

An Open Framework for Integrating Widely Distributed Hypermedia Resources

Stuart Goose, Jonathan Dale, Gary Hill,
Dave De Roure, Wendy Hall
{sg93r,jd94r,gjh,dder,wh}@ecs.soton.ac.uk

University of Southampton
Multimedia Research Laboratory
Southampton, SO17 1BJ, UK

Abstract

The success of the WWW has served as an illustration of how hypermedia can enhance access to large repositories of distributed information. However, current distributed hypermedia systems offer limited hypermedia functionality which cannot easily be augmented or simply applied to existing applications and data formats. This paper describes a model and prototype, based on the Microcosm system, of a reflexive and extensible distributed hypermedia framework for managing hypermedia application development, sharing and re-use.

1. Introduction

The use of hypermedia-based navigation as the basis of a rich, distributed information environment is now widely accepted, as illustrated by the success of the World Wide Web (WWW) [2]. However, the hypermedia model used by the WWW and other contemporary systems are very simple, offering only limited linking functionality and oriented towards the delivery of information authored specifically for use within these systems, rather than providing a range of services to people working with distributed information.

The system described in this paper aims to provide a flexible and extensible system for the provision of distributed hypermedia functionality, which will allow users to exercise control over the way in which they choose to use and share information. The model is a development of that used by the Microcosm system [9], an open hypermedia framework developed at the University of Southampton. The Microcosm system is based upon a set of autonomous, communicating processes which together provide complete hypermedia functionality. This modular system allows straightforward introduction of alternative functions and is clearly well suited to distributed operation.

The flexible hypermedia architecture that we have developed extends the current Microcosm architecture to support scalable, distributed operation, whilst retaining the extensibility of the existing system. Unlike most distributed hypermedia systems, it is designed to provide an underlying hypermedia layer to a user's preferred

environment. The system will allow hypermedia functionality to be accessed from a wide range of information resource bases, provide access to shared resources made available by other users and allow users to integrate distributed resources into a single, cohesive information space. We will show how the model developed will support collaboration and information sharing between users, and describe how the distributed architecture may be extended to support a wide range of multimedia types.

2. Existing distributed hypermedia systems

Many existing systems, such as the WWW, although open in terms of the platforms and protocols supported and the ability to handle distributed operation, do not explicitly support a flexible hypermedia model. The WWW in particular suffers due to its reliance on the Hypertext Markup Language (HTML) [3] for the provision of linking. This prevents hypermedia navigation from taking place within documents stored in other formats. The WWW hypermedia model embedded within HTML is restricted to links marked up in documents, which makes it difficult to provide linking as a third party activity and to extend the available functionality seamlessly. It is possible to implement a variety of more advanced functions for the WWW community through the use of CGI scripts that execute on WWW servers, but this is restricted to dynamic generation of HTML-based documents. The latest generation of WWW browsers promise further advances through the use of executable objects such as Java applets [17], and third party extensions such as the Distributed Link Service [6] offer alternative linking facilities, but this approach is not fully satisfactory as a basis for a seamless information system.

The Hyper-G system [1] whilst superficially similar to the WWW, offers some advantages due to the separation of link information into link databases. This facilitates the application of links to a wider range of document formats, for example images or postscript documents, in a common manner. The hypermedia model supported, like that of the WWW, is based on simple point-to-point links, although it is augmented by an integrated search engine. The Hyper-G

server design provides some degree of scalability through the use of replication across servers and its use of separate link management promotes the maintenance of link integrity, but it does not provide general extension facilities to include alternative link models or to incorporate other information systems.

Hyperform [19], goes further by providing a fully extensible hypermedia server, or hyperbase. This provides a set of object classes to allow the hypermedia services provided to be tailored to particular requirements and facilities to integrate external tools into the system. This clearly provides scope for the development of a wide range of hypermedia functionality. The integration facilities suggest that inter-operation between Hyperform servers would also be possible, but this is not explicitly supported.

The Virtual Notebook System (VNS) [5,15], provides an open hypermedia architecture based on the use of an underlying database system, thus allowing the development of custom programs that utilise the hypermedia facilities. The system also supports interaction between different VNS servers, allowing workgroups to be supported by dedicated servers whilst still permitting the sharing of information between individual workgroups. VNS provides less scope for extending the hypermedia model and integrating other information services than Hyperform does, but it has better support for controlling shared access to information.

Although all the systems described support some form of distributed operation, none currently provide the degree of flexibility which a distributed version of Microcosm could provide. Initial work on a distributed version of Microcosm [12] refined the original communication architecture to demonstrate the feasibility of fine-grained distribution of hypermedia functionality. This system enabled users to 'publish' their local Microcosm processes (link databases, search engines, etc.) for incorporation into remote users' configurations. It allowed a wide range of distributed configurations to be supported, but was best suited to small scale, local sharing of information. It lacked both the ability to manage logical configurations of processes and the facilities for one user to easily manage processes spread across multiple platforms.

Popular systems, such as the WWW, have illustrated a demand for global access to information, and for many people the WWW also served as an introduction to the concept of hypermedia; demonstrating its potential in providing an integrated and unifying paradigm for information organisation. As demands on the WWW grow, it is becoming apparent that the simple hypermedia model is not adequate for large scale information management, navigation and discovery. Open hypermedia systems, such as Microcosm, offer more flexible facilities that can overcome these problems. The following sections describe

the distributed model for Microcosm that has been developed, the way in which it can facilitate more flexible access to distributed information, and the initial development of a system which implements this model.

3. A reflexive model for the construction of distributed hypermedia applications

A *hypermedia object* is the smallest element that is managed by a hypermedia system, that is, a document or a process. This is the model that has been adopted by the WWW; everything is considered to be a document and even processes themselves (apart from the WWW infrastructure) exist to generate other documents. As the granularity of this system exists at the document level, it is not difficult for users to bind documents through links into sharable information resources, but it is difficult to logically group these individual resources into sharable components that can be associated with larger repositories that are organised by context.

To provide a model that retains the desirable properties that exist at the document level and that also allows hypermedia objects to be associated with larger information resources that possess context, the concept of a *hypermedia application* was introduced. A hypermedia application is an entity that comprises a resource base of hypermedia objects (documents, link databases and processes). Documents can consist of any media type, while link databases describe the link structure that exists between documents and processes define the set of operations that can be performed on the information. To allow users to share information between each other, applications possess the characteristic of reflexivity, that is, applications can make references to other applications. In this way, users can make references to other applications and associate them with their own hierarchical structures to build information resources which can seamlessly

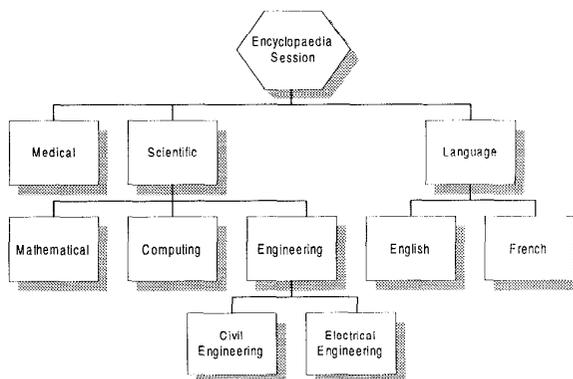


Figure 1 : Application Integration to Form Distributed Hypermedia Resources

comprise a number of diverse and distributed applications.

Before a hypermedia application can be considered for publishing to a wider environment, a number of distributed issues need to be addressed:

- *Ownership*. Hypermedia applications should be able to be attributed to a particular user, or set of users, who are considered to be its authors. These owners are responsible for managing their applications and making them available to the global community.
- *Sharing*. Applications should be sharable between users, to facilitate information dissemination, reuse and co-operative working.
- *Permissions*. The interactions between users and applications need to be protected by a set of permissions that allow users to control access to the applications that they have made sharable.

Applications can be imported into a *hypermedia session* to allow users to construct larger information resources. Users can have multiple sessions which represent the information resources that they are interested in, either for work or recreational purposes. These collections permit users to apply group operations to logical sets of information, for example, issuing queries to all applications within a session.

The hypermedia framework proposed is illustrated in figure 1 *User 1* has set-up a number of sessions within their environment and at the top level, *session 1* contains two applications; *application 1.1* is local and has been authored by *user 1* (denoted by a solid box) and *application 1.2* is remote and has been imported from another user (denoted by a dashed box). *Application 1.2*, in turn, contains two more applications; one which is local and one which is remote to this application. Although *user 1* has four applications in *session 1*, *application 1.2.1* and *application 1.2.2* are presented transparently to the user as *application 1.2*. This means that when *user 1* makes use of the resources in either *application 1.2.1* or *application 1.2.2*, they perceive that they are utilising the resources of *application 1.2*. This is advantageous because it abstracts communication and connectivity issues from the user, who can concentrate on building hypermedia applications and creating link structures between them.

Application 2.2 in figure 2 is present in two forms; a shared view from the perspective of *user 1* (maybe a co-author) and an owner's view from the perspective of *user 2* (probably the author). *User 1* has developed *application 2.2* and *application 2.2.1* and has made it sharable, through some mechanism, via *application 2.2*. *User 2* has decided that they wish to make use of this application and have incorporated it into *session 2* of their environment. The operations that *user 1* can perform on this application will be determined by the permissions that have been exported

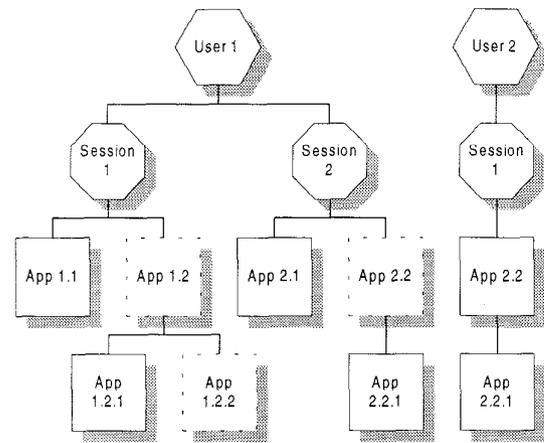


Figure 2 : Application Integration and Sharing in a Distributed Environment

with it.

An example of this abstract model in context is give in figure 1. *Encyclopaedia* is a session that is comprised of many applications which describe its behaviour, both in terms of information content and functionality. Such an application is a dynamically extensible entity within which new applications can be plugged or unplugged at any time. Because the application is not monolithic, different interpretations can be generated from restricted or expanded views on the same sets of applications. For example, a *Junior Encyclopaedia* might reuse some of the constituent applications of *Encyclopaedia*, but restrict other applications that are considered to be of either an adult nature or are too advanced for children to understand.

This hypermedia model provides for a very flexible and extensible model that allows a user to author multiple and different views on the same sets of applications, depending upon the context within which they are being placed. Moreover, these contextual hypermedia applications can be shared across a number users, thus promoting information dissemination and reuse on a global scale.

3.1. A distributed framework for integrating hypermedia applications

The mapping between the abstract hypermedia model into a physical representation is achieved through an event-driven process framework. A physical hypermedia application realisation tightly encapsulates a collection of processes (and associated preferences and configurations), an arbitrary amount of data together with the optional characteristic of reflexivity, and can be expressed as:

$application ::= \{process\} \times \{data\} \times [\{application\}]$

Such an application can itself be considered as a self-contained, logical entity and the application paradigm provides an abstraction for modelling communities of co-

operating processes. Each discrete process publishes a suite of services, usually operates on specified data objects and responds to a designated set of events. As an application comprises executing processes it can be considered as a dynamic object which can respond quickly to those events that occur within the system.

Due to their well defined and modular nature, such processes predispose themselves to being distributed over interconnected networks. To deliver this reflexive model for hypermedia applications provision for flexible communication between widely distributed asynchronous processes is essential.

3.2. The Heterogeneous Communication Model

The Heterogeneous Communication Model (HCM) [10,11] is an open framework that we have designed to facilitate the integration of asynchronous processes so that they can function effectively within a heterogeneous, distributed, multi-user environment. Distributed systems such as this have an inherent reliance upon the underlying communication substrate and it is therefore crucial that this layer exhibits a range of desirable qualities:

- *Scalability*. The ability to add processes to the system with minimal impact and a reduction in the number of intermediate steps a message has to take.
- *Openness*. Openness allows the system to integrate with other third-party applications, thus allowing applications to benefit from improved interconnectivity and interoperability.
- *Flexibility*. By promoting a “plug and play” rationale, the functionality of the system can be dynamically augmented or diminished through the addition or removal of further processes.

By incorporating these characteristics into the design of the HCM, we have sought to embody many of the requirements of a distributed communication layer and to provide an open framework above which a powerful distributed hypermedia system can be developed.

3.3. Process addressing

For direct communication to take place, one process must be able to uniquely identify or address the process where it wishes to send a message. The addressing scheme used within the HCM is based upon a metaphor found commonly in computer filing systems. Each file in a filing system can be uniquely addressed by using the fully qualified path name, for example, `/papers/multimedia/ieeemm96.eps`. It is possible to issue a query along with some criteria that returns a list of matching file names. The matching criteria is generally expressed by using *wildcards* and regular expressions. A subset of these techniques have

been combined to produce a flexible process registration and addressing mechanism. Although the filing system metaphor may imply that process registrations are hierarchical, the addressing scheme does not impose any rigid structure. This scheme is explained in greater detail in [10,11]

3.4. An architecture to support widely distributed process interaction

One of the key principles of distributed computing is that any processing need not be performed locally, but on the most appropriate machine available. It was decided that the design of the HCM would embody this principle from the outset and the infrastructure to support this activity required three core components.

A central module, called a *router*, is responsible for coordinating all communication between participating processes wishing to pass messages across machine boundaries. The router acts as a bureau where processes, using the addressing scheme as the vehicle, can dynamically advertise and withdraw their services and also post messages to other registered service providers.

A *process manager* is also started with each session to govern load balancing across the network and the distributed invocation of processes. Via its user interface, processes can be remotely configured and managed.

A single *daemon* process exists to serve the local network *domain* (an administrative term that is used to describe a set of logical machines where HCM processes can execute). Due to the fact that processes can execute anywhere within the local network, routers are required to register their network location with the daemon. As the daemon provides the sole, fixed point of contact within the domain, processes can query it to determine the network location of routers. The dynamic allocation and registration of network connections removes any explicit network location dependencies from processes within the HCM infrastructure.

The HCM supports a variety of alternative configurations; single user, multi-user and distributed multi-user environments.

A conceptual view of the third configuration is presented in figure 3. Here, *user 1* in *domain 1* initiates an action causing the daemon of *domain 2* to be contacted, requesting a communication transaction between the two routers; once a virtual connection has been established, negotiation and information transfer can occur. The process registration entries declared as *published* (that is, services that are available to other users) are then exchanged between the two routers. This activity alleviates superfluous message passing between routers. The scalable framework for providing widely distributed, peer-to-peer

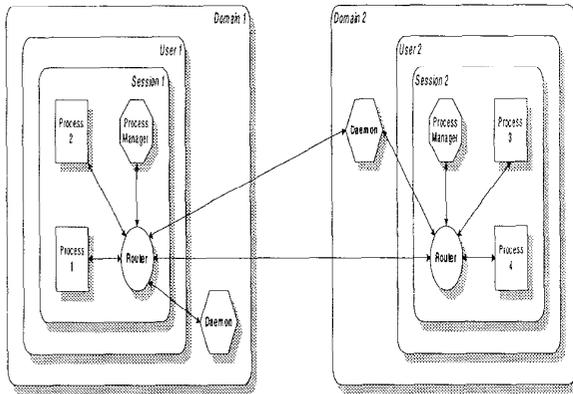


Figure 3 : Interaction Between Processes in a Distributed, Multi-User Environment
communication can now be exploited.

3.5. Microcosm: The Next Generation

Microcosm: The Next Generation (Microcosm TNG) is a prototype that retains the core philosophy of the Microcosm system and has been developed in response to the driving requirement for extending the principles across a distributed environment. Within the Microcosm hypermedia model there are a number of core elements that work together to provide a minimum level of hypermedia functionality:

- **External link storage.** The concept of link separation is well documented and its advantages clear when applied to open hypermedia [1,16]. A *linkbase* process maintains the link database and handles all link queries and creations. Multiple linkbases can exist to break document link connectivity into logical groupings.
- **Presentation and browsing.** Displaying links and allowing users to browse through links is handled by a range of *viewers*, primarily one for each class of media (text, bitmap, sound, etc.).
- **Link selection.** The results of link queries that have been submitted by a user are handled by the *available links* process. This presents the links to the user for selection and then invokes the appropriate viewer to render the document.
- **Link authoring.** To allow links to be created within a link database, an intermediary process (the *linker*) performs preprocessing on the new link to determine its details (for example, the start and end anchor positions, the type of the link, the description of the link, etc.). Once this information has been collected, the completed link is forwarded to the relevant linkbases for creation.
- **Document organisation.** The document database (or *docuverse*) was introduced to provide a layer of abstraction between the documents themselves and the

filing systems upon which they reside.

The Microcosm TNG prototype has been developed using the HCM and has taken the core concepts identified previously and modelled them as discrete, distributed processes. These processes (linkbases, viewers, available links, linkers, docuverses, etc.) were described in terms of the messages that they send and receive and the processing that they perform. Processes are considered to be peer entities as communication can occur with other processes asynchronously. The distribution of both processes and messages is handled by the HCM subsystem. In this way, the naming of a process is kept independent of its location.

Resources are distributed through the application entity to which they belong. Once a user has authored an application, they can publish it to the world at large. The Microcosm TNG system provides a mechanism for allowing remote applications to be included transparently within a user's session. Also, since an application can comprise both data and processes, this inclusion not only enhances the functionality of the user's session, but also increases their available dataspace; all actions (for example, link queries) will be propagated automatically to all of the user's local *and* remote applications.

In figure 4, for example, the two users, Jonathan and Stuart, both have sessions running in their respective domains, which could be distributed across the Internet. In a strict local session, all link queries and link creations only apply to the applications that are bound to a particular session (for Jonathan, this is the *Tigers* application). Jonathan and Stuart can create their applications in isolation, importing local documents and media, and making links between them.

However, imagine that Stuart wishes to create a new application, called *Cats*, which comprises both the *Tigers* and *Lions* applications. This new application is made in his local session and he forms a connection to the *Tigers*

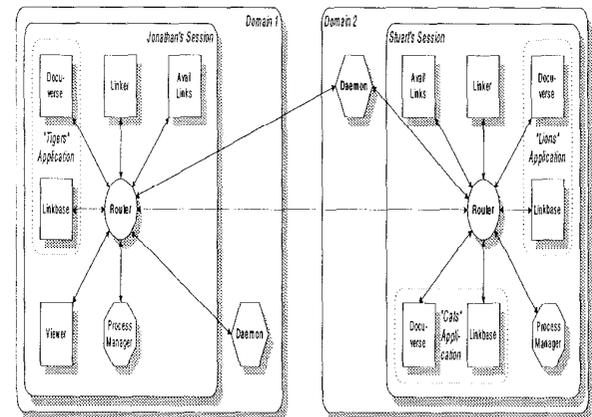


Figure 4 : Microcosm TNG Users Sharing Applications Across a Distributed Environment

application through domain the address of Jonathan, which he obtained previously via a resource discovery agent. Once connected, through the HCM subsystem, Stuart has access to both applications and he can create links between them and import new documents as appropriate.

Upon its completion, the *Cats* application can then be published for other people to connect to and can be reused in a similar fashion as indicated previously. It is important to note that when users access this application, connection to the *Tigers* and *Lions* applications, subsequent communication and the manipulation of hypermedia links are handled transparently.

From the work that we have undertaken so far, the distributed open hypermedia system that has been developed using our reflexive hypermedia model above the HCM shows great promise for the future. The prototype has illustrated that the distribution of users, data and processing is possible and advantageous, and that future work lies in refining the prototype for information sharing and reuse of applications on a global scale.

4. Service layers

The provision of integrated and extensible multimedia handling within Microcosm TNG is provided through two support service layers (figure 5); a *presentation layer* and a *quality of service layer*. The presentation layer is responsible for determining how different media formats are managed and presented dynamically and consistently to the user across multiple platforms, and the quality of service layer is responsible for negotiating bandwidths with the underlying network layer to support the transport requirements of the different media formats.

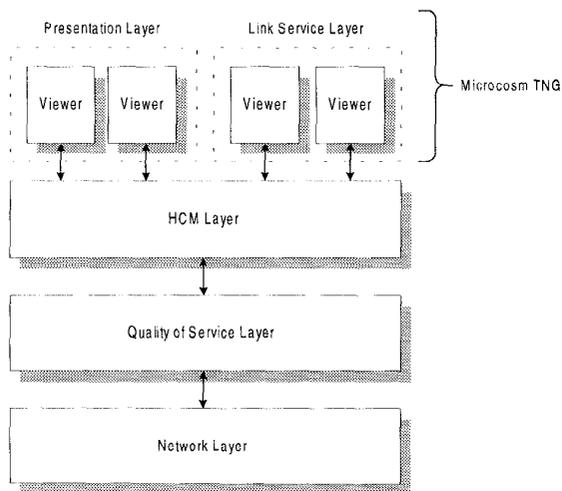


Figure 5 : Service Layers to Support Multimedia Extensibility

4.1. Heterogeneous presentation layer

A viewer is responsible for providing an interface for allowing a user to manipulate a particular data type and also to create and traverse links. The flexible architecture of Microcosm facilitates the easy integration of new applications; through the introduction of new viewer modules, extensible multimedia support has been achieved. Although a trivial task to integrate the array of existing Microcosm media viewers within Microcosm TNG, moving into a distributed environment forces the issue of heterogeneous user interfaces to be confronted.

Through the careful engineering of portable code, a process can be implemented enabling it to be compiled for multiple platforms. Yet, it is very difficult to apply the same techniques for the implementation of user interfaces as no Graphical User Interface (GUI) standard is supported across the majority of machine platforms.

A solution to this problem that we have developed is to split a process into two separate, but communicating physical processes; a Process Handler (PH) and an Interface Handler (IH) [7]. The PH deals with the generation of user interfaces and the processing of data, and the IH deals with choosing a suitable viewer for a given media format and handling user input data from the viewer and preprocessing it into a format suitable for the PH.

A prime requirement of any GUI-based process is the ability to either construct new user interfaces dynamically in response to certain data conditions, or to allow user interfaces to include dynamic data. Therefore, the construction of a user interface must occur within the PH to give the system the flexibility needed to be able to construct interfaces dynamically. As the IH is now unaware of the data that passes through it, it can be reused for each PH, thus decreasing the development time for processes and systems.

To abstract the PH from specific interface details, a platform and GUI independent display description language is required. This enables the PH to build and subsequently communicate new user interfaces to the IH in terms that are descriptive but abstract. HTML is such a language that is currently in wide-spread use within the WWW. When a user interface is to be displayed, a HTML viewer (such as Netscape or Mosaic) can be invoked by the IH to render the interface into GUI-dependent terms.

The use of HTML means that processes can be developed far more rapidly, since less programming effort is required in producing the interfaces. However, due to HTML's reliance upon the configuration abilities of the viewer that is rendering it, there exists a lack of seamless integration for new media formats. A superior level of integration is currently under investigation using Java technology. By providing a generic media viewer that can

work within a Java-aware browser and is aware of the underlying hypermedia layer, each new media type can be handled by a separate applet. Such an applet needs to possess enough information about a media format to be able to render the data and also to present a manipulation interface to the user.

4.2. Synchronisation support

Various researchers have confronted the issues of specifying and delivering synchronised multimedia documents, for example, [4,18]. This is an important area for the future development of multimedia information systems. Of the many proposed solutions to this non-trivial task, the most favourable is to provide support amongst the layers of the operating system. Existing mechanisms such as semaphores, message passing and clock scheduling have proved inadequate on their own, so mechanisms such as triggering have been introduced to complement these existing services. These problems are made worse when considering what types of mechanisms are required for synchronous co-ordination within a distributed environment.

Alternative high level approaches may be sought through designing the media viewers in such a way that they may communicate and hence, collaborate. A generic or composite media viewer could assume responsibility for playing and rendering all media formats, thus giving it absolute control over the rate at which each medium may progress. Research is currently being followed to extend the IH to act as a composite media viewer that directs the synchronisation of the different media format viewers from a script describing the interdependencies that exist between them.

4.3. Quality of Service layer

The guaranteed delivery of continuous media across a network is an issue that serious multimedia information system designers must be concerned with. The multimedia and networking research community have been addressing this central problem for a significant period and over this time have proposed many worthy solutions [13].

Microcosm TNG was developed to enhance the scalability, reliability and performance of the hypermedia system through distributed processing practices. It is envisioned that the technologies alluded to above for providing negotiation and guaranteed bandwidth delivery, can serve the higher layers of our model.

Preliminary experimentation is also currently being conducted to establish how best exploit the advantages of ATM within the context of a layered communication architecture.

5. Future work

The initial implementation of the model described in the previous sections has proved the viability of the communication architecture and the Microcosm TNG prototype has allowed preliminary experimentation. This prototype currently supports multi-user interaction and application sharing across a network, but users are restricted to a single session. As the prototype system is extended and refined, there are a number of additional areas that need to be addressed:

- *Security.* A comprehensive model for security is necessary in order to prevent unauthorised access to shared information. As mentioned earlier, the richness of the permissions system will determine how flexible the system is when sharing hypermedia applications between users.
- *Integration.* A strength of Microcosm has been its ability to integrate with almost any third party application, as demonstrated by Davis [8]. To develop this into a distributed environment, we are currently experimenting with emerging technologies such as Common Object Request Broker Architecture (CORBA) [14], which displays great promise for improved distributed and heterogeneous application inter-operability.
- *Agent technology.* The opportunities afforded by incorporating agents within the Microcosm TNG system can assist in the areas of resource discovery, link maintenance, navigation assistance and third-party application integration.

6. Conclusions

The success of the WWW in using hypermedia facilities to provide access to vast amounts of distributed information provides firm evidence for the viability of hypermedia navigation as a component of an information management system. However, we have shown how few existing hypermedia systems provide all the facilities necessary to act as the basis for such a system.

The distributed hypermedia model outlined in this paper is designed to act as a modular framework within which a flexible and extensible information system may be developed. It achieves this through its ability to be extended through several levels:

- The system is based on a peerless communication model which allows a wide range of information resources to be integrated within a single, cohesive architecture without requiring the user to address each resource independently.
- By altering the component processes that are

incorporated, a user is able enforce alternative views upon the same information.

- By providing a complete underlying hypermedia link service, the system allows a variety of existing applications to be integrated into the hypermedia functionality.
- The ability to package resources into logical groups which may then be published for other users to utilise allows scalable access to information.
- The ability to utilise widely distributed components allows users to build a system which makes the best use of available resources.

The utility of this logical architecture is demonstrated by the Microcosm TNG prototype that has been developed. There are clearly a number of issues that must be addressed when attempting to provide a robust distributed system such as that envisaged, some of which are described in the Future Work section. We are currently extending the prototype system to address these issues.

7. References

- [1] ANDREWS, K., KAPPE, F. and MAURER, H., Serving Information to the Web with Hyper-G, *Computer Networks and ISDN Systems*, 27(6), pages 919-926, 1995.
- [2] BERNERS-LEE, T., CAILIAU, R., GROFF, J. and POLLERMANN, B., *World-Wide Web: The Information Universe*. In: *Electronic Networking: Research, Applications and Policy*, 2(1), Meckler Publishing, Westport, USA, pages 52-58, 1992.
- [3] BERNERS-LEE, T. and CONOLLY, D., *Hypertext Mark-up Language Specification 2.0*, 1995, http://www.w3.org/pub/WWW/MarkUp/html-spec/html-spec_toc.html.
- [4] BUCHANAN, M.C. and ZELLWEGGER, P.T., Specifying Temporal Behaviour in Hypermedia Documents, In: Lucarella, D., Nanard, J., Paolini, P., *Eds. The Proceedings of the ACM Conference on Hypertext, ECHT '92*, Milano (November), ACM, pages 181-190, 1992.
- [5] BURGER, A. M., MEYER, B. D., JUNG, C. P. and LONG, K. B., The Virtual Notebook System. In: *Hypertext 91*, Proceedings of Third ACM Conference on Hypertext, San Antonio, Texas, USA (December), ACM Press, pages 395-402, 1991.
- [6] CARR, L., De ROURE, D. C., HALL W., and HILL, G. J., The Distributed Link Service: A Tool for Publishers, Authors and Readers, *The World Wide Web Journal Issue One*, Conference Proceedings of the Fourth International Word Wide Web Conference, Boston, USA, O'Reilly and Associates Inc., pages 647-656, 1995.
- [7] DALE, J., Achieving User Interface Heterogeneity in a Distributed Environment, 1996, <http://wwwcosm.ecs.soton.ac.uk/TNG/papers/m96-1.htm>.
- [8] DAVIS, H. C., KNIGHT, S. J. and HALL, W., Light Hypermedia Link Services: A Study of Third Party Application Integration. In: *ECHT '94 Proceedings*, Edinburgh, Scotland (September), ACM Press, pages 41-50, 1994.
- [9] DAVIS, H. C., HALL, W., HEATH, I., HILL, G. and WILKINS, R. J., Towards an Integrated Information Environment with Open Hypermedia Systems. In: D. Lucarella, J. Nanard, M. Nanard and P. Paolini, *Eds. ECHT '92*, Proceedings of the Fourth ACM Conference on Hypertext (November 30-December 4), Milan, Italy, ACM Press, pages 181-190, 1992.
- [10] GOOSE, S., The Design Of A Generic, Yet Customisable, Distributed Communication Framework, 1995, <http://wwwcosm.ecs.soton.ac.uk/TNG/papers/m95-4.html>.
- [11] GOOSE, S., Extending The Heterogeneous Communication Model: A Scaleable and Seamless Architecture to Support Widely Distributed Process Interaction, 1996, <http://wwwcosm.ecs.soton.ac.uk/TNG/papers/m96-2.html>.
- [12] HILL, G. and HALL, W., Extending the Microcosm Model to a Distributed Environment. In: *ECHT '94 Proceedings*, Edinburgh (September 18-23), Scotland, ACM Press, 32-40, 1994.
- [13] LITTLE, T., GHAFLOOR, A., Synchronisation and Storage Models for Multimedia Objects, *IEEE Journal on Selected Areas of Communications*, 8(3), pages 413-427, April 1990.
- [14] Object Management Group, The Common Object Request Broker: Architecture and Specification, OMG Technical Document, Number 91-12-1, Revision 1.1, December 1991.
- [15] SHIPMAN, F. M. III, CHANEY, R. J. and GORRY, G. A., Distributed Hypertext for Collaborative Research: The Virtual Notebook System. In: *Hypertext '89 Proceedings*, Pittsburgh, USA (November), pages 129-136, 1989.
- [16] SMITH, K. E. and ZDONIK, S. B., Intermedia: A Case Study of the Differences Between Relational and Object-Oriented Database Systems, *OOPSLA '87 Proceedings* (October), pages 452-465, 1987.
- [17] Sun Microsystems Inc., <http://java.sun.com/>.
- [18] STEINMETZ, R., Synchronisation Properties in Multimedia Systems, *IEEE Journal on Selected Areas in Communications*, 8(3), pages 401-412, April 1990.
- [19] WILL, U. K. and LEGGETT, J., Hyperform: Using Extensibility to Develop Dynamic, Open and Distributed Hypertext Systems. In: D. Lucarella, J. Nanard, M. Nanard and P. Paolini, *Eds. ECHT '92*, Proceedings of the Fourth ACM Conference on Hypertext, Milan, Italy (November), ACM Press, pages 251-261, 1992.
- [20] WILKINS, R. J., De ROURE, D. C., HALL, W. and DAVIS, H. C., The Role of Agents in Multimedia Information Systems. In *Proceedings of the Intelligent Agents and the Next Information Revolution*, Manchester, UK (May), pages 14-23, 1995.