Towards A New Authoring Methodology for Large-Scale Hypermedia Applications

Ian Heath, Wendy Hall*, Richard Crowder#, Gary Wills# , Jim Ballantyne%*

*Department of Electronics and Computer Science
#Department of Electrical Engineering
University of Southampton
SO17 1BJ, UK
Tel: +44 (0)23 8059 3255
%Post Office Research Group
Concept 2000
250 Farnborough Road
Farnborough
GU14 7LU, UK
E-mail: ih@ecs.soton.ac.uk

Abstract

As the amount of information technology increases, managing information resources, so that the correct people can find the information easily, becomes a critical issue. Hypermedia systems are considered one solution to this problem as they provide a means for representing higher level relationships between the underlying information. However, the amount of information available electronically is increasing at an accelerated rate. Using standard hypermedia authoring techniques, the effort required managing and maintaining large-scale hypermedia systems is enormous. Hypermedia authoring in the large requires new methodologies if it is going to be feasible.

This paper presents a new model for building and structuring large-scale hypermedia applications. It describes a case study that explored the delivery of hypermedia information in an industrial environment on a small scale. Models and techniques developed for that case study were then refined and augmented so they could support the construction of large-scale hypermedia systems. In order to support such endeavors a new linking model is presented that allows the author to explicitly represent abstract concepts contained within the underlying information and interconnect them in some meaningful manner. An example usage of this linking technique is then presented.

KEYWORDS: Industrial strength hypermedia, generic links, link clusters, authoring methodology.

Introduction

The provision of information on a large scale brings with it a whole new set of problems when considering how such information systems are to be constructed. Current hypermedia structures are mainly concerned with the representation of relationships between fixed locations within documents. This approach works well if the relationships are structural in nature (for example, relating entries in a table of contents to the documents containing the relevant text) but becomes cumbersome if the things being represented are more abstract.

However, new authoring methodologies are being developed that provide alternative strategies to explicit link structures. Microcosm [9] introduced the generic link that provides a context free link available from any occurrence of the source anchor, providing a many-to-one link. While this does indeed reduce the amount of authoring effort required, it still effectively uses an explicit link model.

This paper explores a simple model for representing document clustering using the Microcosm system and implemented for an industrial application. By examining the original case study, developed for Pirelli Cables, Aberdare, experience gained using this approach is then used to extend the model to provide a richer environment for representing these relationships. This model has been developed as a part of the Factory Information Resource Management (FIRM) project also in co-operation with Pirelli Cables. We then describe how these ideas have been extended and applied in another case study using information held by the Post Office Research Group (PORG) and how they might be generalised to provide a new authoring methodology for industrial hypermedia applications.

The Aberdare Case Study

Our initial case study involved the design and implementation of a hypermedia system for industry [4][5] using a cable packaging line at Pirelli Cables, Aberdare, Wales as the target environment. All aspects of the work required to produce such an application were being considered, from capturing resource information, through the structuring of it and finally the delivery of the information within a factory environment. The application developed was targeted in particular at maintenance and operator personnel. The Microcosm [9][11] hypermedia system was chosen to hold the information set as it provided a unique environment for the development of the tools and techniques required to support such an undertaking.

Capturing the electronic information was a large undertaking as all the information was only available in paper form. Once captured, the documents were stored as simple text files or raster images, with no structure or markup.

A main objective of the project was to interconnect all of the relevant information about an item, which might be spread over the documentation. The user could click on an electrical part in a circuit diagram and access all the other diagrams that contained that part, along with other information such as location diagrams, parts lists, etc. Such a collection of information was termed a *cluster*.

Microcosm generic links [9] were used to provide the clusters of information for any item as desired and proved to be very effective. They were easily implemented for textual information. To achieve the same effect in the drawings, extensions to the hypermedia system had to be introduced that allowed the association of textual names with areas on an image. An author would then name each mechanical or electrical part on each diagram using the name chosen to represent the cluster. The user was then able to click on a part in a drawing and access the same resources as were available from the text documents.

To aid in the development of these clusters, tools were developed that could automatically parse several classes of documents and derive the required generic links. Also, the use of raster images was dropped in favor of CAD drawings. These had the advantage of containing structure that could be interrogated and used to automatically label the drawings, which helped reduce the effort required generating the clusters. Figure 1 shows an example of this. The user has selected the object named "Y145" from the electrical drawing and Microcosm has returned a list of destination documents that can be reached from that term, by searching the available link databases.

The case study was well received when trials were conducted at the Aberdare site. The acceptance of the system and the approaches used to structure the information led to a second project, FIRM, which looked at similar problems with information in industry but on a much larger scale.

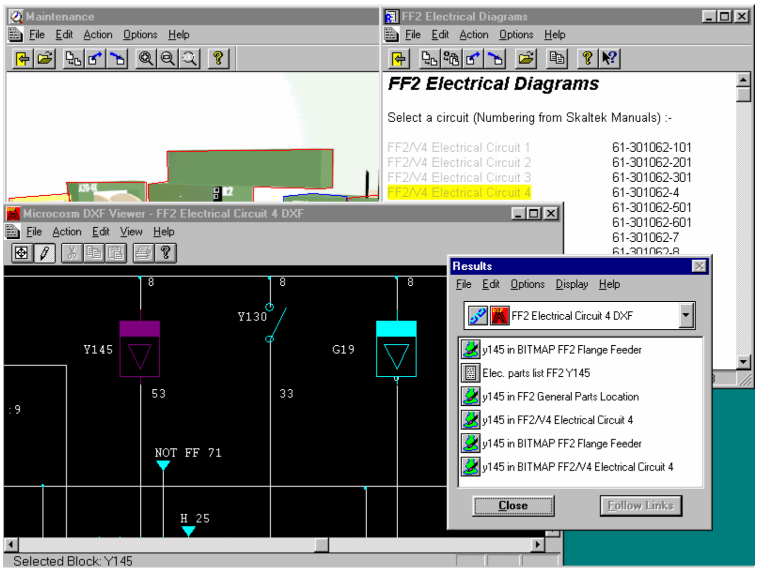

Figure 1: The Generic Link Clusters within the Aberdare Case Study

The FIRM Project

A research project, FIRM [6][7], is being undertaken in order to examine the problems of dealing with capturing, structuring and delivering information on a factory or organisation wide basis. When considering the provision of information on this large a scale, many new problems become apparent. The amount of information that could possibly be accessed becomes vast. There were a large number of information sources that had to be brought together, covering not only specific equipment but also general practice and company policy.

The Aberdare case study outlined above highlighted some of the problems with constructing an industrial hypermedia system but was only concerned with one particular piece of plant equipment. The information was manageable, even though it originally existed in paper form. It was also easily controlled, as the hypermedia application was very self-contained as it dealt only with information generated by the author.

The requirement for industrial strength hypermedia is well known [18]. If hypermedia is to be used in such an environment, then a great deal of work is required to integrate the resulting system with existing business processes. Many organisations will already have methodologies in place to control the flow and quality of information (e.g. ISO 9000). These processes will have to be embraced by the hypermedia system if it is to be adopted. This means any proposed model must be simple to maintain and implement whilst providing a real benefit for the organisation as a whole.

Controlled access to the information becomes of paramount importance. The original case study simply dealt with operators and maintainers of one machine. When considering the use of the system factory-wide, a large range of users with vastly differing computing skills have to be supported. Also, the range of information available will be greater, coming from all aspects of the company. Filtering of the information based on various attributes of the user (for example, their job function or specific requirements) is essential to stop the user from being overwhelmed with information.

The FIRM Authoring Model

In order to make the development of very large-scale hypermedia applications possible, there has to be a move away from monolithic information sets that try and contain all of the required information. A prime example of this is the information set finally produced for the original Aberdare case study. The final statistics involved approximately 150Mb of data and 7,000 links connecting some 600 documents. Although some of this was automatically generated it was becoming unmanageable.

Also, when considering the requirements for delivering and maintaining information on a factory wide scale, the ability to reuse the same information wherever it is required becomes critical. For a production line there may be many cases where the same component (for example, a drive) is used in many different locations, forming part of a larger assembly. If this information is to be accessible electronically, there must be methods for sharing the information between the different assemblies (or *instances*). To help achieve these aims, a new authoring methodology is required.

Automatic Structural Linking

One of the conclusions of the original Aberdare case study was that a great deal of work was required in order to manually link the relevant information. The use of generic links where possible reduced this effort but the building of large-scale hypermedia systems was still considered a large undertaking. Therefore, research was undertaken into the development of tools to automatically or semi-automatically derive structure from the information itself.

There have been many approaches to this task. The information might be marked up in some structured format, such as SGML [15] and then links derived from that structure using tools such as Dynatext [10]. EPM [16] describes a more detailed study of the development of links in a large-scale application based on structured information. However, work on the original case study lead to the desire to use electronic versions of documentation as source material wherever possible. Ideally, the originator of the particular piece of equipment would supply this information wherever possible. Thus, the FIRM system must support the use and structuring of third-party information supplied in a variety of formats. Any tools must deal with these formats and so it was decided to develop a simple, ad-hoc, solution to generate simple, structural links. Most of the links created would be simple navigational shortcuts, providing access to information as described below.

- The table of contents (or index) to the relevant section.
- A list of tables (or figures) to the relevant table (or figure) and their reference in the text.
- Explicit references to other procedures, data sheets, etc., be it on the local system, or on an external network such as the WWW.

The Microcosm system supports the creation of links between documents without modifying the underlying content, and so is able to more readily use the information as supplied, without conversion. To automatically generate such structural links, tools were developed to process the structure of the documents. As the documents were generated electronically, the structure was derived from the word processor markup. Due to the fact that the documentation supplied was designed for display on the page and not the screen, some conversion work was undertaken to make the resulting information more readable.

First, the manual was dissected into small nodes by dividing it into separate sections. Tables of contents were automatically generated and linked to these nodes. Due to the programmable nature of the word-processing environments (such as Microsoft Word), it was possible to develop the required tools using the in-built macro languages, providing complete access to the underlying representation of the document and control over the format and structure of the generated hypermedia nodes.

Automatic generation of links is made easier by the use of templates/procedures and guidelines for construction of the documents. In practice, several templates/procedures and guidelines are required i.e. for manuals, memos, reports, specifications, etc. In many companies these will already exist, either as part of the company's quality procedures or just good practice.

Modular Hypermedia Applications

A standard principle of software engineering suggests that a large programming problem can be more easily managed by decomposing it into smaller modules that are more easily dealt with. This approach can be applied to the information domain, so that large information systems are broken down into smaller hypermedia applications, or Modular-Hypermedia Applications (MHAs), each of which represents a fraction of the whole. For example, one MHA could be created describing a drive unit used on various production lines whereas another describes a particular Programmable Logic Controller (PLC).

In Microcosm, an application is defined to be a collection of links documents and process that, when operated together, can be considered an information system. The MHA model extends this definition to also include a set of child MHAs that provide additional information. Each MHA is considered to be self-contained (i.e. it does not contain explicit references to documents that are not a part of the MHA). However, an MHA can use generic links to provide context free linking from anywhere within the resulting application.

If explicit links are required between MHAs then they are defined within the MHA that is the parent of the source and destination. This is because the parent MHA is considered to provide the context in which the two child MHAs are being used. An example of the hierarchical nature of MHAs is shown in Figure 2. Here the "Extruder" MHA also uses a "PLC Manual" MHA and a "Drive Manual" MHA to supplement the available information. As the "Extruder" MHA is providing the context for the use of the other two MHAs, it also holds any cross-MHA links (e.g. link "A3").

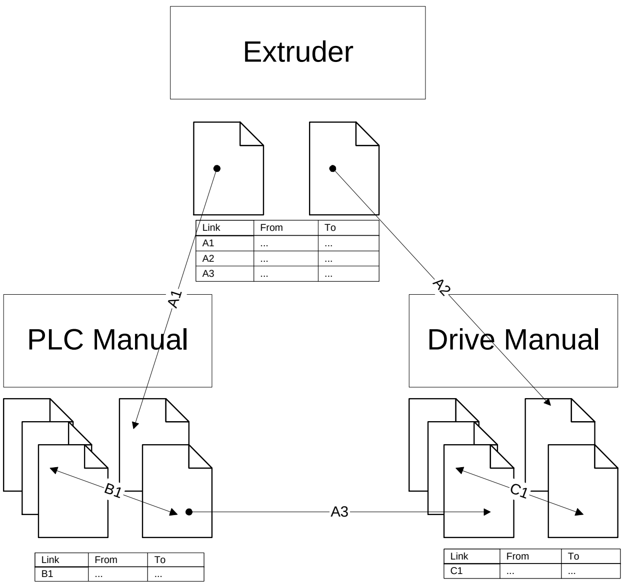

Figure 2: An Example of MHA Relationships

Information reuse becomes critical [12] as the size of hypermedia systems becomes ever larger. This approach has the benefit in that once an MHA has been developed, it can then be used in many different contexts. Rather than dealing with reuse at the document level, the author can reuse entire hypermedia collections.

A second benefit is that the modules can be managed centrally which helps increase the maintainability and quality of the final application. If the content of a module is updated, all uses of that module will also be updated automatically.

Finally, the authoring effort required for building applications becomes less, as they can be built largely from pre-existing information modules. In the future, manufacturers will be able to supply the documentation in the form of modules that can then be integrated into the corporate information base and used wherever needed.

This model is similar to that used by HM-Card [22], which defines a data model to hold the hypermedia information in units called S-collections. These S-collections are the basic operational unit in HM-Card, and can contain pages, links and other S-collections that are stored in an underlying database.

Link Clusters

The original Aberdare case study used generic links to represent higher level information about an element of the line. This had the advantage that it was easy to implement as it used features already present in the hypermedia system. Also, it was easy to extend the set of documents contained within the cluster, all the author had to do was create another generic link with the source anchor referencing the cluster's name and the destination pointing to the new material.

However, this implementation had several drawbacks. Only one term could be used as a lead-in to the cluster, as the same term was used to provide the interconnectivity. If several terms were used to describe the same idea (for example, a component on a drawing may be identified as "Y145" and on a parts list is identified by the manufacturer's part number, "62-1221-4598") then a copy of the cluster had to be created for each individual term required. The visibility of clusters was also important. All clusters were visible, regardless of who was using the information system. For example, clusters containing information about the physical components of the machine were accessible to any of the operators that used the application, which might cause confusion. Again this was due to the fact that the clusters were represented as links and, as such, had no control over their visibility. Finally, each cluster was an information island. It contained a set of documentation relevant to the particular concept it represented but that was all. In some cases it would have been useful to be able represent relationships between the various clusters (for example, "Y145", "Y148" and "Y150" are all components located within a "Flange Feeder").

All of the problems outlined above could be traced back to the original implementation chosen to represent the clusters. As clusters were represented as links and did not exist as first class objects [2] within the information system, it was difficult to provide the additional functionality required to solve these problems. Clearly, the representation and implementation of the clusters would have to be changed.

Electronic thesauri [1][17] provide an established model for representing clusters of related information and the semantic relationships between them. A thesaurus entry contains one or more *lead-in*s, which are terms used to locate entries within the thesaurus. Also, the entry has a set of *related documents*, which are defined to be somehow associated with the entry in question. This representation has the advantage that the terms used to locate a cluster are separate from the documents related to that cluster, which was not the case using the link representation, and so was chosen to be a part of the new cluster model.

There has been much work in the development of better methods for representing the interrelationships between semantically similar documents. Classification schemes and semantic querying [8] can be used to navigate the hypermedia system via queries. A Description Logic [3], such as GRAIL [13] provide languages that allow the definition of entities, and rules for the complex relationships between them.

Aquanet [20][19] was a browser-based tool that allowed users to represent information in order to explore its structure. The user could define both the underlying structure and the graphical appearance of the knowledge. Viki [21] took this spatial structuring one stage further by providing an more informal environment that supported emerging structure. The browser was able to spatially parse the knowledge structure and suggest different ways to both represent and display the information Trellis [14] used petri nets to represent the interrelationships between items of information and also provided programmable browsing semantics.

However, care must be taken when considering the use of any of these models due to practical considerations of the target application (in this case, manufacturing industry). Rather than building a system explicitly for knowledge representation, the link clusters are simply providing another view of the underlying information.

The simplest method for representing relationships was to define an explicit connection (or *link*) between the clusters. Each link was a directional relationship that was typed to allow the development of future querying tools. The description generated for traversal of the link in each direction was configurable (for example an *information* link from A→B would generate "More Detail on ’B’" when viewed from A, and "Other related information in ’A’" when viewed from B). The author was free to extend the set of available cluster link types.

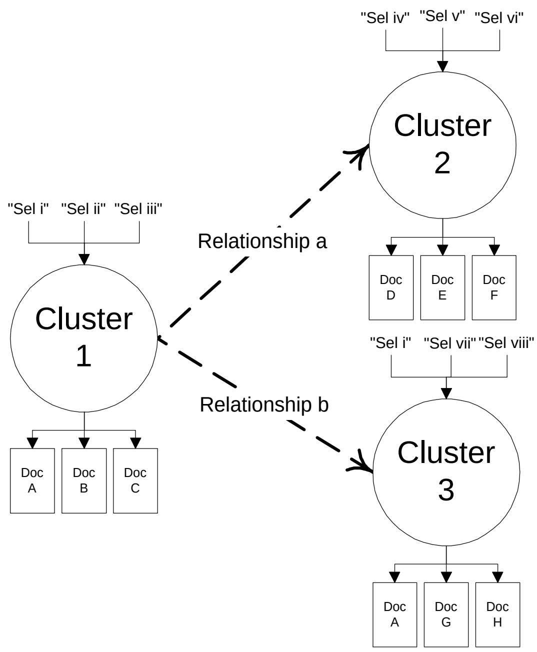

Figure 3: An example of the Link Cluster Model

The final model for link clusters used in the FIRM application is shown in Figure 3 and consists of the following:-

- A set of source anchors, termed *lead-ins*, that connect from the information to a particular link cluster (e.g. “Sel 1” as shown in Figure 3). A lead-in may be generic, i.e. not bound to a particular location within a document. A number of different clusters may use the same lead-in.

- A set of destination documents that are somehow associated with this cluster (determined by the author). Documents may be referenced from several clusters (e.g. both Cluster 1 and Cluster 3 in Figure 3 refer to Document A).

- A set of arbitrarily typed relationships between this cluster and other clusters.

A simple permission model has been developed for the clusters to provide a means for controlling their visibility. Users of the system were assigned to groups, which initially reflected the different job descriptions within the company (for example, *maintenance*, *operation*, *sales, etc*). Access control lists were attached to the clusters, allowing the author to control access by specifying an arbitrary list of permissions for users and/or groups. Thus, clusters that represented maintenance information could be tagged to be visible only if the user was a member of the *maintenance* group.

A simple relational database was used to store the cluster definitions, using tables to represent the relationships, lead-ins and related documents. For the initial implementation, this database was "compiled" to a Microcosm linkbase, as this required the minimal amount of development work to test the concepts.

Each lead-in became a generic link which, when traversed, caused another linkbase query to be executed to return the relationships and related documents. Cluster relationships were also represented as special links that would cause the destination cluster to be displayed, again by sending a linkbase query. The related documents were "normal" Microcosm links that caused the relevant information to be displayed. Thus, the user would navigate the link clusters in exactly the same way as they would navigate structural links.

Using this approach, it was possible to experiment easily with the link cluster model. However, there were problems, the biggest being that the links have to re re-compiled every time the underlying cluster database changes. For the system to be truly usable, the cluster database would have to be queried "live", and the results presented somehow. Fortunately, the Microcosm system provides an open method for extending the functionality of the information environment. Developers are able to write process, termed filters, which can extend or modify the system in some way. A cluster database filter was developed for just such a purpose. Again, the results were presented alongside the normal linking interface but this time the cluster links were being dynamically generated. Thus it was possible to modify the underlying database and have those modifications instantly available to the user.

Authoring and Visualisation of Clusters

The Post Office Research Group (PORG) is watching technologies develop to try to deduce how and when they might be relevant to the UK Post Office. Charged with researching and disseminating information about technology to the Post Office businesses it is essential to answer these questions in a relevant and timely manner.

To assist in the search for such technologies PORG have sought an approach to research that will allow the staff of PORG to be aware of such technology and to share with colleagues in an easy to understand fashion what this means for PORG and the Post Office. The results of the approach are documented in the PORG Research Plan it is in the plan that the various initiatives are formally linked with the business objectives. The production of the plan decomposes into a number of elements:

- Identify the Business Objectives of the Post Office
- An identification of technologies from a Post Office point of view.
- Identify candidate 'technologies' for future research
- Identification of candidate technologies is a selection process relating 'selected' business objectives and 'selected' technologies

To give insight into what technologies are associated with others and to give an indication of areas of high activity, PORG have adopted a simple but effective method based on 'bubble diagrams'. Despite their simplicity, bubble diagrams offer a rich means of representing technologies and how they relate to each other. Using this device it has been possible to make considerable headway with a paper based exercise to sorting and categorising the technologies of interest to PORG and the Post Office. An example of a bubble diagram is shown in Figure 4.

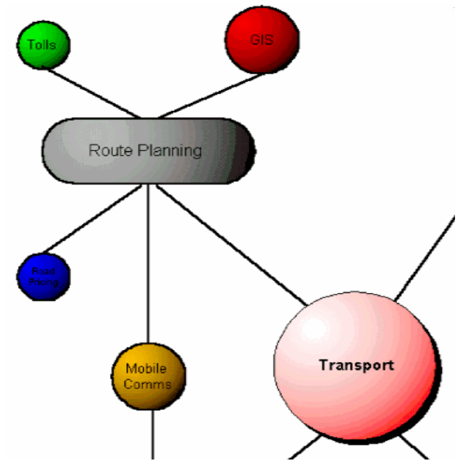

Figure 4: A sample bubble diagram derived by PORG

There was a desire to allow members of PORG to interactively view and navigate these diagrams and so a trial was undertaken to build an electronic version of them. Each bubble within a diagram maps directly onto a link cluster along with the relationships between it and other bubbles. To provide added value to the diagrams, a set of link clusters was developed to represent the individuals within the research group. By using information regarding their research interests, each member of the group could be connected to the relevant research clusters. Diagrams containing various views of the research clusters were imported into the system and connected to the actual clusters using generic lead-ins. Finally, the technical reports and information pages for the various group members were imported, and indexed to allow full text retrieval.

An example of navigating the link clusters is shown in Figure 5 and Figure 6. Here, the user has opened a diagram and followed the link from the picture of "Route Planning" to the actual cluster (shown as Cluster 42 in Figure 5). Here all of the connections can be seen, along with the related documents, which in this case simply point back to the original diagram. However, the user has decided to find out about a particular member of the research group who has registered an interest in this cluster. By choosing the appropriate connection, the cluster relating to that individual is displayed (Figure 6). The user is then able to see the homepage of that group member, and also their other interests. It is also possible to search the document database for documents that may contain their name. This is achieved by using the concatenation of the text-based lead-ins as a search term for a full text retrieval engine.

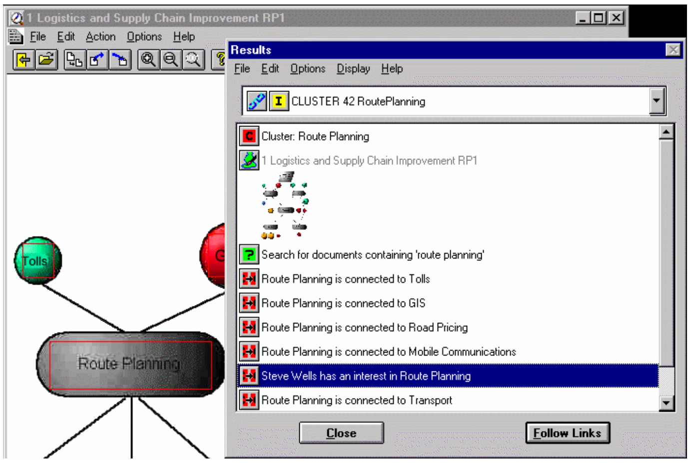

Figure 5: Accessing the Cluster Database

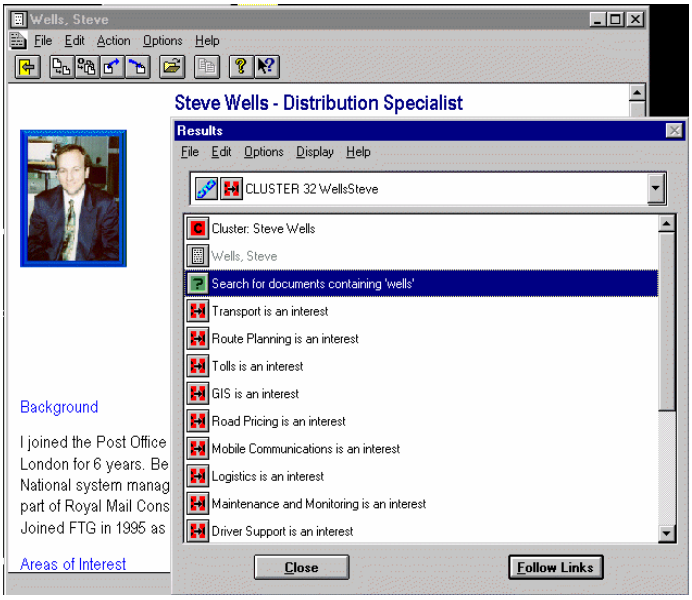

Figure 6: Navigating the Clusters

The Authoring Environment

A large amount of work has been undertaken to build an environment that supports this style of application building. Rather than make fundamental changes to the Microcosm system, a separate authoring environment is being developed.

This environment consists of a front-end that supports the development of information modules, and a database that is used to store the module and cluster structure. This information is then *compiled* into a Microcosm application that can then be used as normal. Once Microcosm has finished executing, the environment parses the generated information looking for any changes made.

This approach has other advantages, namely it is not limited to generating one single representation of the underlying data. It is possible to build other "compilers" that can generate output for other hypermedia systems. Currently the environment has support for generating versions for Microcosm and the Web (although only as a limited subset).

Towards a Methodology

The design and implementation of FIRM has been targeted at solving problems encountered in the delivery of large-scale information to an industrial user. However, it is envisaged that the techniques and tools that have been developed are applicable in many different domains. By examining the work to date, a general methodology begins to emerge.

To create a large-scale hypermedia system the MHA approach is used. First, the large information resource is hierarchically decomposed into smaller, manageable MHAs which can then be developed as separate information systems in their own right. The way in which the information is divided up is dependent on the subject matter, but in the industrial domain, a division at the physical component layer (e.g. drive, PLC etc) seems appropriate. This has several advantages. The information developed will be applicable within other similar application domains, promoting reuse. Also, the cost in developing such information systems may be prohibitively expensive for the users of such equipment, but in the future it is envisaged that the supplier of the equipment will perform this task as with current documentation provision. These MHAs can be developed in any hypermedia system the FIRM environment is able to import from. Currently, this is only the Microcosm system but it is envisaged that other import modules will be developed as the need arises and that, in time, delivery will be via the Web.

Once the MHAs have been developed, they must be integrated into the target information system. Again, in the industrial domain, this is seen to be another task that the system integrator would perform. As each MHA is a self contained information system, the integrator's responsibility is mainly to provide the personalisation information required to describe how the system integration was performed, linking to the manufacturers' information as appropriate. Again, this procedure is completely dependent on the domain, and the following description is based on an industrial application.

As an initial step, the MHAs are imported into the FIRM authoring environment. The target information resource is constructed by combining the individual MHAs together into a final hierarchy. To achieve this, MHAs that are relevant to a particular subsystem are collected together into a larger MHA that contains references to the sub-components along with any additional documentation and links required providing the "glue" (for example, the "Extruder" in Figure 2). This process is continued, collecting together larger and larger MHAs until the final resource is built. This "super" MHA is then delivered as a complete information system to the client, along with the completed production equipment. A similar process is then followed at the client's location. The MHA supplied by the integrator is integrated into the FIRM environment as a sub-component of the factory information system.

Cluster links can be used during the integration stage to provide a rich set of high-level abstract links that bridge between the MHAs. Because a cluster link can have many source anchors, it can be used to provide a coherent linking structure where different terminology is used to represent the same idea. For example, documentation supplied by the drive manufacturer will refer to the drive by its model number. However, other documentation may refer to the same drive by its location (for example, "the Extruder Drive"). A single cluster link can be created that allows both of these terms, plus any others that are appropriate, to be connected to one or more relevant destination documents.

Conclusions

The MHA model provides a means by which large-scale information systems can be more easily constructed and maintained. The FIRM information system delivered to Pirelli Cables contains over 1250 nodes and in excess of 3600 links utilising 450Mb of underlying resource information comprising a mixture of text, graphics and video. This information was, however, spread amongst 22 MHAs with a single MHA containing, on average, 56 nodes and 160 links. These smaller MHAs are much easier to maintain than the data for the Aberdare case study where there was only effectively a single "MHA" with 600 documents and 7000 links.

Link clusters provide a high-level way of representing the interconnections between abstract ideas or concepts relating to information in the system. By promoting them to first class objects within the information system, they become the primary means for interconnecting related documents of interest rather than structural linking used normally.

The provision of links between the clusters means it is possible to represent complex interrelationships between ideas. The user is able to use the clusters as another way of browsing the available information, but at a conceptual level. Also, controlled access to the information is possible via the addition of permissions to the link clusters. This allows the hiding of areas of the information space that may be confusing to a particular type of user.

The design of the link clusters described here has evolved from that implemented in the original case study. There is still a great deal of *structural* linking used, to provide all of the within-module hypermedia structure. Even though most of these structural links can be automatically generated, it may be possible to use other representations of structure and content to further aid in the construction of large-scale systems.

The PORG example outlined above applies cluster linking to text-based media. However, it is possible to represent abstract relationships between documents of *any* media type. Microcosm has support for defining hypermedia anchors in a rich array of document formats both graphical and text based. These anchors can then be used as lead-ins to the clusters for navigation. Also, the link clusters have only been applied within the Microcosm system. The model proposed is not limited to one particular hypermedia environment, and should be applicable anywhere.

To test how this approach scales, trials are being conducted using a much larger and richer information set. Here, information relating to company assets is stored within many disparate large document databases. The clusters are to be used as a means for defining and browsing relationships that would be difficult, or even impossible to see if the raw data is examined. As there is such a large amount of data, clusters can no longer be authored manually. Tools are under development to allow an author to define rules by which the clusters can be defined. The underlying data is then processed and cluster definitions and interconnections are derived using these rules.

There is still a great deal of work required in order to develop the model presented in this paper. At present, the design and creation of the link clusters and their relationships is performed manually by the author. This gives the author complete freedom to design any particular structure they see fit, but will lead to problems in the future as the information system grows in size. As the number of clusters increases, it will become more difficult for the author to keep track of all the network of clusters, reducing their maintainability. Tools, such as automatic link and document clustering systems [23] will have to be employed in order alleviate this problem.

ACKNOWLEDGMENTS

The authors acknowledge the EPSRC[1] for funding the work under grant numbers GR/H/43038 and GR/L/10482, and Pirelli Cables, Aberdare and Eastleigh for allowing us to use their sites for the case studies.

The authors would also like to thank the Post Office for the use of their information for the final case study.

References

1. Aichison, J and Gilchrist, A. Thesaurus Construction - A Practical Manual, Aslib, 1972.

2. Beitner ND, Microcosm++: Development of a Loosely Coupled Object-Based Architecture for Open Hypermedia Systems, Ph.D. Thesis, University of Southampton, 1995.

3. Borgida A, Description Logics in Data Management, IEEE Trans Knowledge and Data Engineering 7(5), 1995, pp. 671-782.

4. Crowder RM, Hall W, Heath I, Bernard R, Gaskell D**.** A Hypermedia Maintenance Information System. IEE Computing and Control Engineering Journal, Vol. 7 No. 3 June 1996 pp121-128.

5. Crowder RM, et al. ITE - Multimedia Information Systems as an Operational Interface within the Advanced Manufacturing Environment, Final Report. SERC Grant : GR/H 43038 University of Southampton.

6. Crowder RM and Hall W. Industrial Information Resource Management. In Conference Proceedings Integration in Manufacturing, Galway, Ireland, 2-4 October 1996, pp206-214.

7. Crowder RM, Wills GB, Heath I, Hall W. The Application Of Hypermedia In The Factory Information Environment. IEE 5th International Conference on FACTORY 2000, Cambridge, UK 2-4 April 1997.

---

[1] Engineering and Physical Sciences Research Council

8. Cunliffe D, Taylor C, Tudhope D. Query-based Navigation in Semantically Linked Hypermedia, Hypertext '97 Proceedings (Southampton, England, April 6-11, 1997), ACM Press, pp. 87-95.

9. Davis HC, Hall W, Heath I, Hill GJ, Wilkins RJ. Towards an Integrated Information Environment with Open Hypermedia Systems. In Fourth ACM Conference on Hypertext (ECHT '92) (Milan, Italy, December 1992), D. Lucarella, J. Nanard, N. M. and P. Paolini, Eds, pp. 181-190.

10. DeRose, SJ and Durand DG. Making Hpermedia Work, Kluwer Academic Publishers, 1994, pp 56-57.

11. Fountain AM, Hall W, Heath I, and Hill G. MICROCOSM: An Open Model for Hypermedia with Dynamic Linking. In Rizk A, Streitz N & Andre J. eds. Hypertext: Concepts, Systems and applications. The Proceedings of the European Conference on Hypertext, INRIA, France, pp 298-311. Cambridge University Press.

12. Garzotto F, Mainetti L, Paolini P. Information Reuse in Hypermedia Applications. Hypertext '96 Proceedings (Washington DC, March 16-20, 1996), ACM Press, pp. 93-101.

13. Goble C, Haul, S Bechhofer. Describing and Classifying Multimedia using the Description Logic GRAIL. SPIE Conference on Storage and Retrieval of Still Image and Video IV San Jose, 1996

14. Furuta, R, and Stotts, PD. Programmable Browsing Semantics in Trellis. Hypertext '89 Proceedings (Pittsburgh, Pennsylvania, November 5th-8th 1989), ACM Press, pp. 27-42.

15. Goldfarb, CF, The SGML Handbook, Oxford University Press, 1990.

16. Hsu LH, Liu P and Dawidowsky T. a Multimedia Authoring-in-the-Large Environment to Support Complex Product Documentation. Multimedia Tools and Applications 8(1), 1999, pp. 11-64.

17. Lewis, PH, Davis, HC, Dobie, MR, Hall, W. Towards Multimedia Thesaurus Support for Media-Based Navigation. Proceedings of the First International Workshop, IDB-MMS '96 (Amsterdam, The Netherlands, August 1996), pp. 83-90.

18. Malcolm KC, Poltrock SE, Schuler D. Industrial Strength Hypermedia: Requirements for a Large Engineering Enterprise. Hypertext '91 Proceedings (San Antonio, Texas, 1991), ACM Press, pp. 13-24.

19. Marsall, CC, Two Years Before the Mist: Experiences with Aquanet. Proceedings of the ACM Conference on Hypertext (Milano, Italy, Nov 30 - Dec 4, 1992), ACM Press, pp. 53-62.

20. Marshall, CC, Halasz, FG, Rogers, RA and Janssen, WC. Aquanet: A Hypertext Tool to Hold Your Knowledge in Place. Hypertext '91 Proceedings (San Antonio, Texas, 1991), ACM Press, pp. 261-275.

21. Marshall, CC, Shipman III, FM, Coombs, JH. VIKI: Spatial Hypertext Supporting Emergent Structure. ECHT94 Proceedings (Edinburgh, September 18-23, 1994). ACM Press, pp.13-23.

22. Maurer, H. HyperG and HM-Card. In: HyperG now Hyperwave. Addison Wesley 1996, Chapter 27.

23. Weiss, R, Vélez, B, Sheldon, MA, Manprempre C, HyPursuit: A Hierarchical Network Search Engine that Exploits Content Link Hypertext Clustering. Hypertext '96 Proceedings (Washington DC, March 16-20, 1996), ACM Press, pp. 180-193.