

Media-based Navigation with Generic Links

Paul H. Lewis, Hugh C. Davis, Steve R. Griffiths, Wendy Hall, Rob J. Wilkins

The Multimedia Research Group
Department of Electronics and Computer Science
University of Southampton, England, SO171BJ.
E-mail phl@ecs.soton.ac.uk

ABSTRACT

Microcosm is an open architecture hypermedia system in which documents remain in their native format and link information is held in separate link databases. This has facilitated the introduction of generic links which, once authored from a text string to a destination anchor, may be followed from any occurrence of the text string in any document. The generic link provides substantial reductions in authoring effort for large hypermedia systems, but the limitation of the generic link to text string source anchors needed to be addressed.

This paper describes extensions to the Microcosm architecture to create MAVIS, Microcosm Architecture for Video, Image and Sound, in which generic links may be used from both text and non-text media. This development makes it possible to navigate through non-text media using content as the key and, through the facilities of the dynamic link, content based retrieval is also available. Examples of content based navigation with image, video and sound are presented.

KEYWORDS: Open Hypermedia, Content Based Navigation, Information Retrieval.

INTRODUCTION

The provision of tools for retrieving or for navigating through non-text information lags a long way behind the provision of tools for handling text based information. For the users of hypermedia systems, the tasks of navigating through large textual information bases or retrieving and displaying specific pieces of text are well supported by text retrieval and text processing tools. For the authors of these hypermedia projects, tools are available for creating the documents and establishing the required links. In the majority of cases, the extension to non-text information has been handled by reverting to text as the basis for retrieval and navigation. Thus retrieval and navigation with images in multimedia systems has been tackled for many years by the use

of associated and manually assigned key terms or the use of manually authored point to point links.

By contrast, the development of hypermedia techniques which capitalise on the direct use of content of image, video or sound as the basis for retrieval and navigation are in their infancy. Research which can contribute to progress in this area comes not only from the multimedia community, but also from image database research, research in image processing and computer vision and also research into sound and speech processing.

The hypermedia system user, with an image on the display, might like to be able to select the image or a sub-image and ask for other images or sub-images which are similar in some way to the particular selection. If the system finds these images or sub images by using the content of the selection directly to match in some way with the images or sub-images in the hypermedia system, we refer to this as content based retrieval (CBR). If we ask for further information relating to the selection and following a link to that information using the content of the selection as the key, then this is referred to as content based navigation (CBN).

Many papers have described image and video databases which use a range of image properties for content based retrieval[?, ?, ?]. Typically they are for use in specific application domains. Two examples from many possible application areas are the approach developed by Barros et al[?] which involves the use of spectral distribution information for content based retrieval of satellite images and the VIMSYS approach[?] which allows the specification of domain dependent image features and the development of semantic models based on the features and their relationships. These in turn may be used to specify the requirements for a retrieval. Some systems provide content based image retrieval tools of wider applicability, rather than being tuned to a specific application domain; for example, the Query By Image Content, QBIC, system developed by IBM[?, ?] which uses combinations of colour, shape and texture for image retrieval and the Manchester Content Addressable Image Database[?].

Although these and other papers address the problem of image and, in some cases, video retrieval by content, either with or without semantic considerations, few papers have explicitly addressed the problem of content

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

Hypertext '96, Washington DC USA

© 1996 ACM 0-89791-778-2/96/03..\$3.50

based retrieval within a hypermedia environment and fewer have addressed the general issue of content based *navigation* through hypermedia information. Notable exceptions include the Miyabi system[?], under development at NEC, which is a hypermedia system in which non-textual information may be selected using queries presented in the same form as the medium to be retrieved. However, it appears that the system is designed to retrieve similar items to the query, that is to offer CBR in a hypermedia environment, rather than to offer a general media based link following mechanism (CBN).

In this paper we describe recent developments to the Microcosm open architecture hypermedia system which provide content based retrieval and navigation facilities for non-text media. This has been achieved by extending the generic link and dynamic link facilities in Microcosm to work with non-text as well as text based information. The new system is modular and extensible so that, as more powerful techniques for the media based processes are developed, the system may be enhanced.

In the next section the Microcosm hypermedia system is summarised and then we examine the aims of the current developments. This is followed by a discussion of the extensions to the Microcosm architecture to create MAVIS, Microcosm Architecture for Video, Image and Sound. A description and demonstration of the use of the system for navigating from image, video and sound is presented and in the conclusion we consider some planned extensions to the system to provide enhanced navigational facilities.

MICROCOSM

Microcosm[?, ?, ?, ?] is an open architecture hypermedia system which allows users to browse and query large multimedia collections. It has been under development for the last five years by the Multimedia Research Group at the University of Southampton. Versions of the system are being used widely in the United Kingdom and elsewhere for a range of industrial, commercial and educational hypermedia applications. Its particular merits may be summarised as follows:

- The ability to integrate information produced by almost any third-party applications.
- All documents remain in the native format of the application which created them. No embedded markup is used.
- Link information is held in separate link databases (linkbases) which hold details of the source anchor, the destination anchor and other attributes such as link descriptions.
- Links may be made to or from any medium including those held on read-only media such as CD-ROM.
- Through the use of filters, Microcosm allows a number of different actions to be taken on items selected, so that much more functionality is available than simply clicking buttons to follow links. In Microcosm, the user makes a selection of the item of interest, then chooses from a list of possible actions.

- In certain circumstances, source anchors may be generalised to offer a spectrum of link possibilities from explicitly defined point to point (specific) links to dynamically generated anchors. The three main link types are:

Specific links. The user is able to follow the link only after selecting the anchor at a specific location in the current document.

Generic Links (providing tools for CBN). The user is able to follow the link after selecting the given anchor at *any* point in *any* document. Generic links are of considerable benefit in that a new document may be created and immediately have access to all the generic links that have been defined previously.

Dynamic Links e.g. Computed links (providing tools for CBR). Not all links need to be created by an author. For example, an inverted index of words in text documents may be used for following links, automatically made available from the index. Relationships between pieces of information may also be computed, based on statistical analysis of the content.

The original Microcosm architecture consisted of three layers: the storage layer containing the linkbases, indexes and documents in native format, a link service layer consisting of a range of filters, a filter manager and a document control system and an application layer consisting of a series of *viewers* which might typically be third party packages such as Microsoft Write, a bitmap viewer or for example, the Autocad package.

The link information held in the linkbase includes details of the source and destination anchors for the link. For specific (point to point) links the important information to represent the source anchor is the location (document and position) of the anchor. The content or boundary of the selection is also stored to allow specific links to be highlighted, for example as buttons. To follow a specific link when a selection has been made, the location of the selection must match with the location in the source anchor(s) of links in the linkbase to identify links which may be followed.

Content: the key to generic links

By contrast, when a generic link is authored, only the content of the source selection is stored in the source anchor. The original location of the source anchor is not retained. At the link following stage, only the content (and not the location) of the user selection is matched with the content section of generic link source anchors and the user may follow links wherever a match is found. Thus generic links may be followed from any place in any document where the selection matches the content of the source anchor of the link. Dynamic links also make use of content either from the user selection or both user selection and source anchor.

The versatile generic and dynamic link facilities in Microcosm greatly reduce the authoring load associated with the creation of large and complex hypermedia applications. They provide a mechanism for replacing multiple links to the same destination with a single generic

link which also reduces the problem of maintaining the integrity of the application. However, in the versions of Microcosm prior to the current project, generic and dynamic links were only available from text because the link following mechanism relied on text string matching between user selection content and source anchor content. The nearest we came to implementing generic links in pictures or video was to give a selected area a text tag. The content of this text tag, rather than the content of the picture was used as the basis for following generic links.

Specific links could be made from non-text media and in this case the content section of the source anchor was used to keep details such as the bounding box of the selection site in an image or video frame. This information was used to generate a button (i.e. highlight a specific link) in the document. The absence of content details in the source anchor of the link prevented the use of generic and dynamic links from non-text information in Microcosm.

THE AIM OF THE MAVIS PROJECT

The aim of the MAVIS project is to extend the Microcosm architecture to facilitate the use of generic and dynamic links from all media types, not just text. This is not simply a question of making the selection content available at the authoring and the link following stages because the problem of matching is substantially more complex than for text. For example, we would like to be able to author a generic link from an object in an image to another point in another document, perhaps some text about the object. If it is a versatile generic link, we should be able to follow the link from any instance of the object in any image in any document. Even if the object is as simple as a cube, it may appear quite differently in different images, either because the camera position or the lighting has changed.

We need to extract media based representations of the selection which are as invariant as possible. Ultimately we need to recognise the object using prior knowledge of how it appears in the image. We may then recognise it anywhere in other images to enable us to follow the generic link effectively. With the current state of image understanding, such an approach would have to be application specific. However, there are some general features which can form the basis for useful generic link following mechanisms for images, video and sound. For example, we could use colour distribution, texture, outline shape or combinations of these, for generic links in images and video and the pattern of the sound wave for speech.

The matching between the query selection and the link anchor must be fuzzy since exact matching is clearly impossible in non-text applications. Measures of similarity are required to assess the quality of a match and to decide if a link should be followed. It has been argued that, even for text, lexical similarity matching has problems because of the difficulty in capturing the semantics[?]. This is even more of a problem for non-text media, but

we believe that general purpose tools for CBR and CBN with non-text media can provide useful additional navigational aids.

During the development of MAVIS our aim has been to ensure that the authoring and link following processes from the users' viewpoint are not significantly different from earlier versions of the system, and to the naive user the interfaces are exactly the same. When authoring a link, for example from an image, the author makes a source selection, perhaps by dragging a rectangle across the image or sub-image and indicating the basis of the link (whether it is to be based e.g. on shape or texture or some weighted combination of features). When following a link, the selection is made in a similar way and then generic link following is requested as the appropriate action. However, in MAVIS, all of the viewing processes are now responsible for ensuring that anchors contain the content of the selected area within the document. The major changes concern the architecture of the linkbases and the selection content processing. When the selection and source anchors have been matched, the link following mechanisms are once again unchanged from the standard versions of Microcosm.

THE MAVIS ARCHITECTURE

Definition of a signature

Microcosm allows users to construct and manage links. These links are represented as relations between two or more anchors. In the majority of cases the anchors form selections from within documents. A link therefore allows the user to relate a selection in one document to a selection in another. MAVIS extends this concept by introducing the idea of a signature. A signature is a representation of the content of a selection, used in both link generation and link following.

In the case of a text anchor for a generic link in Microcosm, only one form of signature was required; a normalised representation of the original text anchor content. MAVIS, by contrast, allows many forms of signature to co-exist. This is because media such as images can be matched in many different ways. For example, in one application it may be the colour of the selection which the user wishes to form the basis of the matching. On another occasion, it may be the shape or the texture, some special feature or a combination of signatures. Thus, MAVIS allows multiple processed representations (signatures) of the original anchor data. For example, the MAVIS system currently provides support for several colour, shape and texture based signatures for images and video. In addition, it supports links in sound using a prototype signature based on the sound frequency spectrum. Each type of signature is handled by a separate signature module in MAVIS.

The Architecture

The MAVIS architecture takes the form of a framework containing a collection of signature modules which are invoked by a Microcosm process via a media table and its associated software. To understand the architecture we will consider the following three areas of functional-

ity, the Microcosm process, the signature modules (the low-level processing functionality), and the media table sub-system (the glue between the signature modules).

The architecture is illustrated in figure ??.

The Microcosm Process The MAVIS framework is implemented beneath the Microcosm architecture. The functionality of the MAVIS framework is determined by the Microcosm process which invokes it. For example, if the process is a Microcosm linkbase (for CBN) then the MAVIS framework is being used to generate and manipulate links. If it is the Microcosm content-based retrieval (CBR) process, then the framework manipulates computer generated relationships between anchors (an anchor can represent an entire document). This host process acts as a buffer between Microcosm and the MAVIS framework. It uses the media table (MT) to call the appropriate signature modules.

The Signature Modules Each signature module is responsible for all the processing for one particular signature type in one particular media type. Thus, for example, one signature module is responsible for all the processing associated with a colour signature for images which is based on the RGB colour histogram. The signature module is the only part of MAVIS which needs to be aware of the format of the signature for which it is responsible. The process using MAVIS is able to operate on anchors for any media types supported by its available signature modules.

Each signature module contains functions to perform the following:

- extract the signature from a selection in the media type for which it is designed
- store the signature in the signature files and update the signature indexes
- retrieve signatures using the index
- match two signatures, generate a measure of similarity and order signature matches

Note that the signatures, rather than the raw selection content, are stored and processed at the *link authoring stage*, in order to reduce substantially the amount of processing required at the link following stage. Because of the variety of structures which may be used to represent different signatures, each signature module maintains its own indexing system appropriate to that particular signature.

The processing algorithm implemented within a signature module may require access to parameters which the user is able to change. This is facilitated by providing each module with the ability to store its own private data and to generate its own user interface (allowing the user to examine and change the private data).

The main signature modules implemented so far are summarised as follows:

- A colour histogram module which uses an RGB colour histogram as the signature. This may be used for generic

links from images and video frames. The signature itself is a 768 bin histogram representing 256 colours in each of the red, green and blue colour planes. A 12 bin histogram, obtained by averaging across groups of 64 bins, is used as the initial index term and a nearest neighbour matching scheme is used on the index to find the initial group of indexed signatures which are similar to the query selection. The full signatures are then used to find a more accurate estimate of the similarity relative to the maximum possible difference between signatures. The matches are presented in descending order of similarity.

- A shape signature module which uses rotation, scale and translation invariant moments. Again this may be used for image and video. The image selection is converted to grey scale and thresholded according to user specified settings of the module's private data before the moments are calculated. Three low order moments are used for the initial index and once again the nearest neighbour technique is used for finding matches.
- A texture signature module based on texture statistics. Also for use with both image and video. This module currently uses the statistical geometrical features developed by Chen et al[?]. The image selection is treated as a grey scale image for extraction of the signature and nearest neighbour techniques are used for the matching.
- A Fourier signature module which uses the Fourier transform (FT) of the selection from a digital sound recording. The module performs an FT on the raw sound data to produce a 500 bin frequency histogram covering the range 1 hz to 15000 hz. This is stored and processed in a similar manner to the colour histogram module described earlier. A 50 bin histogram, obtained by averaging across groups of 10 bins in the signature, is used as the initial indexing term.

The Framework (Media Table) The signature modules are controlled via the Media Table (MT) and its associated control software. This MT subsystem is responsible for invoking required signature modules and for loading and unloading them as necessary. Modules are only loaded when instructed to by the calling process, a mechanism which provides a fully scaleable solution to the problem of limited system overheads.

The MT sub-system also provides an interface between the user and the available modules. It allows the user to select which modules should be active or passive at any point and to specify or change the private data for a module. It enables the user to specify either individual signatures or weighted combinations of signatures as the basis for CBN or CBR. The MT ensures that only modules capable of handling a particular media type are called when an anchor of this type is processed.

The relations between the processing algorithms and the media types are embodied within the signature modules. The MT sub-system maintains a list of currently available signature modules within the MT file. The MT interprets this file each time it is executed, providing the user with the ability to dynamically change the current

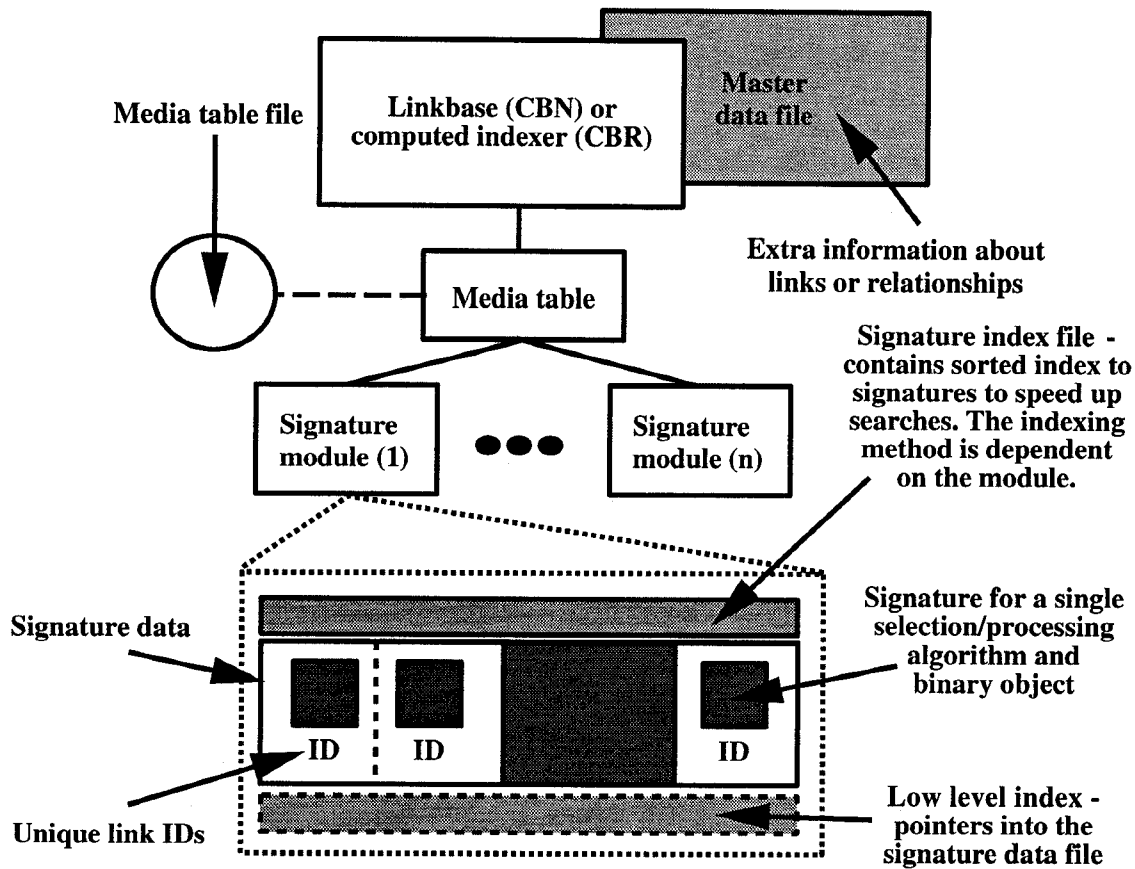


Figure 1: Module Framework for CBR and CBN

CBR and CBN functionality.

The dynamic nature of the MT allows new signature modules and, if necessary, new media types to be added at any time. If a particular set of objects require a specific matching algorithm, then a suitable signature module can be developed and plugged in (although it should be stressed that developing effective signature modules is a crucial and non-trivial task). Such a mechanism makes provision for signature modules that contain, for example, domain specific knowledge (geared towards a specific matching process).

Each module can have more than one entry within the MT file. This allows multiple instances of the same module to be concurrently available with different settings. For example a module, which thresholds grey-scale images producing a signature of the result, can be available with different settings for the threshold value, one with a low threshold value allowing very dark objects to be distinguished from a black background and a second with high threshold to distinguish only very light objects.

Module implementation

Microcosm allows users to partition their data into separate *applications*. The MAVIS architecture adheres to this paradigm allowing users to set up collections of modules to be used for each application. This allows ap-

plication specific processing algorithms to be integrated in a seamless manner. For example, if an application has a large collection of images that require a specific processing algorithm to extract image features then this can be encapsulated in a single signature module. The ability to set specific modules to operate for specific applications has a second side effect. This is that the signature indices are also partitioned, only the indices for signatures relating to the current application are searched or extended.

The process of writing a signature module is a complex low-level task. This complexity has been considerably reduced in the following ways:

1. A library containing all code common across every signature module is provided. The module author links to this library to gain access to the shared functionality. In terms of code this library accounts for approximately 60% of each module. The remaining code is module specific, usually dealing with the dedicated processing of the signatures themselves (e.g. signature comparison).
2. A skeleton source code file containing stubs for all module specific code, complete with examples for each stub is provided. The author is required to customise the contents of the stubs to allow the module to handle their particular signatures and index formats.

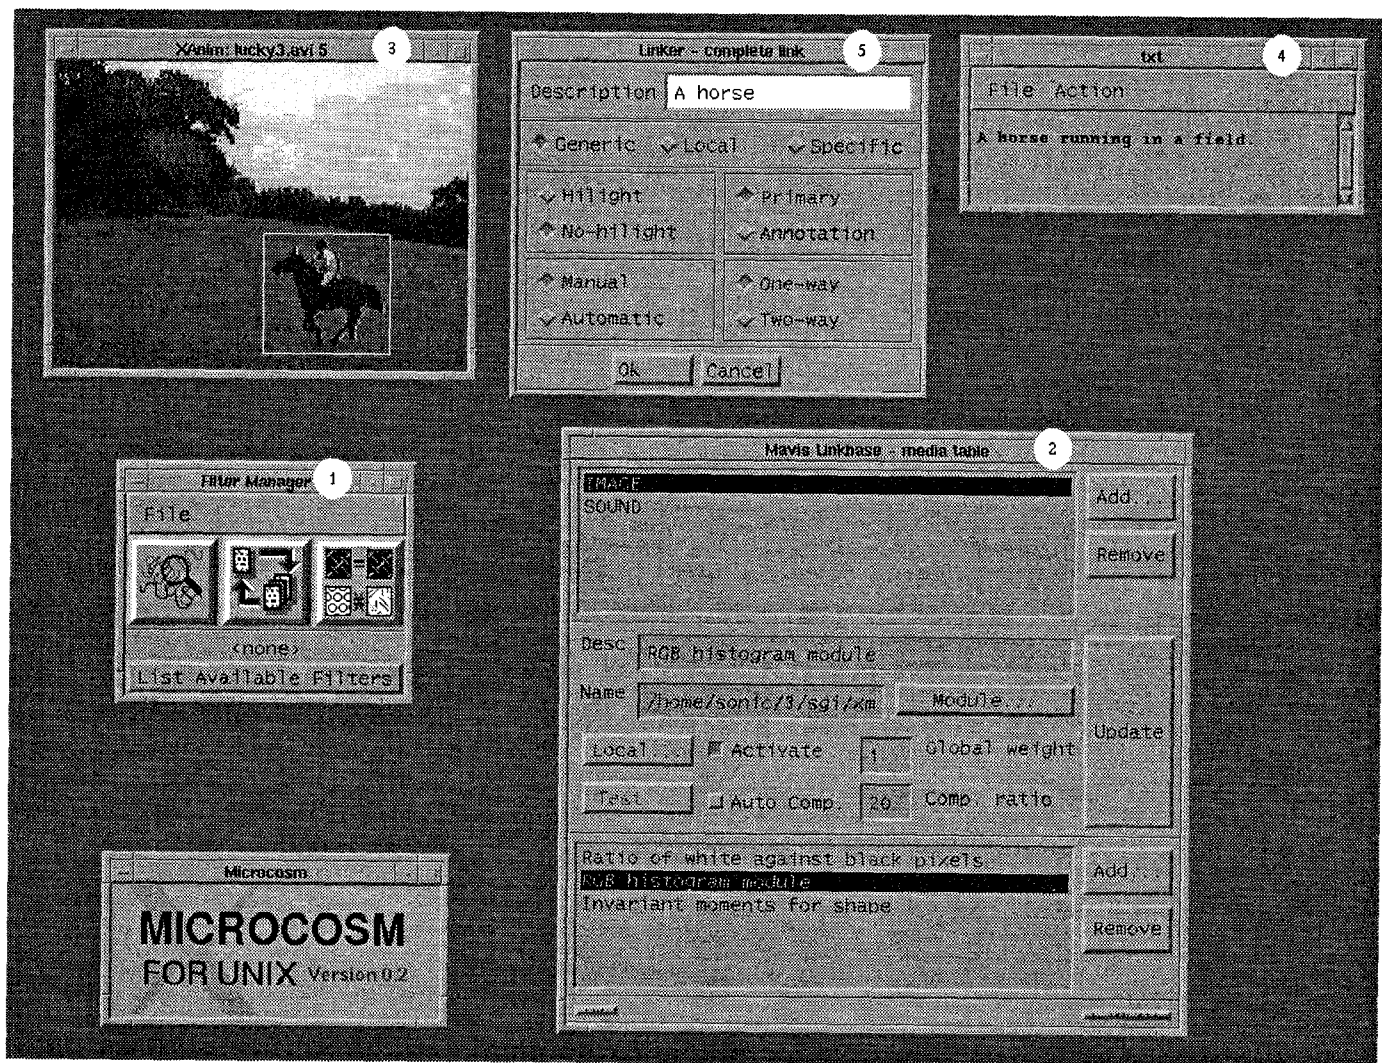


Figure 2: The Link Authoring Process with Images

3. A fully documented API is provided for the construction of new modules.

The author can choose to use the skeleton code and common functionality library or completely re-implement a signature module from the beginning. Such a process would be used if the author was integrating a third-party process into the architecture as a signature module. The new module would then be acting as a shim between the MAVIS system and the third-party application.

The MAVIS Architecture In Use

The MAVIS architecture is used for two major activities. Firstly, the signature indices can be extended by adding new anchors (the authoring process). Secondly, the signature indices can be searched to retrieve a set of the closest matches to a given link anchor (link following). These two scenarios are described and illustrated in the following sections.

Authoring Links When an anchor is created, by the user creating a link (for CBN) or a Microcosm process automatically indexing data (for CBR) it has to be added to the currently active signature indices. Each signature module is responsible for extending its own index. The Microcosm process that is using the MAVIS framework calls each of the currently active signature modules through the media table (MT). Each signature module is passed the original selection data and processes it to form a module specific signature. It then indexes this and stores both the index term and signature in its own indices. When this process has finished, each module that is related to the current selection media type will have recorded a signature and its associated index term. Any module that does not handle the current media type or has been deactivated by the user will not be called by the Microcosm process.

From the user's point of view the link authoring process is similar to conventional link authoring. In figure ?? a link is being authored from an image to a text doc-



Figure 3: The Link Following Process from Video

ument. It involves selecting the Microcosm linker and the MAVIS Linkbase (1 on the figure) and selecting the required modules via the media table (2). In this example, a colour histogram module has been selected. The link is made by highlighting an area in the image as the source anchor selection (3) and a destination anchor is selected in the text document (4). The End Link action is chosen to activate the complete link pop-up (5) which enables the type of link to be set, together with a description. Choosing *ok* then adds the link to the linkbase to complete the authoring process.

Link Following Signature module indices are searched when the user attempts to follow a link from a selection (CBN) or requires documents containing a similarity match with the selection (CBR). The Microcosm process using the MAVIS framework calls each signature module that is active and is related to the media type of the current search anchor. Each signature module is given the data from the original search anchor. This is processed and used as a key for a search of the modules

indices. The module stores a signature for each anchor indexed by the system and maintains a sorted index for these signatures, allowing it to speed up searches.

Each active module processes the search anchor data, producing a signature. This is then processed to form an index term. This index term is used by the module to locate the nearest n matching signatures (stored in its indices). The module returns a list of the nearest matching signatures to the Microcosm process. It then combines these together to form one overall list of the closest matching signatures for the given search anchor and active signature modules.

An example of generic link following is shown in figure ???. Here, a digital video sequence has been stopped on a particular frame and a selection made in the frame in an attempt to find and follow any generic links. The MAVIS image clips filter, the MAVIS linkbase and the Microcosm Image Links filter (1 on the figure) are required. The image clip pop-up (2) allows the user to select an area from the screen and apply a Microcosm

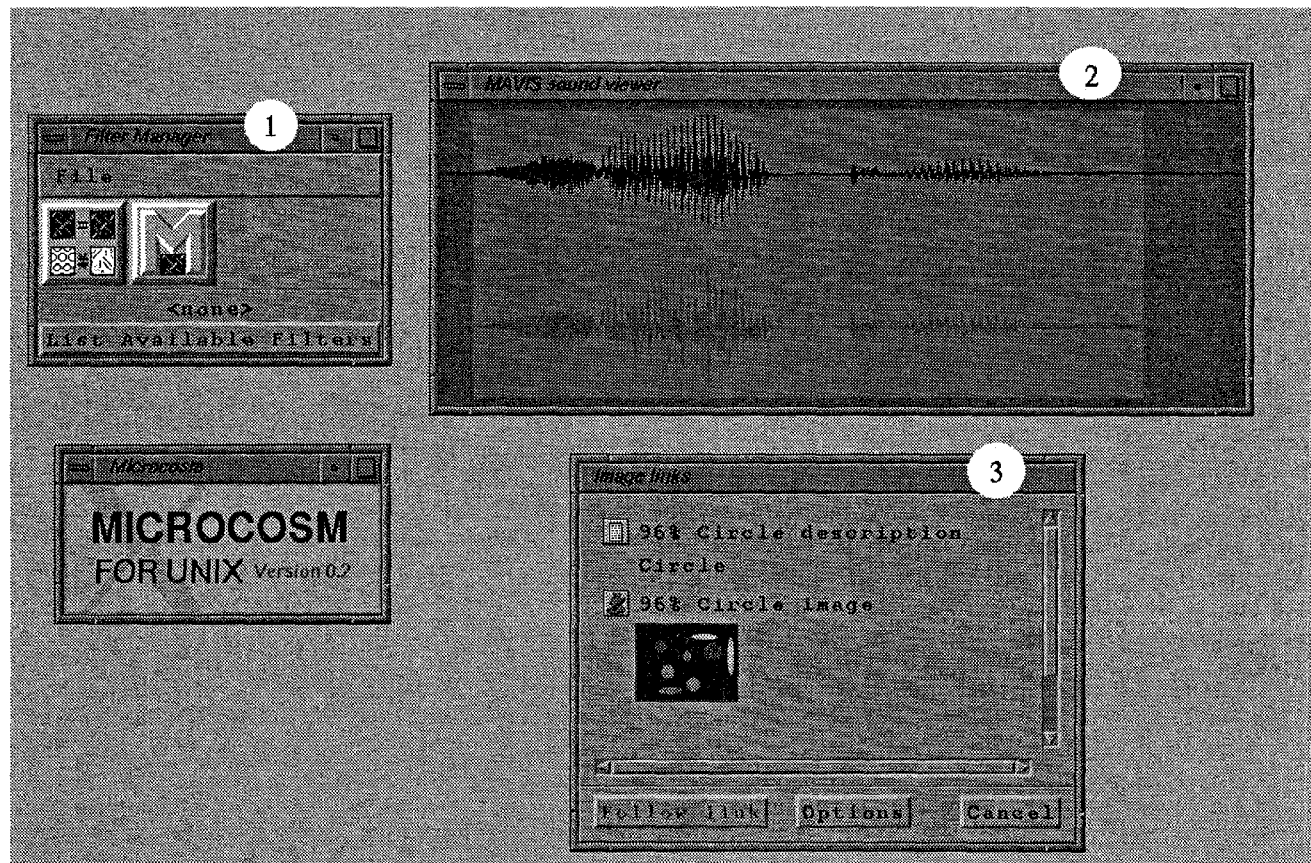


Figure 4: The Link Following Process from Digital Sound

action to that selection. In this case we wish to try and follow a link on the selection. A third party movie player (3) is showing the video and has been paused at a frame where we wish to search for links. By selecting a region within the movie display, the linkbase is searched and a link matching the content of the selection is found and shown by the Image Links filter(4).

Link Processing in Sound The sound (FT) module has been used as a demonstrator for the audio media type. Single words have been articulated and selected at the authoring stage as the source anchors of generic links, by dragging a rectangle over the raw digitised sound waves. In figure ?? the word *circle* has been articulated again. The MAVIS linkbase and Microcosm Image links have been selected (1 in the figure). A sound viewer (2) contains the digitised sound waves for the word *circle* recorded in stereo. Highlighting the appropriate selection and performing a follow link results in two links shown by the Image links filter (3). One link is to a text description and one is to an image of circles.

PERFORMANCE ISSUES

The time required to author or follow a link depends on both the time required to make the selection and the time required to process it. In the examples given above, using the colour histogram module, the selection process consists of clicking a few menu items, dragging

a rectangle over the required region and clicking on the action required. It takes no longer than making a similar selection in text. However, to make a selection for shape matching may sometimes involve the user in specifying a non-rectangular selection or invoking a segmentation algorithm to delineate the desired object. This involves more user operations and we are currently working on fast interactive tools to assist with selection delineation in these circumstances.

The time for processing the selection is mainly determined by the indexing and matching strategy used in the individual signature modules. Initially we have used k-nearest neighbour matching which, for the maximum sizes of link base used so far (about 400 entries), provides a very fast response on the Silicon Graphics Indy platforms we are using. However, this approach is $O(n)$ where n is the number of links and we are now implementing tree based indexes which will scale better to very large link collections.

CONCLUSIONS AND FUTURE WORK

In this paper we have described extensions to Microcosm to facilitate content based retrieval and content based navigation by extending the generic link to work with any media types, not just text. In contrast to closed systems, the Microcosm/MAVIS system is an open system providing a link service for documents of any type,

authored using a wide range of end-user applications software.

The MAVIS extension relies on the development of media specific signature modules which use selection content as the basis for signature generation. The signatures are used for indexing and matching source anchor selections and user query selections. Examples of experimental hypermedia systems which we are developing using the MAVIS software include a product catalogue for a large retailer which contains product images and text, and an archaeological information system containing large numbers of artefact images and associated documentation.

Current modules use statistical pattern matching techniques for signature matching to provide basic general purpose navigational tools. In addition to the performance related developments mentioned earlier, we plan to extend the range of signature modules available and draw on recent research in image, video and speech recognition to provide more powerful and robust content based retrieval and navigation. Some future modules will require prior knowledge of expected objects to provide more versatile but applications specific navigation.

ACKNOWLEDGEMENTS

This research has been undertaken with the support of SERC grant number GR/J53614. Thanks also to Nick Beitner for the equestrian video and to Mark Dobie for technical support.

REFERENCES

1. BARROS, J., FRENCH, J., AND MARTIN, W. System for indexing multi-spectral satellite images for efficient content based retrieval. In *SPIE Proceedings, Storage and Retrieval for Image and Video Databases III* (1995), pp. 228-237.
2. BEITNER, N. B., GOBLE, C. A., AND HALL, W. Putting the media into hypermedia. In *SPIE Proceedings, Multimedia Computing and Networking* (1995), pp. 12-23.
3. BERNSTEIN, M. An apprentice that discovers hypertext links. In *Hypertext: Concepts, Systems and Applications* (France, November 1990), A. Rizk, N. Streitz, and J. Andre, Eds., pp. 212-223.
4. CHEN, Y., NIXON, M., AND THOMAS, D. Texture classification using statistical geometrical features. *Pattern Recognition* 28, 4 (1995), 537-552.
5. DAVIS, H., HALL, W., HEATH, I., HILL, G., AND WILKINS, R. Towards an integrated information environment with open hypermedia systems. In *Fourth ACM Conference on Hypertext (ECHO '92)* (Milan, Italy, December 1992), D. Lucarella, J. Nard, N. M., and P. Paolini, Eds., pp. 181-190.
6. FALOUTSOS, C., BARBER, R., FLICKNER, M., HAFNER, J., NIBLACK, W., PETKOVIC, D., AND EQUITZ, W. Efficient and effective querying by image content. *Journal of Intelligent Information Systems* 3 (1994), 231-262.
7. FOUNTAIN, A., HALL, W., HEATH, I., AND DAVIS, H. Microcosm: an open model with dynamic linking. In *Hypertext: Concepts, Systems and Applications* (France, November 1990), A. Rizk, N. Streitz, and J. Andre, Eds., pp. 298-311.
8. GUPTA, A., WEYMOUTH, T., AND JAIN, R. Semantic queries with pictures: The VIMSYS model. In *Proc. 17th International Conf. on Very Large Data Bases* (Barcelona, September 1991), pp. 69-79.
9. HILL, G. J., WILKINS, R. J., AND HALL, W. Open and reconfigurable hypermedia systems: A filter-based model. *Hypermedia* 5, 2 (1993), 103-118.
10. HIRATA, K., HARA, Y., SHIBATA, N., AND HIRABAYASHI, F. Media-based navigation for hypermedia systems. In *Hypertext '93* (Seattle, Washington, USA, November 14-18 1993), ACM Press, pp. 159-173.
11. JAIN, R., MURTHY, S., CHEN, P., AND CHATTERJEE, S. Similarity measures for image databases. In *SPIE Proceedings, Storage and Retrieval for Image and Video Databases III* (1995), pp. 58-65.
12. OAKLEY, J., SHANN, R., DAVIS, D., AND HUGUEVILLE, L. A database management system for vision applications. In *5th British Machine Vision Conference* (York, U.K., 13-16 September 1994), BMVA Press, pp. 629-639.
13. RAMESH, N., AND SETHI, I. Feature identification as an aid to content based image retrieval. In *SPIE Proceedings, Storage and Retrieval for Image and Video Databases III* (1995), pp. 2-11.
14. SMOLIAR, S., AND ZHANG, H. Content-based video indexing and retrieval. *IEEE Multimedia* 1, 2 (1994), 62-72.
15. TREAT, H., ORT, E., VO, M., JANG, J.-S., HALL, L., TUNG, F., AND PETKOVIC, D. Searching images using Ultimedia manager. In *SPIE Proceedings, Storage and Retrieval for Image and Video Databases III* (1995), pp. 204-213.