriented repository, since the structures stored are mainly trees and graphs. The one used in MIPS is ObjectStore. The HyTime Engine is linked with the Web Builder and the Presentation Manager to form the HyTime Process, which runs on the Unix application server as a Software Bus component.

6.2 Selection and Delivery

The work done under this heading was concerned with the selection and retrieval of information from external sources and its delivery to the Presentation Subsystem. The objective, as stated in the Technical Annex, is

"Specify, design, and implement tools to accomplish the selection and retrieval of information coming from distributed heterogeneous multimedia databases."

The work proceeded from both theoretical and pragmatic viewpoints. From the theoretical view came the query language and data model specified in [17]. It was decided to use this deliverable to investigate a potential model theoretic approach to implementing the SRT subsystem. At the same time, a pragmatic engineering approach was considered as an alternative. On completion of [17], the Consortium took a decision to adopt the pragmatic engineering approach. [17] provided a useful model for consideration in implementing the system, but as described in Section 9, it was felt to be inappropriate to the scale of the MIPS project. The results of this work have therefore not been used directly within the MIPS project, though the algebraic specification is used as the basis for the implementation by STI of a Heterogeneous Database Integrator (HDI). The pragmatic approach is described in [21] and [22].

The technical and user documentation is provided in [25] and [26].

6.2.1 Data Model and Queries Description

This specifies and models a multidatabase system with a transparent access functionality from the user's point of view. This approach is applied to transparent querying document and heterogeneous relational databases, although in a strict sense the data model constructed only combines heterogeneous relational data since the document answers are converted to relations prior to any integration with any other data. Mathematical proofs are provided in support of the correctness and completeness of the model.

6.2.2 Selection and Retrieval Tool (SRT)

The input to the SRT is a query, initiated by a user browsing the HyTime Web and either written entirely by an Application Builder while creating the HyTime Application description or created by the user through the GQT.

The output of the SRT is a consolidated answer to that query, composed of data retrieved from local and remote heterogeneous databases. This answer is passed to the Web Builder which integrates the answer into the HyTime Application Description.

The SRT has been written to work with any MIPS application by drawing all domain specific information from the Knowledge Base System (KBS). The KBS is configured for the given application.

The task of the SRT, therefore, is:

- a) to transform an abstract user query expressed in terms of the concepts from the ontology into a set of concrete subqueries, appropriate to the databases to which they are targeted (the B&C sub-module). This is done in conjunction with the KBS. Semantic information is retained during this synthesis to ensure that the returned data can be built into an answer which corresponds to the original query. Each subquery is passed to the EDA.
- b) to query those databases and retrieve appropriate answer data (the EDA sub-module). Here, each subquery generated by B&C is targeted to the specified data source. This is effected through one of a set of separate submodules - one for each data source. On receipt of the answer to the subquery, the data is passed to the ACA.
- c) to consolidate that retrieved data into a single answer that reflects the semantic intention of the original query (the ACA sub-module). This includes, ensuring that data is returned, that data is in the right format (applying appropriate conversions if necessary), combining the data, ensuring that there is neither too little nor too much data. The semantic information gathered by the B&C sub-module assists this task. This consolidated answer is passed to the Presentation Subsystem through the Web Builder.

Communication between modules within the system takes place in a project-developed Internal Representation Language (IRL) using the SoftwareBus, dialects of which are used for communication between the EDA sub-module and the B&C and ACA sub-modules.

6.2.3 Web Builder

The purpose of the Web Builder is to connect the Selection and Retrieval subsystem with the Presentation subsystem. This consists of two main functions:

- to pass a query embedded in the HyTime Application Description (HAD) by the Application Builder from the Presentation subsystem to the Selection and Retrieval subsystem;
- to receive the query answer from the Selection and Retrieval subsystem and pass it to the Presentation subsystem, and to integrate it into the navigation and presentation structures of the HAD.

In order to integrate the returned data, it is necessary to know the structure and format of the answer. Normally, this will be provided in a template by the Application Builder in the HAD itself. However, where the query is generated from the GQT (or where the returned data does not fit the supplied template), the KBS will generate a new template. In this way, the Web Builder itself is independent of the MIPS application.

7. Demonstrator

The MIPS demonstrator was concerned with the travel and tourism industry in Greece. The two end users involved in this task were Ginis Vacances S.A., one of the top travel agents/tour operator in Greece and the Corfu Initiative, a Corfu island based non-profit organisation devoted to promoting the image of Corfu both locally and abroad.

The MIPS demonstrator for Ginis Vacances was aimed as a tool to assist prospective customers in selecting an area to visit. It was decided that the MIPS demonstrator present the islands of Corfu, Crete and Rhodes which are the major market for Greek tourism. In the function of Ginis Vacances as tour operator the demonstrator was designed to show package holidays, therefore promote the selling of those packages. With respect to the travel agent operations, the demonstrator was designed to act as a consultant to tourists and generate interest to visit a place.

For Corfu Initiative, the MIPS demonstrator was designed exclusively for Corfu island to function as an information kiosk. An example screen from the demonstrator is shown in Figure 6.

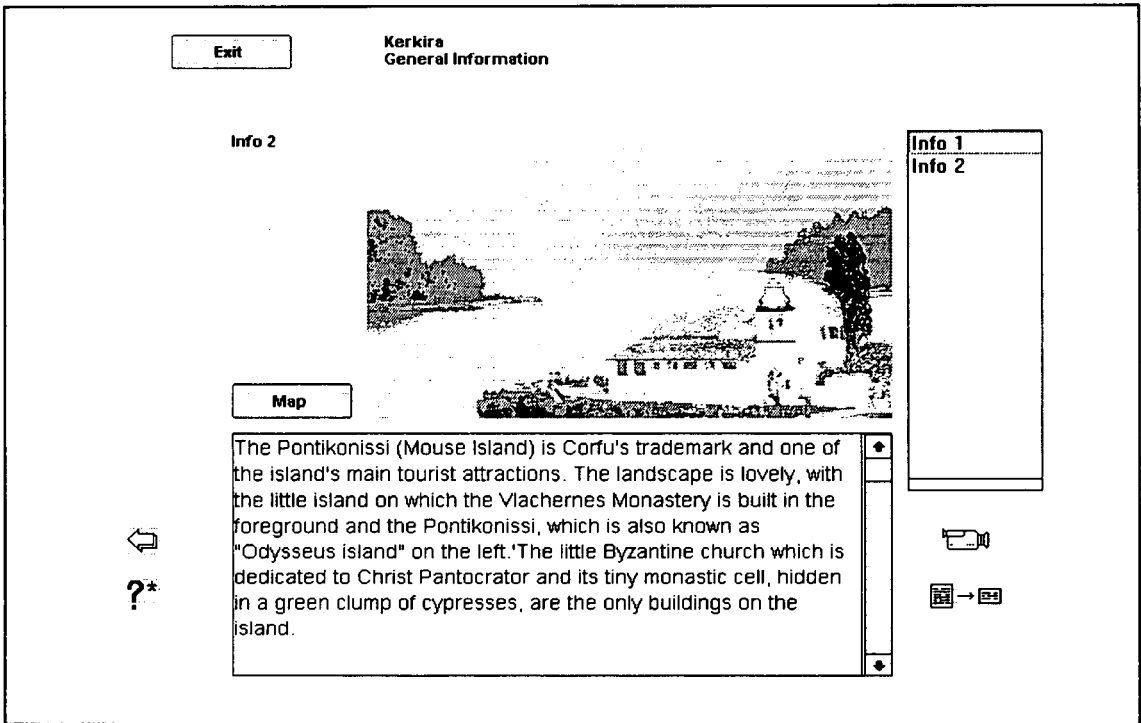


Figure 6 - An example screen from the Corfu MIPS demonstrator

The primary language for the demonstrator was Greek with English being the secondary one. Databases to be accessed contained information about accommodation, going out (food, entertainment), transportation and sports.

The sources of information were as follows. For text, an agreement with a publisher was reached (Adam Hellenic Reproductions SA) which has several quality publications for Corfu, Crete and Rhodes. Additional material was provided by end users. For images and narration, in-house material and the publisher's material were the main sources. The rights to use music performed by artists were purchased. For video, agreements with public organisations were reached while also local shooting was done, particularly in Corfu. Material for Corfu was the responsibility of the Corfu Initiative whereas material for Crete and Rhodes was collected by Ginis Vacances.

Corfu Initiative organised a large publicity campaign in Corfu for the MIPS demonstrator and secured collaboration from several key organisations concerned with tourism. Ginis Vacances provided its in-house material and collected assets from hotel owners.

A first version of the demonstrator design was ready in May 1994. This was produced with reference to the original Demonstrator Design document [7] in cooperation of Ginis Vacances. There were two requests for functionality not available originally in the MIPS toolkit, namely the incorporation of a "back" facility (to take the user back to the page last visited), and the addition of a table viewer. The second version of the demonstrator design was ready in early September 1994. Throughout the design phase extensive references were made to the MIPS CD-ROM with respect to the functionality of the MIPS toolkit. Asset collection had started from March 1994. The asset collection for Corfu was completed in October 1994 while their digitisation ended in mid-November 1994. The digitisation was done by Epsilon. Asset collection by Ginis Vacances for Crete and Rhodes finished in March 1995. To assist the data entry effort, specialised applications were developed by Epsilon in a PC database and delivered to Ginis Vacances.

In addition to including the assets in the MIPS overall architecture and demonstrating Scenario 4 operation the assets and application design were also used for creating an installed version of the MIPS demonstrator using an off-the-shelf tool as a test of user acceptance and appreciation.

Following a visit to Corfu in June 1995 by the MIPS project team, reviewers and Commission Project Officer, Corfu Initiative formally launched the MIPS demonstrator at a press conference in the Hilton Hotel. The demonstrator was well received by the users and is currently installed at several prime locations including Corfu Airport and the centre of Corfu Town as part of a tourist information centre (Figure 7).




Figure 7 - The information kiosk in Corfu Town with the installed demonstrator

8. Contribution to the state-of-the-art

So far in this report, the system requirements and the structure of the system developed have been described. This section describes those aspects of the development which are original contributions to information technology. The main reason for multipartner collaborative projects such as MIPS existing, and the advantage that they have over a single organisation undertaken the research and development themselves, is that previously independent groups of specialist researchers are brought together to integrate previously independent technologies. No single organisation has the skills required for this; no small organisation has the financial resources to support it; and the organisational cultures of large organisations usually prevent the independent innovation, and the subsequent engineering of the resulting distinct technologies into a single system. The architecture of MIPS embodies the results of this systems engineering task, where the design trade-offs which are understood in the languages of each contributing speciality must be broadened and translated to be re-evaluated over the whole set of technologies in the engineered system. To allow the development of an integrating architecture, an integrating framework had to be developed at a sufficiently abstract level to represent the component views on the system. [34] presents a layered model of the semantics of information retrieval tasks which extends the prevailing existing 3-layer model at the basis of most database system architectures (physical, logical and conceptual layers) to include two additional layers of intensional (between logical and conceptual to account for multiple heterogeneous data sources) and semantic (between conceptual and the user in order to account for the user's interpretation of the presented information) to produce a 5-layer model. This model remains one of the main supports of the integrated architecture of MIPS.

Beyond the integrating architecture, advances on the state of the art have been made by research in three areas:

- Access to heterogeneous data sources
- Hypermedia application description and multimedia data
- User Interfaces and HCI issues.

Each of these will be briefly described in the following sections, indicating where further information on each topic is available.

8.1 Access to heterogeneous data sources

The problem of heterogeneous data source access has previously been considered as a single problem, and several competing solutions have been suggested: universal schema, ontology, directory etc. Within MIPS we investigated the economic basis of the problem and the costs of the different solutions to identify when each was appropriate. The main variable was the number of data sources to which access was required. If this number was less than about 10, and the data sources were designed to be consistent and under the control of a single manager, then an approach based on a single universal schema could be applied. If these few

data sources are all on a local area network in an office where there is sufficient bandwidth available for extra traffic, and the traffic is effectively free, then it is also possible to broadcast all queries to all data sources in order to locate all possible answers. Indeed, in this situation, if the data sources not only include different relational databases, but also text information retrieval systems, each word in the IR systems can be encoded into the schema as an attribute, and searched over to retrieve all documents containing a target word. Since only small local IR systems will be used, the set of false positive targets retrieved, but not desired, will be small and the user can further constrain queries to focus them more. Similarly, other limitations on this approach do not apply in this environment: object oriented and deductive databases are not currently significant market opportunities in these circumstances so the lack of extensibility of the approach to them is not a problem; legacy databases with conflicting schema do not exist in these newly computerised SME's, so they do not give rise to the problems they do in larger organisations with IT histories; text IR systems are small and little changing, so the cost of indexing is negligible. This small scale solution to data was explored in [17] where the properties of a rich mathematical model of it were investigated.

In extreme contrast to this class are the hundreds of thousands of servers which have been installed on the World Wide Web while the project has been underway. The Web is too large to allow any universal schema, and communication costs are too high to allow any form of broadcast query. The solution on this large scale is to allow the use of directory services, indeed distributed directory services in order to distribute the cost of their maintenance. X500 is another directory service in this class which has been implemented to address the problems of large number of data servers.

Between these two solutions lies the case of a multimedia publisher accessing between 10 and 75 data sources across a combination of local and wide area networks where communication costs are significant. In this case, the constraints on the processing of the query itself (time to reply, cost of communications, cost of data source access, quality of service) may vary with individual tasks and queries but are sometimes going to rule out the directory service. Equally, the access costs are too high to allow query broadcasting, and the use of external data sources prohibits the use of universal schema. For this category the KBS ontology approach is appropriate, of using a generic set of terms through which individual schema and queries may be matched, and allowing the specification of processing constraints on each query to control quality and cost of service (see [32]). This is the solution that was adopted in MIPS in order to address the medium scale distribution of non-uniform legacy systems of different paradigms whose answers had to be integrated as described in the user requirements for the multimedia publisher task in tourism and other domains found in the user requirements.

A query language had to be developed which was expressive enough to capture the special properties of multiple media in the data, and to allow access to multiple heterogeneous data sources which varied not only in their schema, but also in their paradigm (relational, text, object oriented, deductive). The present standard query language SQL was designed not only

to be sufficiently expressive to capture relational queries, but also sufficiently concise to be written by hand by users experienced with the target database schema. Since its inception, its use has changed so that it is now mainly used by users who have no knowledge of the target database schema, through forms interfaces. Once the use of an intermediary programme between the user and the target query language is acknowledged, the need for the language to be concise or writable by humans goes away, and other features can be introduced into the language to address constraints on the processing of the query (time, cost, quality of data), possible multimedia features of the data (pictures of Paris, text about the Paris Metro), and a conceptual expressiveness which is sufficient to capture not only the relational paradigm, but also text retrieval systems, object-oriented, and deductive databases can be introduced. The Query Language resulting in the project from this reasoning is the IRL(QD) which is produced by user interface tools, and passed through the SRT to gather data from the External Data Sources (EDS). The structure of the IRL developed was designed to be of the first order logic language with second order syntax class of languages which are used as the output of natural language interfaces so that natural language input agents could be added later when they were available to supplement the forms interface provided by the General Query Tool (GQT). Since then the development of IRL(QD) agent interface languages of this class have moved towards standardisation through the ARPA Knowledge Sharing Initiative project, and the development of the Knowledge Interchange Format (KIF) and the Knowledge Query and Manipulation Language (KQML). In December 1994 the Object management Group (OMG - a 470 plus member software consortium) adopted the Common Object Request Broker Architecture (CORBA) as a standard specification for interoperability in heterogeneous distributed computing systems. ARPA is currently negotiating with OMG to adopt KQML too as a standard. In February 1995 ANSI approved a project to produce a draft American National Standard for KIF and has assigned this project to X3T2, the ANSI committee on data interchange. Following the development of IRL as an expressive query language, it has provided the basis of several arguments used in these forums to ensure that the required functionality is incorporated into the developing standards.

After a query in IRL(QD) has been passed to the SRT, it is transformed by the KBS into a tree structure where terminal nodes are mapped to external data sources (see [32]). When the project started, this ontological mapping approach was believed to be original, although during the course of the project several other projects have published research using a similar approach, reducing this achievement from a first proof of the approach to one of resolving detailed design options. This ontological mapping permits a syntactic method for capturing the semantics of conventional database representation languages by grounding each symbol either in an external data source attribute for relational databases; a document for text retrieval systems; an object for object-oriented databases, and a database term for deductive databases or in terms of other symbols lower down the hierarchy for non-terminal symbols. It is this tree that the B&C component of the SRT uses as its representation on which processing constraints can be applied efficiently for query generation and from which a simple data combination model can be generated called the intentional structure (see [36] and [37]). Once the data are returned from the external data sources, the ACA part of the SRT performs two tasks (see [31]). Firstly it applies transformation rules to the returned data to resolve any

semantic heterogeneity in the returned data, so that it is all reduced to a common semantic form (e.g. the same units of currency etc..). Secondly, ACA applies its combination engine utilising only simple set operations grounded in the relational calculus to the data according to the instructions in the intentional structure. The novel achievement in this process is not the use of an ontology to map from a defined language to the terms used in external data source schema, but the derivation from this of a simple combination model which can efficiently be implemented within the known and optimised technology of the relational calculus.

8.2 Hypermedia application description and multimedia data

The conventional way to produce a multimedia application is by authoring it in a proprietary tool into a discrete hypermedia document which the user can then browse. However, the content, presentation information and navigation information are all stored together in the document so that the content cannot be used elsewhere. The format of the document is proprietary to one company and it cannot be transformed into an interchange format. The user has no access to documents outside the hypermedia document which was originally authored.

The ISO standard for hypermedia information HyTime was chosen as the representation to be used in MIPS to overcome the lack of interoperability resulting from proprietary formats. The HyTime standard was approved as the MIPS project started, and the development of a generic HyTime engine to interpret any HyTime Document Type Definition (DTD) is a novel achievement of the MIPS project, although several others have been developed in parallel.

Details on the application of HyTime within MIPS can be found in [41], [42] and [43].

The open platform independent format used in MIPS was defined as a HyTime compliant DTD which supports the separate encoding of navigational and presentation aspects of hypermedia documents in order to allow the content of a document to be used in many different presentations, but only be stored once.

One construction included in the DTD was the HyTime Application Description (HAD) template concept which allows hypermedia presentation generation from external datasets by the instantiation of data into templates specifically written for them. The result is that an application document can be written to include queries with corresponding templates at a node instead of the node contents, presentation and navigation information. When the query is evaluated, the returned data can be instantiated into the templates to construct the complete node for presentation.

When the data returning as the answer from a query do not exactly correspond with a pre-written HAD template, dynamic multimedia layout rules can be used to construct new templates appropriate with the data available from existing primitive templates. In this case, not only are templates on a node instantiated, but new nodes or web fragments are generated including links between nodes. By this mechanism, an application document can be grown

by issuing queries to external data stores, once the initial primitive templates and template manipulation and construction rules have been written (further information on presenting dynamically expandable hypermedia can be found in [35] and [40]).

8.3 User Interfaces and HCI issues

When describing heterogeneous data source access, the increased expressivity of the IRL(QD) over conventional query languages to include multimedia aspects and processing constraints as well as a conceptual description was discussed. It was established that a query tool is appropriate for users without experience in the tool or target domain to construct queries rather than expecting them to write queries. Established forms query interfaces could be written to exploit the increased expressiveness of the IRL(QD), but these would still be very restrictive. Forms query interfaces usually only allow variations in the parameters within query formulae (e.g. to set the departure and return dates of a package holiday) and do not allow the whole conceptual structure of queries to be varied (e.g. changing operators, functions and tables to be searched). Therefore for each variation on the conceptual structure of a query, a new forms interface would have to be written. This contrast between the extremes of a richly variable hand written query or command language, and a fixed forms or menu interface has been much considered in HCI. The user class we are considering in MIPS is not expected to develop a rich enough knowledge of IRL(QD) or the underlying ontology to be able to write queries by hand. However, a conventional forms interface would be too restrictive, so a new form of query tool needed to be developed with more flexibility, yet which would support the users tasks without needing them to adopt an additional burden of learning a command language.

The approach adopted is to provide a query tool which allows the user to write a large set of queries for a particular task which vary in the operators and functions to be included in the formulae, but also in the medium types to be retrieved, the security and processing constraints on the query itself (time and cost of retrieval), and the layout of the final report generated from the retrieved data. In order to construct such templates for each user task an analysis of the user's internal representation of the task is required. From this an external representation of the task can be defined in the task template language which is consistent with the user's internal representation. The consistency of the internal and external task representation should minimise any additional learning burden on the user due to task structure or performance (see [39] for a complete description of the developments associated with task templates and the query tool).

The task template includes more than just a task description since it also includes a definition of possible dialogues. The graphical user interface design used is generated from these templates. The dialogue structure has been reduced to minimise each interaction, so that each screen establishes one clause for a query. Although this approach minimises the possibility of omitting a clause, it does pose a memory load on the user to remember the set of clauses already considered as a query is progressively built up. To overcome this burden, the user is provided with a natural language version of the query as it is progressively built up. This

feedback is generated from the task template and can include icons to clearly describe the query and minimise the memory burden on users (see [33] for a complete description of the developments associated with progressive query construction and natural language feedback in the query tool).

Hypermedia systems usually contain a limited set of screen objects which define a strict style within which a page designer can define each node. For graphical interface designers working within GUI toolkits this constraint results in the “look and feel” of each GUI development product swamping any attempt to develop a house style for products (e.g. MS-Windows or Motif dominates ICI or Ford in the corporate identity of the GUI). To overcome this, the Multimedia Presentation Tool developed in MIPS had a unique appearance designed by a graphic designer so that the graphics design set the direction for the software design rather than the other way around. The resulting design was for the presentation of multimedia assets in a seamless window with common interaction mechanisms, rather than just through viewers for each format as some WWW browsers support. The MIPS Presentation Tool resulting from the implementation of the design is used both for presenting hypermedia application documents & retrieved data, and for presenting query templates in the query tool in order to provide a consistent user interface to MIPS. During implementation, when conflicts arose between the design need for a feature and the resulting software complexity or cost, trade-offs were always biased towards the design need. The design framework that was developed to create the look and feel of MIPS is intended to be extensible not only to new interface objects and information types, but also to new design and user tasks. This design and development process was facilitated by a design team at Trinity College Dublin consisting of a wide range of skills from graphic designers, audio-visual technicians and software engineers to an experienced paper publishing manager (see [38] for a complete description of the developments associated with design and implementation of the MIPS Multimedia Presentation Tool).

9. Results availability

There are several results which have arisen from the MIPS project which the partners intend to commercialise and exploit ranging from licensing the MIPS technology and toolkit and its constituent parts to using MIPS as the basis of an information product production platform. These are described below.

9.1 MIPS Toolkit

9.1.1 Cartermill International Ltd

Cartermill International (previously Longman Cartermill) brings together consulting competencies in new media and information management services and a publishing portfolio with an impressive coverage of key reference information worldwide. The Company's publishing and consulting activities are specialised in the following product markets: Science, Technology and Industry, Health and Social Care Reference, Business Intelligence and Current Affairs.

The Company's publishing activities bring together many leading brands in the form of CD-ROMs, online databases, journals, directories and reference books. Long established titles include the FT Yearbooks - the FT Mining Yearbook dates back to a Mining Manual first published in the year of Queen Victoria's Golden Jubilee in 1887 and published annually ever since. The Current Affairs list can boast individual titles of great history and distinction such as the veteran Annual Register, first published in 1758, the oldest continuing reference work of its kind in English, and Keesing's Record of World Events, a monthly journal first published in 1931. The past decade is now available on CD-ROM. New book series include Spicers European Union Policy Briefings.

Other key published product includes the Medical Directory, also available on CD-ROM, with entries on every doctor registered to practise in the UK, Current Research in Britain (CRIB), co-published with the British Library, providing immediate access to over 62,000 research projects from academic and related institutions in the UK, while the BEST (Building Expertise in Science and Technology) databases provide comprehensive coverage of key competencies in the UK's universities and government research establishments. BEST also covers a growing number of European countries and has an exclusive agreement with France's CNRS (Centre National de la Recherche Scientifique) to cover the research activities of over 10,000 researchers in CNRS laboratories.

Our CD-ROM authoring, production and distribution facilities are mainly used on in-house information products but are also available on a commercial basis to third parties and are directed at creating CD-ROM products using information from online databases, existing CD-ROM and paper based sources. In addition to production, services include market analysis and design, production of documentation, sales and marketing, and post-sales support.

Cartermill considers the MIPS technology as a means of re-engineering the publication process to maintain and grow the market share to achieve 80% growth between 1995 and 2000. When these figures are viewed against the background of the shift in Cartermill's delivery media from 80% hard copy in 1995 to a projected 75% electronic (of which 25% will be network publishing) by 2000, the pivotal role of MIPS-based applications in the Company's future is clear. Making these transitions and growing successfully almost entirely depends on the effective use of MIPS along with some comparable and complementary technology in the short to medium term.

There are two basic ways in which Cartermill, as a user, envisage exploitation of the MIPS toolkit:

- development of products and services based on the tools developed; the tools resulting from the project will allow us to offer services that our competitors cannot, and to incorporate new product features into our publications that will give them strong competitive advantage.
- additional knowhow in developing products and carrying out contracts; our customers expect us to be able to deliver a wide range of information processing solutions outside the pure text area. Access to multimedia technology expertise in general will allow us to keep our skill at the necessary level to satisfy these customer demands.

Cartermill is keenly interested in the publishing opportunities arising from the use of the MIPS toolkit in the development of new products and services within its broad areas of interest. The initial focus is on current and short-term opportunities in the environmental information market and the need in this segment for multimedia and information obtained from multiple sources. The role of MIPS will be as a dynamic interactive aid to the authoring of multimedia presentations which can then be frozen for delivery to end users either on-line or by CD-ROM.

Based on the success of this initial application we anticipate using the MIPS toolkit in areas such as medical, financial and business publishing. These professional publishing areas will benefit especially from the ability to deliver information from multiple sources in an effective way for non-expert end-users.

9.1.2 The Danish Technology Institute

The Danish Technological Institute, with its 1250 employees of which 800 are consultants, organized in 22 departments, each with its specific area of expertise, is Denmark's largest supplier of technological services including industrial management. The institute is a self governing nonprofit private organisation. Two thirds of the Institute's turnover is covered by income derived from clients, while the remaining one third originates from programme and project support from the Danish Ministry of Industry.

The institute's most significant task is to assist SMEs in their development in human, economical and technical areas. The institutes consultants carry out more than 15,000 paid assignments annually, and more than 20,000 participants benefits from our courses every year. To carry out this task properly a high level of competency is maintained within all areas that can affect the rentability and survival of the manufacturing enterprises.

One of the institute's principal tasks is to assist companies in their implementation of new technology. New technology alone is, however, not sufficient. Successful company development implies a fine balance between management strategies, technology, staff training, motivation, organisational development and working environment. The Institute therefore operates with integrated solutions adapted to individual requirements. The Institute's activities make use of forms of actions such as: Industrial R&D, Consultation, Testing, Training and Information services. The service to enterprises and other professional organizations often comprises a combination of these activities.

DTI is convinced that multimedia is extremely important in the future and will be used in most functions in the modern industrial company!

DTI are offering to assist industry and commerce to benefit from the MIPS results in a number of different ways:

- 1) by making available licences for the use of the MIPS tool-kit through the joint exploitation structure as an product to:
 - Suppliers of Value Added Services, in particular those related to on-line databases.
 - Software developers of industrial information solutions.
 - Industrial manufacture and service companies having the need of a wide range of heterogeneous information
 - OEMs
- 2) by developing applications and application shells with the MIPS-Toolkit in co-operation with VARs for:
 - Technical manuals
 - Preventive maintenance manuals
 - QA manuals
 - Various other Instruction/training applications specially in areas where the information comes from various heterogeneous databases.
- 3) by organising and conducting MIPS training courses/seminars in corporation with the other MIPS partners
- 4) by providing consultancy services to Scandinavian industrial companies

(SMEs) in their potential and strategic use of multimedia based on the components and knowledge acquired in MIPS.

- 5) by using MIPS internally, for one or more of the following:
- Building applications for training and educational purposes (CBT in areas where heterogeneous information from many sources is important.)
 - Utilising the MIPS-tool-kit in our existing information providing services for the Danish industry by our DTI EU-Infocentre, DTI dissemination-centre and the 15 regional Technological Information Centres (Special services to SMEs)
 - Configuring MIPS-services to some of our 28 departments in their access to various heterogeneous databases and thereby offering improved services to the Danish SMEs
 - Integrate the MIPS tool-kit with DTI products, to provide- multimedia facilities and integrated information from heterogeneous databases
- 6) by tailoring, populating and re-configuring parts of the demonstrator application to suit:
- Scandinavian Charter tourism enterprises
 - Scandinavian Travel agencies
 - County tourist offices

9.2 MIPS Modules

9.2.1 Sema Group

In the context of the MIPS project (ESPRIT 6542), Sema Group developed a toolbox based on the ISO standard HyTime (ISO/IEC 10744:1992).

The MIPS HyTime toolbox is made of five main components:

- an SGML Engine,
- a HyTime Engine,
- a HyTime-compliant DTD,
- a Presentation Manager,
- a Web Builder.

The MIPS SGML Engine is a C++ class library to manipulate SGML documents as object-oriented structures. It imports SGML-compliant files and transform them in trees, it allows client programs to access any SGML information and navigate around the document

tree. This is ideal for EDMS applications requiring global analysis of SGML documents.

The SGML Engine is currently implemented above ObjectStore, the OBMS from Object Design; its import facility uses the Sema SGML parser Mark-It.

The MIPS HyTime Engine is also a C++ class library. It is implemented above the SGML Engine. Supporting the powerful addressing and hyperlinking facilities defined by the HyTime standard, it allows applications to browse and process HyTime hyperdocuments made of a set of SGML documents. Cross-references can be internal or external, they can locate text or multimedia assets. The MIPS HyTime Engine address resolution facility is extensible by means of applicative hooks.

The MIPS DTD has been specifically designed to encode in an open platform-independent way hypermedia application descriptions. The MIPS descriptions are sets of HyTime-compliant SGML documents organized in two parts: a navigation structure and a presentation structure.

The navigation structure is described as a set of typed information nodes connected by hyperlinks; an information node groups together all the multimedia assets about a domain entity; the hyperlinks encode the domain relationships of interest for the intended end-users.

The presentation structure determines the way the information is rendered to the end-users and the behaviour of the application. The main components are window models and scripts. The window models determine the look and feel of the application user interface, while the scripts define the application reaction to user interaction.

The MIPS DTD also supports the concept of template, which allows generic description of application fragments; these generic descriptions are then instantiated before or during runtime; the instantiation process is parameterised by datasets encoded as nested relational tables. These tables can be written by hand or generated by evaluating some queries, possibly embedded in the description itself.

This unique facility can be used to build template-based report generators which will integrate the contents of relational tables within a hypermedia presentation.

The MIPS Presentation Manager is a generic playback system able to interpret any MIPS hypermedia application description. Whenever it traverses an hyperlink and reaches an information node, it builds a window according to the window model specified in the description and associates with the different basic viewers the available multimedia assets. Whenever the end-user clicks on a mouse-sensitive area of the window, it retrieves and executes the script specified in the description for the concerned area; the script can either modify the current view or traverse a hyperlink in order to reach another information node.

The MIPS Presentation Manager is implemented as a C++ application above the generic

MIPS HyTime Engine. It currently runs on SunOS 4.1.3 with a presentation dispatcher running on MS Windows 3.1.

The MIPS Web Builder is a reusable component of the MIPS Presentation Manager; it implements the instantiation process defined by the templates. The instantiation process generates new description fragments according to the template definition and replace within the new fragment

placeholders with values found in the dataset used as parameter for the instantiation. The Web Builder can be called either before playback or during playback to generate the needed description fragments.

9.2.2 Rutherford Appleton Laboratory

RAL has recently established a company called 'Web Spyder' to market multimedia products. MIPS based consultancy services are being marketed through the Web Spyder identity to provide income for investment in the development of products.

RAL, and particularly Systems Engineering Division, has experience in building and producing products in the Knowledge-Based System field. These have been used in various projects, both funded by ESPRIT and by Commerce and Industry.

RAL is rapidly developing its business interests, from a base of almost nothing 5 years ago to a turnover of £20 million per annum. The market sectors in which RAL are involved cover the full spectrum from financial services through to heavy engineering, and are world-wide. All the sectors require the facilities of the MIPS toolkit, and all specifically require sophisticated access to heterogeneous distributed databases such as provided by the SRT and KBS modules.

Based on its development of the MIPS KBS module and their involvement in the SRT RAL are offering consultancy services to solve multiple distributed heterogeneous data source access problems for medium to large sized companies in the UK. The consultancy will include using the KBS and non-EDA SRT components being marketed through Web Spyder.

RAL are also seeking opportunities to take the hypermedia design and presentation technology into further national and European R&D projects in order to refine the use of the HyTime engine in the area. RAL would like to develop partnerships for joint projects to develop technology, product descriptions and product positioning when a market evolves in this area by about 2005.

In addition RAL support the commercial needs of client organisations by providing technical innovation from their own resources or by introducing them through European universities of fellow members of ERCIM. This role of innovation brokerage as a basis for establishing strategic alliances with commercial developer companies, and then supporting future product development, enhancement and migration is one for which RAL is ideally suited.

9.2.3 STI

STI intend to commercialise the External Database Access module with further developments to produce and market a Heterogeneous Database Integrator (HDI).

EDA has been designed with a twofold objective: as a genuine MIPS product, and yet as a MIPS independent commercially exploitable product. For such a purpose further development - outside MIPS - will be carried out to provide EDA with an appropriate user interface, and API Generator tool and an information (answers) integrator based on the results of T4.1. This is hereby described.

EDA will be provided with a user interface and an HDI much in the same way as specified in [17]. Such an independent system will be targeted immediately to customers in the healthcare area so that they can query databases either locally or across a network. For such a purpose further development will be carried out to provide EDA with a windows-based user interface to enable users to parameterise it: determining the databases to access thus activating or disabling the corresponding EDA submodules, establishing the communications environment they are using.

The HDI will, in principle, provide transparent access to distributed heterogeneous relational and document databases being further extended to incorporate deductive and object oriented databases.

STI will develop an API Generator tool following the initial studies carried out in MIPS which will be incorporated to the EDA to access those databases not provided with APIs.

In addition of being integrated with EDA and commercialised with it, the HDI will be exploited independently. Agreements are already established with middleware vendors to integrate HDI into their products. Contracts have been signed with two Spanish hospitals and the product is expected to be finalised and running in the hospital environment in one and a half years time.

9.2.4 Heriot Watt University

Heriot-Watt University has one of the leading database research groups in the UK. In addition to the MIPS project, work is currently progressing on deductive object oriented databases, active databases, and parallel database systems. These are funded both by the EU and by the UK EPSRC.

Under MIPS Heriot-Watt has developed the BACA module which takes a high-level query expressed in concepts familiar to the user, and progressively refines it to a level at which it is solvable by a set of subqueries aimed at target databases. This refinement is aided by information stored in the KBS. The fully refined queries are handled by the EDA module which receives the responses from the target databases.

BACA also takes the answer fragments from the EDA, transforms these in an appropriate way and assembles them into a form to be input into a HyTime web in the WB.

Being an educational establishment, the main market for the exploitation of MIPS by Heriot-Watt is in education where the continuing growth in student numbers in the UK requires the exploration of innovative teaching methods. This will require flexible and sophisticated querying of heterogeneous databases, as provided by the BACA module.

UNILINK is the University's industrial liaison arm through which commercial exploitation of MIPS will be channelled. This includes Heriot-Watt's role in any joint venture undertaken by the consortium, the development of specific new applications (eg with the Edinburgh International Festival), commercial exploitation of innovative teaching applications, and future exploitation of further research work.

The University is in close contact with other research institutions. In particular, it is involved in the UK Teaching and Learning Technology Programme, through which innovative teaching programmes are created and evaluated. In this way, Heriot-Watt has the opportunity to explore the development of MIPS in the educational field.

Through Heriot-Watt's other related research work, it is intended to use BACA in other research projects, and to develop the GQT/BACA/EDA/KBS as a stand alone query tool along with the other MIPS partners concerned.

Heriot-Watt is particularly interested in discussing with industrial concerns the potential for joint projects which can capitalise on their experience and expertise in all aspects of innovative approaches to databases.

9.2.5 Trinity College Dublin

TCD as a 400 year old university situated in the centre of Dublin has strong linkage to the very vibrant software sector of the Irish economy. TCD has considerable experience in bringing research projects to market as there exist some 25 "campus companies" employing over 150 people producing products which mostly started as research projects. In addition a number of former campus companies now trade in their own right outside the college walls, many in partnership with international software companies such as Sunsoft.

The MIPS Presentation Tool (MPT) developed by TCD in MIPS is a seamless and transparent presentation mechanism for multimedia assets. It runs in a MS Windows environment and displays text, graphics, animations, graphs and video as well as playing audio files. The MPT was developed by TCD in conjunction with Cartermill International.

It represents a particular solution to the problem of presenting different multimedia assets encoded in different formats without making the screen full of clutter with consequent confusion for the user. The different assets are presented in seamless windows so that the who

presentation appears to be in a single window This means that the user has a single method of interaction and not a variety of menus and tool bars and the use of the product is transparent in that the user need not know or worry about the formats used.

The MPT will be used in an academic context as a "front end" to multimedia research for both undergraduate and postgraduate students. This research would include, but not be limited to, research in the area of object oriented authoring tools for multimedia and hypermedia presentation. In this context the MPT will need to be extended and enhanced in order to become an integral part of the new tool.

TCD will also consider partnership with private companies, in the context of a campus company, to enhance and exploit the MPT as a presentation mechanism. In such a partnership the commercial partner would be expected to provide the market placement for the product and would bring their commercial expertise to the partnership for issues such as pricing as a complement to the technical expertise and innovative approaches in TCD.

9.3 MIPS Demonstrator

9.3.1 Epsilon Software

Based on the MIPS results Epsilon Software offer a range of consultancy and system development services in the area of information kiosks, and more specifically in the retailing and services sectors. In the former sector, Epsilon is active in the area of department stores and top-market estate agents. In the services sector, the key areas are tourism-related business and government services.

For sophisticated access to diverse databases MIPS has a wealth of features not found in mainstream commercial off-the-shelve multimedia authoring tools.

An information kiosk based on MIPS has as direct advantage the provision to a customer of information of superb quality with minimal assistance from staff. This leads to higher personnel efficiency and to a reduction of the time necessary to provide information to customers. As an indirect advantage, MIPS may dynamically change the multimedia presentation according to a particular user profile, thereby providing the opportunity to provide additional information tailored to the user's interest.

With respect to the cost of a MIPS application for an information kiosk, indicative figures are shown in the following table.

Service type	Cost for first MIPS workstation	Cost for additional MIPS workstation	Maintenance cost per year
Information kiosk for direct use	30 KECUs	10 KECUs	8 KECUs
Information kiosk for indirect use	15 KECUs	10 KECUs	3 KECUs

10. Conclusions

MIPS has been a highly successful project with the major technical objectives achieved in line with the partner business motivations for investing in the work. The results of MIPS as a toolkit and as a series of modules are expected to contribute significantly to the industrial partner business objectives and the research partner work has made a substantial contribution to the state-of-the-art.

MIPS primary success can be seen as an enabler of integration both in terms of integrating information from heterogeneous data sources and in terms of integrating diverse software development methods and process through the use of the Software Bus.

One of the major technical achievements of MIPS is the development of the first European HyTime Engine.

Other major technical achievements include the development of methodology and proven architecture for open format hypermedia presentation including a hypermedia report generator facility. The reusability of multimedia assets from heterogeneous data sources, and the re-use of the web structure, nodes and presentation mechanisms are highly important to providing flexible solutions to the production of multimedia presentations particularly where information from remote sources has a time dependence.

The architecture, methodology and software has been validated by the development of a demonstrator in the tourism domain which is itself a contribution to the development of that market. The demonstrator has also provided a basis on which the system can be tailored to other applications and re-engineered to meet performance and functional targets.

Seven members of the MIPS team have transferred their skills from academic to industry - mainly small software companies. Three doctorates and one MSc have been awarded to members of the MIPS team for their contributions to the state-of-the-art.

However, by far the major success of MIPS has been the creation of the potential for the partners to bring the benefits of MIPS based solutions and products to their clients.

Appendix 1

Information Available

Information Available

The following information is available on the MIPS project:

MIPS leaflet
MIPS CD-ROM
MIPS Conference Papers

Requests should be sent to:

Jack Donald
MIPS Project Manager
Cartermill International
Technology Centre
St Andrews
Fife KY16 9EA

Tel: +44 1334 477660
Fax: +44 1334 477180
e-mail: jackd@cartermill.com

Appendix 2

References

References

- [1] MIPS Consortium (1992). *Overall Functional Specification of the System*. MIPS Project, D2.1.1.
- [2] MIPS Consortium (1992). *Overall Architectural Design of the System*. MIPS Project, D2.2.1.
- [3] MIPS Consortium (1993). *User Requirements, Volume 1 : Requirements Listing*. MIPS Project, D1.1.1.
- [4] MIPS Consortium (1993). *User Requirements, Volume 2 : Appendices*. MIPS Project, D1.1.1.
- [5] MIPS Consortium (1993). *Technical Requirements*. MIPS Project, D1.2.1.
- [6] MIPS Consortium (1993). *Summary of User Requirements*. MIPS Project, D7.1.1.
- [7] MIPS Consortium (1993). *Demonstrator Design*. MIPS Project, D5.1.1.
- [8] MIPS Consortium (1993). *Product Objectives and Definition*. MIPS Project, D6.1.1.
- [9] MIPS Consortium (1993). *The definition of the MIPS system, Volume 1: The detailed architectural design of the MIPS system*. MIPS Project, D2.2.2.
- [10] MIPS Consortium (1993). *The definition of the MIPS system, Volume 2: The detailed functional specification of the interactive modules*. MIPS Project, D2.1.2.
- [11] MIPS Consortium (1993). *The definition of the MIPS system, Volume 3: The detailed functional specification of the HyTime modules*. MIPS Project, D2.1.2.
- [12] MIPS Consortium (1993). *The definition of the MIPS system, Volume 4: The detailed functional specification of the KBS module*. MIPS Project, D2.1.2.
- [13] MIPS Consortium (1993). *The definition of the MIPS system, Volume 5: The detailed functional specification of the SRT module*. MIPS Project, D2.1.2.
- [14] MIPS Consortium (1993). *The definition of the MIPS system, Volume 6: The IRL definition*. MIPS Project, D2.3.1.
- [15] MIPS Consortium (1993). *The definition of the MIPS system, Volume 7: The MIPS technical glossary*. MIPS Project, D2.2.2.

-
- [16] MIPS Consortium (1993). *User Interface Tools*. MIPS Project, D3.1.1.
 - [17] MIPS Consortium (1993). *Data Model and Queries Description*. MIPS Project, D4.1.1.
 - [18] MIPS Consortium (1994). *Knowledge Acquisition for the MIPS KBS*. MIPS Project, D1.3.1.
 - [19] MIPS Consortium (1994). *MIPS CD*. MIPS Project, D7.2.1.
 - [20] MIPS Consortium (1994). *Presentation and User Interface Subsystem : Development Report*. MIPS Project, D3.2.1.
 - [21] MIPS Consortium (1994). *SRT Subsystem : Development Report*. MIPS Project, D4.2A.1.
 - [22] MIPS Consortium (1994). *External Data Access Subsystem : Development Report*. MIPS Project, D4.2B.1
 - [23] MIPS Consortium (1995). *Presentation and User Interface Subsystem : Technical and User Documentation, Volume 1*. MIPS Project, D3.6.1.
 - [24] MIPS Consortium (1995). *Presentation and User Interface Subsystem : Technical and User Documentation, Volume 2*. MIPS Project, D3.6.1.
 - [25] MIPS Consortium (1995). *Selection and Delivery Subsystem : Technical and User Documentation*. MIPS Project, D4.4A.1.
 - [26] MIPS Consortium (1995). *External Data Access Subsystem : Technical Documentation*. MIPS Project, D4.4B.1.
 - [27] MIPS Consortium (1995). *Demonstrator and Evaluation Report*. MIPS Project, D5.5.1.
 - [28] MIPS Consortium (1995). *Market and Competitor Analysis*. MIPS Project, D6.2.1.
 - [29] MIPS Consortium (1995). *MIPS Exploitation Plan*. MIPS Project, D6.3.1.
 - [30] MIPS Consortium (1995). *MIPS Press Book*. MIPS Project, D7.2.2.

Appendix 3

MIPS External Publications

MIPS External Publications

- [31] Austin, W.J., Hutchinson, E.K., Kalmus, J.R., MacKinnon, L.M., Jeffery, K.G., Marwick, D.H., Williams, M.H., Wilson, M.D. (1994) *Processing Travel Queries in a Multimedia System*. Proceedings of ENTER '94 Conference, Innsbruck.
- [32] Behrendt, W., Goble, C.A., Hutchinson, E., Jeffery, K.G., Macnee, C.A. Wilson, M.D. (1993) *Using an Intelligent Agent to Mediate Multibase Information Access*. Proceedings of the Workshop on Cooperating Knowledge Based Systems, Deen, S.M. (Ed.), Keele University Press: Keele, UK, 27-43.
- [33] Lisa Dunphy, (1995) *Holly: A Domain Independent Query Interface with Natural Language Feedback*, MSc Thesis, Trinity College, Dublin.
- [34] Jeffery, K.G., Behrendt, W., Macnee, C.A., Wilson, M.D., Kalmus, J.R., and Hutchinson, E.K (1994) *A model for heterogenous distributed database systems* Proceedings of BNCOD '94, Lecture Notes in Computer Science Series, Springer Verlag.
- [35] Macnee, C.A., Behrendt, W., Wilson, M.D., Jeffery, K.G., Kalmus, J.R., and Hutchinson, E.K. (1995) *Presenting Dynamically Expandable Hypermedia*, Information and Software Technology, 37 (7) 339-350.
- [36] MacKinnon, L. (in press) *A Model for Query Decomposition and Answer Construction in Heterogeneous, Distributed Database Systems*, Information and Software Technology.
- [37] MacKinnon, L. (in preparation) *Intelligent Query Manipulation for Heterogeneous Data Sources*, Ph.D. Thesis, Heriot-Watt University, Edinburgh.
- [38] McGuigan, R. (in preparation) *A HyperMedia Object Manager for an Open HyperMedia System*. PhD. Thesis, Trinity College, Dublin.
- [39] Riordan, M. (in preparation) *Querying Multimedia Data in Distributed Heterogeneous Information Spaces*, PhD. Thesis, Trinity College, Dublin.
- [40] Wilson.M.D. (1995) *Enhancing Multimedia Interfaces with Intelligence*. In J.A. Vince and R.A. Earnshaw (Eds, *Multimedia Systems and Applications*. Academic Press: London, 23-51.
- [41] Bruffaerts, Albert (1994) *Encoding hypermedia applications with HyTime: The MIPS approach*. Proceedings of SGML Europe '94, 16-19 May 1994, Montreux, pp. 231-242

-
- [42] Bruffaerts, Albert (1994) *The role of HyTime in the MIPS system*. Proceedings of the SGML Users Group Holland Congress, 10 November 1994, Amsterdam, pp. 45-52
- [43] Bruffaerts, Albert (1995) *Using HyTime for complete hypermedia application descriptions*. Proceedings of SGML Europe '95, 16-19 May 1995, Gmunden, pp. 33-38; also in OII Spectrum, Vol 2, No 11, pp. 11-18

Appendix 4

MIPS Consortium

MIPS Consortium

Mr Jack Donald
Cartermill International
Technology Centre
St Andrews
Fife KY16 9EA
Scotland UK

Tel: +44-1334-477660
Fax: +44-1334-477180
email: jackd@cartermill.com

Dr Michael Wilson
Rutherford Appleton Laboratory
Chilton
Didcot
Oxfordshire OX11 0QX
UK

Tel: +44-1235-446619
Fax: +44-1235-445831
email: M.D.Wilson@rl.ac.uk

Ms Ana Martinez
Sistemas y Tratamiento de Informacion, SA
Paseo Castellana, 180
28046 - Madrid
Spain

Tel: +34-1-345-1349
Fax: +34-1-359-6022
email: sysnet@bitmailer.net

Albert Bruffaerts
Sema Group Belgium SA
rue de Stalle 96
1180 Brussels
Belgium

Tel: +32-2-508-5242
Fax: +32-2-512-1499
email: albert.bruffaerts@systems.sema.be

Professor Howard Williams
Department of Computing & Electrical
Engineering
Heriot-Watt University
Riccarton
Edinburgh EH14 4AS
Scotland, UK

Tel: +44-131-451-3430
Fax: +44-131-451-3431
email: howard@cee.hw.ac.uk

Dr Jane Grimson
Department of Computer Science
Trinity College Dublin
Dublin 2
Ireland

Tel: +353-1-7021780
Fax: +353-1-6772204
email: jane.grimson@cs.tcd.ie

Dr Dimitrios Gritsis
Epsilon Software SA
377 Syngrou Avenue
18564 P. Faliro
Greece

Tel: +30-1-9308370
Fax: +30-1-9427719

Mr Kristian Hansen
Industrial Management
Danish Technological Institute
Gregersensvej, PO Box 141
Taastrup DK-2630
Denmark

Tel: +45-4350-4740
Fax: +45-4350-7250
email: krh@dti.dk