

Issues of Information Semantics and Granularity in Cross-Media Publishing

Moira C. Norrie and Beat Signer

Institute for Information Systems, ETH Zurich
CH-8092 Zurich, Switzerland
{norrie,signer}@inf.ethz.ch

Abstract. While there have been dramatic increases in the use of digital technologies for the storage and processing of information, the affordances of paper have ensured its retention as a key information medium. Recent developments in digitally augmented paper provide the potential to embed active links within printed documents, thereby turning paper into an interactive medium. In this paper, we address the issues of information granularity and semantics that arise in integrating paper as a first-class interactive information medium in hypermedia systems and show that the information server is vital in realising the true potential of this vision. Further, we discuss the authoring issues of cross-media information environments and the forms of tools required to support the various categories of authoring activity.

1 Introduction

Earlier predictions of the paperless office in which paper documents would be replaced entirely by digital documents no longer seem realistic [19, 23]. Paper as a medium has many advantages over digital media in terms of how people can work with it, both individually and in groups. It is portable, cheap and robust. It is much more convenient to scan through a book by rapidly flicking through pages than to browse a digital document. Paper also supports forms of collaboration and interaction that are difficult to mimic in current digital worlds. Whereas the focus in the past has been to replace paper, increasingly, the trend is towards integrating printed and digital media, thereby achieving the best of both worlds.

Integration of printed and digital media can be achieved through technologies for encoding active links on paper in such a way that special reader devices can detect and activate these links. There have been various technical solutions proposed to achieve this aim and it is beyond the scope of this paper to discuss these in detail. Generally, we refer to these solutions as coming under the heading of *digitally augmented paper* and would recommend the following references to readers interested in more detail of the various technologies [2, 7, 10, 11, 15].

While the paper, printing and reader technologies required to fully support digitally augmented paper are still in their infancy, various forms of reader devices are now emerging in the marketplace. For example, Sony Ericsson's Chatpen uses a camera to track movement on a page based on a special printed grid

pattern of tiny dots developed and licensed by the company Anoto [1]. Expected developments in the near future for invisibly encoding document positions on paper and cheaper reading devices will surely have a dramatic effect on the spread and use of digitally augmented paper.

The question remains as to the potential uses and models of interaction afforded by digitally augmented paper. How can we fully exploit the potential of paper as a client device and integrate it seamlessly into hypermedia systems? While a number of research projects have proposed different variations of digitally augmented paper and presented their visions for its use, none have really addressed the issue of making paper a *first-class medium* in the context of a hypermedia system. This means that it should be possible to link not only from paper to digital resources, but also from digital resources to paper and even from paper to paper.

The key to incorporating paper into hypermedia systems as a first-class medium lies primarily with the server rather than the client technologies. The server must be capable of dynamically mapping document positions to information objects and vice versa. It must manage not only the links, but also their semantics and have flexible means of classifying and associating both digital and printed resources. We argue therefore that the potential functionality and flexibility of the system will be determined by the underlying information model and architecture.

In this paper, we show how issues of information semantics and granularity impact on the richness of the resulting interactive information environment. However, everything comes at a price, and we discuss the consequences for the content provider in terms of complexity and cost. Indeed, we discuss how it could require a major shift away from the traditional practices of publishers towards the content management solutions that have been adopted in the web publishing world.

To explain the issues and possible solutions, we develop in a stepwise manner an information model for an application server for digitally augmented paper. This model is the basis for the information server that we have developed within the European project Paper⁺⁺ [18] which is investigating concepts and technologies aimed at enriching the use of paper in everyday settings. A detailed description of the Paper⁺⁺ server component can be found in [15].

We begin in Sect. 2 with an introduction to cross-media publishing and simple solutions that have been proposed to link printed material to digital resources. Section 3 then discusses issues of granularity and examines in more detail the means of embedding active links within printed documents to arrive at a highly-interactive environment. In Sect. 4, we address the issue of information semantics and how to provide richer information environments. Section 5 deals with the general issues of the authoring of cross-media material and the different categories of system and publishing tools to support them. Concluding remarks are given in Sect. 6.

2 Cross-Media Publishing

It is common nowadays for publishers to produce a package of related materials on different media. For example, the BBC (British Broadcasting Corporation) often produces books and also websites to accompany television documentaries. As a specific example, for their series on ocean life called *Blue Planet*, they have an associated book and a website with games, quizzes and a fact file on different species [3]. In addition, this series was adopted by the Open University [17] in a course on oceanography and a course text book was published to be used in conjunction with the TV programmes which are also available on video and DVD. We will use this example as a small case study to examine the issues of information granularity and semantics that arise in developing a system that could link these various forms of printed and digital materials together through the use of digitally augmented paper. We first note however that such cross-media publishing is not just something of concern to the large commercial content publishers. It is something that most of us do in producing teaching materials such as lecture slides, handouts, manuals, tutorials, exercises, demonstrator applications and course websites.

How is linking across media currently done? As a first remark, it is important to note that, in many cases, it is not done at all. What might at first appear to be a package of related materials, is often actually a set of quite independently produced materials. For example, we used a children's nature encyclopaedia as a basis for a first demonstrator and user studies experiment since both a book and a CD-ROM were available from the same publisher. However, discussions with the publisher and examination of the material revealed that both were not only marketed as independent products, but also designed and developed by separate departments with little or no cooperation and had different content. The Paper⁺⁺ nature encyclopaedia prototype running on a Tablet PC (a) and on a Pocket PC (b) is shown in Fig. 1. For further information about the user studies see [8].




Fig. 1. Nature encyclopaedia application

Where linking from printed documents is provided, it is usually done through a printed reference either to another printed document or to a digital resource by means of a URL. In the electronic version of the document, URLs may be converted automatically to active hypermedia links based on recognised formats.

In the case of the Blue Planet material, the main linking occurs between the Open University course book and the TV programmes (on VHS or DVD). At the end of a section of text, there is a recommendation to view all or part of a programme. These recommendations are indicated by a video cassette symbol placed alongside the text as shown in Fig. 2.

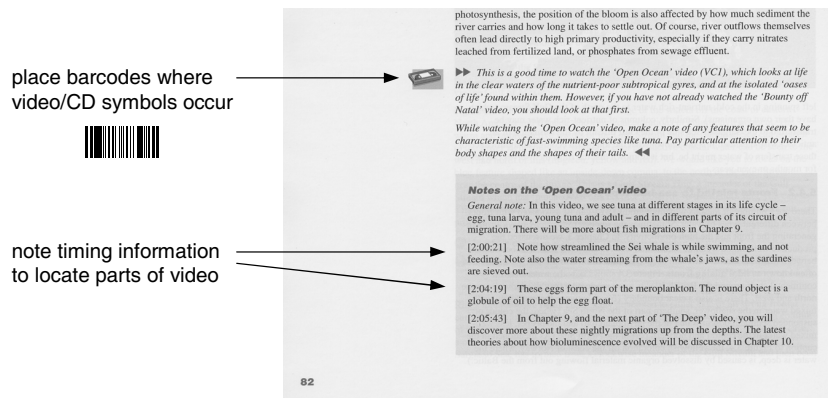


Fig. 2. Page of Open University course book

As a first means of creating an active link from the course book to the video material, we could add a barcode alongside the video symbol as indicated in Fig. 2. The barcode is an encoding of a unique digital resource identifier (ID) and a simple mapping table on the server could be used to map this ID to a multimedia file. In Fig. 3, we show such a mapping and model this as an association **Binds** between a set of **ResourceIDs** and a set of **Resources**. Using normal barcode reader technology, we could play video files when the user swipes or points at the barcode with the reader.

Note that we are using the OM model [14, 16] since the information servers for interactive paper that we have developed are based on this model and its associated OMS data management frameworks [12, 24]. The OM model has a collection construct (represented by the shaded rectangles) to represent semantic groupings of objects and associations between the members of collections are represented by shaded ovals with the cardinality constraints over participation written alongside. The association construct enables relationships to be classified and manipulated directly which becomes particularly useful in supporting link management for hypermedia browsing.

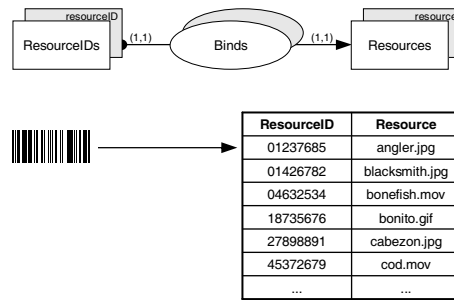


Fig. 3. Simple mapping of IDs to resources

What comments can we make to such a simple linking scheme? The first is that the links are unidirectional, only supporting links from the printed document to digital resources. Secondly, the link paths are restricted to length 1. The activation of a link causes a file to be displayed using the appropriate application, but there is no way to link to further digital resources or to link back to the printed document, or even other printed documents. For example, in the case of the Blue Planet material, it might be useful to have a link to a video and then to be able to follow links to related video clips, entries in the fact file of the website or relevant on-line quizzes and also to images in the Blue Planet book.

These limitations are in many ways due to the over-simplistic information model which simply binds IDs to files. A much richer information model can lead to much richer and more flexible interactive information environments and this is what we want to develop in the remainder of this paper. However, it is important to emphasise that most systems linking printed documents, or indeed other physical objects, to digital resources are based on such a basic mapping model [5, 6, 13, 22]. The focus has tended to be on the hardware and communication technologies and their uses, rather than on the information models. What we want to show is that approaching the problem from an information point of view can take us a step further in realising the potential of *interactive paper*.

3 Interactive Paper

Our goal is to, not only provide links between printed and digital materials, but also support highly-interactive systems that enable the reader to easily move back and forth between the printed and digital worlds.

“Highly-interactive” implies a fine level of granularity of information objects in both the printed and digital worlds. However, most of the Blue Planet material is produced as relatively large chunks of information such as 45 minute TV programmes, sections of text, complete quizzes etc. As described in the last section, the course textbook links entire sections of text to TV programmes. Since the publishers are aware that this level of granularity for video material is

too coarse, they do provide timing information to enable readers to locate parts of a TV programme related to a section of the text or specific themes addressed within the section. To create a richer information environment, a much finer level of granularity must be supported that enables readers to link words, phrases or images to relevant video clips.

Similarly, we note problems with the web-based material in terms of granularity and lack of openness. Quizzes and games are programmed as atomic units and it is not possible to link into or from part of a quiz or game. This might be something useful and worthy of consideration in future publishing scenarios given technologies and tools to support highly-interactive cross media publishing. Also, the “fact files” providing information on species are quite flat, consisting only of complete sections of text with no links. One can imagine this being replaced by a richer database of information about species, their classifications and geographical locations, linked together and also with links to images, video clips etc.

As a first step towards a finer granularity, we could annotate the printed course text with additional barcodes linking pieces of content such as words, phrases and images to digital resources such as texts, images, audio files or video clips. How would the reader know which items were linked? This could be done by simple positioning of barcodes alongside linked objects in the case of images, or possibly with the use of content highlighting and a line connecting the highlighted object with the relevant barcode. However, this would be an unsatisfactory solution with problems arising due to visual interference with texts and images. It would certainly be an inadequate approach if there are large numbers of links. For this reason, researchers have been developing encoding schemes which are invisible, or at least nearly invisible. Work has been done on the use of inductive inks to print special encoding schemes, whereas the solution of the company Anoto [1] uses a pattern of tiny visible dots which appears to the reader as light grey areas that do not create too much visual interference with the document content.

Developing technologies that enable invisible, or nearly invisible, encodings brings us back to the issue of how users know where links are located within a page. In terms of published content with embedded active links, highlighting can be used or even recognised conventions that have arisen from the use of web browsers where users have come to expect content items such as headers, images and underlined or blue text to be possible links. We have also experimented with the use of audio feedback to inform users that they have activated links. With the further development of reader devices, other forms of audio or force feedback from the reader device itself could be used to inform a user of the existence of a link on rollover.

Instead of using a “printed resource ID”, whether it be a barcode or some other form of encoding, we now want to shift to a solution that binds elements of printed material, such as pieces of text or images, to digital resources by means of these invisible encodings. We can do this by associating these elements with logically defined *active areas* of the page as indicated in Fig. 4. The reader

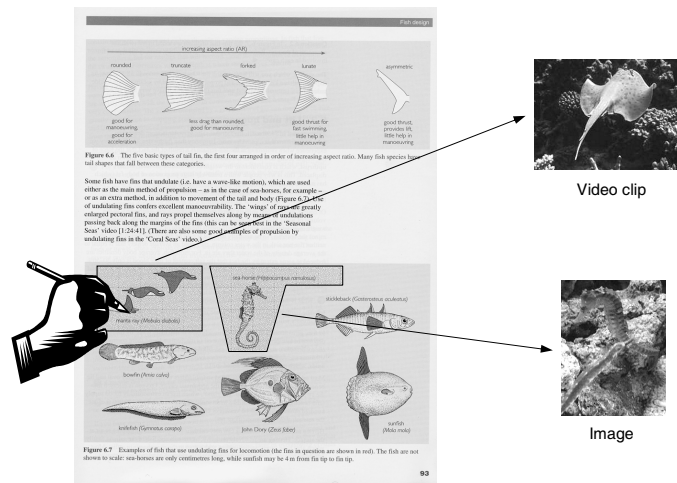


Fig. 4. Active areas

device will deliver an (x, y) page coordinate and the server will determine which, if any, active area contains this point and then bind it to the appropriate digital resource. This implies that entire pages are encoded with position information; for example using a specially designed grid of barcodes printed using invisible inductive ink. Again, a number of technical solutions have been proposed to enable this.

The active areas can be defined as one of a number of simple geometrical shapes or as a complex shape composed from other shapes. In Fig. 5, we show the corresponding part of the information server's schema.

We have replaced the notion of **ResourceIDs** with a general notion of **Anchors** that are bound to **Resources**. As a specialisation of **Anchors**, we have a sub-collection **Shapes** which may in turn be a simple shape or a complex shape (**CompShapes**) consisting of two or more shapes.

We have also expanded the modelling of **Resources**. First note that we now have a **RelatedTo** association between resources and this can be used to increase the link path length to more than 1. This means that the user can follow further links to other related multimedia material. Second, we have categorised **Resources** as **Texts**, **Images** etc. and allow for complex resources which may contain other resources. For example, it is possible that a complex image contains other images as indicated by the example in Fig. 4, or we may want to activate parts of an image. For instance, we may be interested in detailed information about a certain part of the body of a fish such as the fin, or a word concept printed within an image. It is also possible that areas may overlap as arises when a printed text flows over an image. This requires a further extension to the information model to support a layering scheme as shown in Fig. 6.

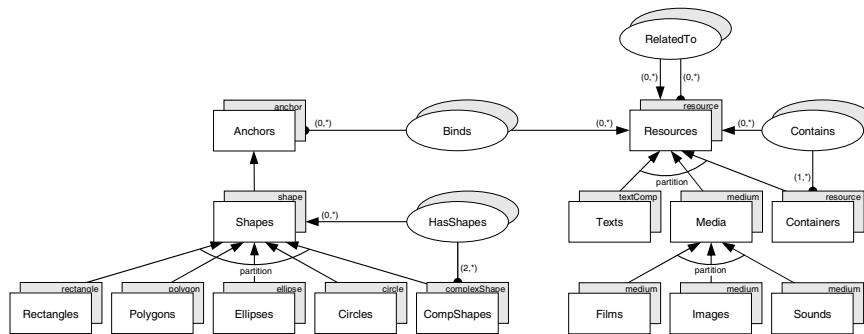


Fig. 5. Linking areas to digital resources

Each shape is associated with a specific logical page layer and there can be no overlapping shapes within a layer. Thus overlapping shapes must belong to different layers and in the case that the selected position lies within two or more shapes, the shape within the uppermost layer is selected. In addition, specific layers may be activated and deactivated and this enables us to generate context dependent results by binding a particular position on a page to different resources according to the current set of active layers. For example, this could be used to provide a “zoom in” effect on images if we were to define the shapes and layers such that repeated selection of a position on a page caused the uppermost layers to be deactivated in turn, thereby moving down to smaller image parts defined on the lower levels. The zoom in/out functionality enables a whole range of new interaction and user interface possibilities in working with static paper, where the digital media can be used to give feedback about the context (current layer) within an image.




Fig. 6. Layers and shapes

To model this, we represent the layers as a totally ordered collection **|Layers|** with a subcollection **ActiveLayers**. Note that support for collections with different behaviours (sets, bags, rankings and sequences) in OM and its associated systems is another motivation for using this model. The association **|OnLayer|** relates each member of **Shapes** to exactly one member of collection **|Layers|**

and we assume that it is not only totally ordered, but that this ordering respects the layer ordering of `|Layers|`. Actually it is possible to derive this ordering automatically within the system, but we model it explicitly here for the sake of simplicity.

OM is object-oriented and has a full algebra over objects, collections and associations. Without going into detail of the OM query language AQL, we show below a query that, given a point (x, y) , will return the uppermost shape of the active layer set that contains that point.

```

first(domain((OnLayer rr ActiveLayers)
           dr (all S in Shapes having S.contains(x,y))
           )
)

```

First, two selections are applied to the association `OnLayer`: we select pairs where the second element belongs to `ActiveLayers` and, in a second step, we further select those pairs where the first element is a member of `Shapes` where the method `contains` returns `true` for the input parameters x and y . Having carried out this reduction of the `OnLayer` association, we then take the `domain` which gives the collection of shapes containing the point (x, y) and belonging to an active layer. Since `OnLayer` is a total ordering that respects the ordering of `Layers`, the `domain` operation will maintain that ordering and `first` will select the uppermost shape as required. Details of the algebra and query language are given in [24].

So far we have achieved a mapping of elements of a printed document to digital resources, but we still do not have the capability to perform an inverse mapping from the digital world to the printed world. The problem is that we have insufficient information about the context of an embedded link within a document. It is not enough to know the coordinates within a page that define an active area. We must also know the page and the document that contains that area. We therefore extend the model with document and page information as shown in Fig. 7.




Fig. 7. Context of links within documents

While this is simple to model, we must think about what this means in terms of the underlying technologies. The encoding system must be amended to include information about the document and page number. With current

encoding schemes, it would be difficult to include globally unique document IDs within a grid-based page position encoding scheme. Within the Paper⁺⁺ project, we have encoded page and coordinate information within the page encoding scheme and handle document IDs separately through either the use of RFID tags or special printed document encodings on the cover page.

With the additional contextual information of document and page, we are able to map digital resources to active areas within documents. For example, we could easily find where a given video or term (printed word or phrase) is used within a document set. This now gives us a basic mapping back and forth between printed and digital media. Note that these back references to printed documents are not just another piece of text within the database giving, for example, a document title and a page number. In fact, the links back to physical paper are built dynamically based on meta information about other documents bound to the same information object. This meta information is already present in the information model (bidirectional associations) and therefore no additional authoring is required. In the following sections, we shall see how this can be extended with the use of derived links and link sharing.

4 Enriched Information Environments

At this point we have achieved a relatively simple two-way linking of printed and digital resources. But can we really say that we have achieved an integration of printed and digital information? The answer is “no” since there is no information about the semantic content of these resources. On the printed side, our objects of interest are simply areas of a page. On the digital side, we simply have content such as texts, images, audio and video without these being related to any conceptual information about the application domain such as oceanography in the case of the Blue Planet example.

To provide a richer information environment, we could introduce information concepts into the server by extending the schema to cover concepts about the application domain as shown in Fig. 8.

We have now partitioned **Resources** into **Contents** to represent the basic text and media files as before and a new collection **Concepts** to cover the representation of application specific information. Essentially, **Concepts** is a way of integrating an application database into the information server of the cross-media publishing framework.

The application concepts part of the schema shown in Fig. 8 actually more or less corresponds to the information currently available in the facts file on the Blue Planet website. Users can select a classification and then one of a list of species in that classification and a simple fact file with basic information about the species including physical description, statistics, habitat, diet and distribution will be displayed. There are no links within this information to maps, images etc. or to related species such as predators. Since OMS supports attributes of type **mime**, it would be an easy matter to include links to multimedia files within the species objects and in this way already achieve a more highly-linked system than

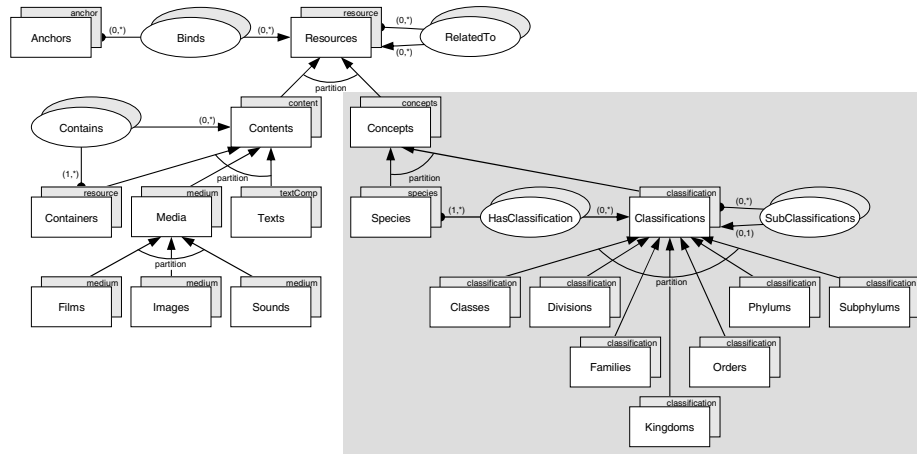


Fig. 8. Adding concepts as resources

is currently available on the Blue Planet website, as well as having links between the printed and digital material. Note that we have extended the basic schema to include general forms of classifications that appear in biological databases.

Extending the application part of the schema to include further information such as geographical locations and with appropriate associations as shown in Fig. 9 would clearly further enrich the information environment.

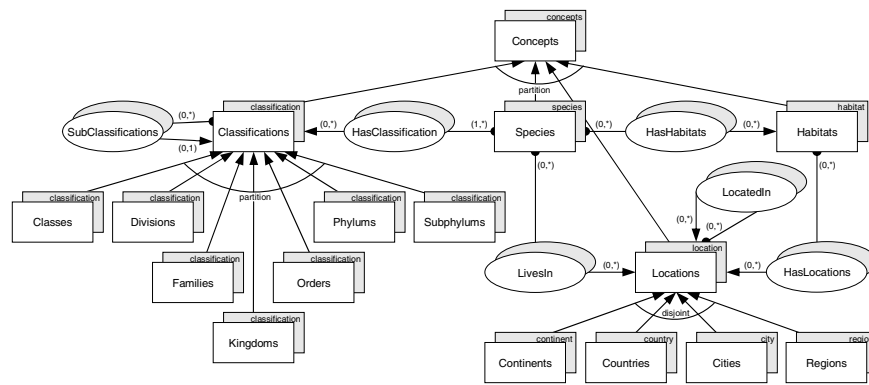


Fig. 9. Extended application model

While the introduction of an application database can potentially provide readers of printed material with access to a wealth of related digital materials, like all database development projects, it needs significant investment to do it

well. Developing a full application database requires domain expertise and lots of data input. Therefore this may not be the solution for all forms of cross-media publishing. However, it is important to point out the potential long-term rewards of such an investment, especially when one considers reusability. Even within the Paper⁺⁺ project, we found that we could reuse all or part of application models. As mentioned previously, our first demonstrator application was based on a nature encyclopaedia [8]. When we came to examine the Blue Planet material, we found of course that the core application schema remained much the same and the difference lay in the data.

The BBC are well-known for the high standard of their nature documentary series and over the years there have been many. Such a nature database could be used for all of these programmes and, of course, would provide a rich source of material that the BBC itself could use in the planning and development of their programmes. In fact, according to our sources, they are in the process of developing such a database and this could provide a valuable source of material if integrated into their cross-media publishing activities in the future.

With the introduction of an application schema which effectively links different concepts of the application domain, note that we can also introduce more powerful forms of derived linking of printed materials. Assume that there is a link from an element of one document to a particular species object. Further assume a link from a second document such as an atlas to a location object. Then through the associations `HasHabitats` and `HasLocations`, a user interested in that species and its habitat can be directed to one or more places within the atlas. Derived linking may be particularly powerful in the case of open authoring systems where users may dynamically add their own links from printed documents to digital resources and links can be shared within user communities.

This leads us to the final aspects that we want to consider here, namely personalisation and customisation. We have already alluded to the use of layering as a potential means of presenting context-dependent information through the dynamic activation and deactivation of layers. In addition, we may want to present different materials to different users according to preference or role. For example, in the case of Blue Planet material, we may want to present different links and content dependent on the age of the user. In the case of young children, we may want to link to less detailed textual information and rather to more images and audio and to link to quizzes appropriate to their age. For educational scenarios, we may want to present different information according to whether the user is an instructor or a student. In particular, links from an exercise sheet or quiz may link to hints in the case of a student and detailed solutions in the case of an instructor.

For this reason, it is important to model users within the system and for that we introduce into the Paper⁺⁺ server a special user concept. Note however that the user information may be shared across applications and hence we show it as a separate component in Fig. 10.

Here we show the main components of the information server that we developed for the nature encyclopaedia application. For each application domain,

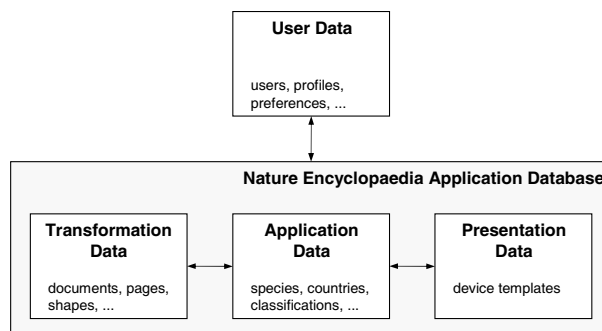


Fig. 10. Information server

there is a *Transformation Component* that defines the active areas and their binding to resources, an *Application Component* that is the domain-specific database and also a *Presentation Component* which manages templates to enable client-dependent delivery of content. For example, we have implemented this application for access from a range of devices including standard desktop browsers and PDAs as described in [15].

The *User Component* may be shared by a number of applications. Note however that it is also possible as discussed above that the *Application Component* could be shared by applications.

5 Authoring and Publishing Tools

When considering the activity of authoring cross-media material, there are a number of fundamental factors to take into account. These factors deal with issues such as the intended usage of the system and the source of materials and they will greatly influence the types of authoring to be supported and hence the approaches and tools to be used. While it is beyond the scope of this paper to describe any particular authoring tool or publishing framework in detail, we want at least to indicate some of the approaches that we are investigating.

The first factor to consider is whether the content is already available in the form of digital and printed materials. If so, the main authoring activity is *link authoring* and a tool is required to support the creation of links between the existing content elements. The second factor is then whether only the publishers can author links or also the users of the material in which case *dynamic link authoring* must be supported. If users can author their own links, the next question is whether there is an *open link authoring* scheme, in which users can link arbitrary materials and not only those provided by a single publisher. For example, a student may wish to link concepts in a set of lecture notes to text books or to websites. Clearly, the ability to freely create links between arbitrary printed materials implies a major shift to a future world of interactive paper.

All link authoring tools require some means of linking positions within a document to resources in the information server. Two problems arise. How can links be marked so that the user is aware of them and how can we ensure the stability of the association between document positions and elements of the document? Any changes to document content will alter the physical position of elements. Dynamic authoring of links by users across existing printed documents is technically simpler in that the document is static. However, in this case, a solution needs to be found to inform the users of link positions. One possibility is to consider combining a highlighter pen with the reader device as described in [2].

In the case of a publisher performing link authoring, access to the document rendering system is required to ensure that the association between document elements and physical page coordinates will be maintained.

First experiences within the Paper⁺⁺ project have shown that existing authoring “solutions” tend to be restricted in flexibility (e.g. by handling only rectangular shapes) and often lack support for context-sensitive information such as user-dependent anchors or personalised links [21]. The authoring tool we are developing is flexible by supporting any kind of shapes and multi-layer authoring for complex figures. Further, it enables the authoring of multi-user documents. These multi-user documents contain anchors (shapes) linking to different resources based on the user currently working with the document. Finally, we are investigating a better integration of the reading device into the overall authoring process. The user should not only be able to define the active page areas (anchors) by marking them on screen but rather generating them directly on the physical paper by using the reader to draw a shape’s outline.

The explicit authoring of links can certainly be time-consuming and it may be based on the knowledge or experience of a single individual. Especially in the case of open, dynamic link authoring systems, it is desirable to share links among communities of users, thereby enabling users to benefit from the knowledge of others. In addition, user profiling and access paths can be used to generate virtual trails, suggesting links to users based either on their own or the community’s experiences similar to the ideas described in [4]. These are topics that we have partly addressed in the context of other projects (e.g. a personal assistant for access to web databases [20]) and we are currently integrating these ideas into the Paper⁺⁺ information server.

So far, we have considered only link authoring. If the content does not already exist, then the authoring activity will consist of both *content authoring* and *link authoring*. In this case, the content can be developed with the resulting hypermedia system in mind and tools can be used that generate both printed and digital documents along with the links between them. In such a case, instead of using traditional authoring tools such as a word processor, a content management approach of the form now prevalent in website engineering could be employed. For example, within our research group, we are developing flexible and powerful content management solutions that manage content, structure, layout and view objects and dynamically create device and context dependent

delivery of information using an XML/XSLT framework [9]. We are now exploiting these systems to generate, not only digital documents such as HTML for web browsers, but also printed documents with embedded active links.

We have briefly outlined the various forms of authoring activity that need to be supported and indicated the sorts of tools and technologies that are required. However, it is clear that, in many cases, the different forms may all be required within a single application domain. For example, in a university learning environment, the instructors may want to generate on-line and printed material which is linked together, but additionally students will want to add and share links both within this set of materials and also to external sources.

6 Concluding Remarks

A study of the literature on digitally augmented paper, reveals that the “information” aspect often tends to be ignored or underplayed. This shows itself in terms of the very basic, over-simplistic models that are often used on the server side to implement the mapping between the printed documents and digital resources. However, the information model of the server has a major role to play in facilitating the integration of printed and digital materials.

We have presented a case study in cross-media publishing to show how a range of extensions to the information model of the server can lead to more flexible, customisable systems that in turn yield much richer information environments. This provides the potential for paper to be integrated as a first-class medium into hypermedia systems.

The information model presented forms the basis of the server that has been implemented for interactive paper in the context of the European project Paper⁺⁺. Using the server and various prototypes of the reader and printing technologies, we have already been able to demonstrate the potential uses of interactive paper and to carry out preliminary user studies. We are now in the process of developing various authoring tools both for open, dynamically linked systems and also for major content providers.

Acknowledgements

The work reported in this paper is part of the European project Paper⁺⁺, under the Disappearing Computer Programme (IST-2000-26130). We thank the other members of the project for their contributions and feedback on our work.

References

1. Anoto AB, <http://www.anoto.com>.
2. T. Arai, D. Aust, and S. Hudson. PaperLink: A Technique for Hyperlinking from Real Paper to Electronic Content. In *Proceedings of ACM CHI'97, Conference on Human Factors in Computing Systems*, March 1997.

3. The Blue Planet Website, BBC, <http://www.bbc.co.uk/nature/blueplanet/>.
4. V. Bush. As We May Think. *Atlantic Monthly*, July 1945.
5. PaperLink, code-it AG, <http://www.code-it.ch/>.
6. CueCat, DigitalConvergence, <http://www.crq.com/>.
7. M. Dymetman and M. Copperman. Intelligent Paper. In *Proceedings of EP'98, 7th International Conference on Electronic Publishing, Document Manipulation, and Typography*, April 1998.
8. D. Frohlich, E. Tallyn, N. Linketscher, B. Signer, and G. Adams. Reading Augmented Paper: Children's Experiences from a Simulation Study. Technical Report HPL-2001-308, HP Labs, 2001.
9. M. Grossniklaus and M. C. Norrie. Information Concepts for Content Management. In *Proceedings of DASWIS 2002, International Workshop on Data Semantics in Web Information Systems*, Singapore, Republic of Singapore, December 2002.
10. D. L. Hecht. Printed Embedded Data Graphical User Interfaces. *IEEE Computer*, March 2001.
11. C. Kafka. DataGlyphs Bridge Paper and Digital Worlds. *Docu World*, 1998.
12. A. Kobler, M. C. Norrie, B. Signer, and M. Grossniklaus. OMS Java: Providing Information, Storage and Access Abstractions in an Object-Oriented Framework. In *Proceedings of OOIS'2001, 7th International Conference on Object-Oriented Information Systems*, Calgary, Canada, August 2001.
13. P. Ljungstrand, J. Redström, and L. E. Holmquist. WebStickers: Using Physical Tokens to Access, Manage and Share Bookmarks to the Web. In *Proceedings of DARE'2000, Designing Augmented Reality Environments*, Elsinore, Denmark, April 2000.
14. M. C. Norrie. An Extended Entity-Relationship Approach to Data Management in Object-Oriented Systems. In *Proceedings of ER'93, 12th International Conference on the Entity-Relationship Approach*, Arlington, USA, December 1993.
15. M. C. Norrie and B. Signer. Web-Based Integration of Printed and Digital Information. In *Proceedings of DIWeb'02, 2nd Workshop on Data Integration over the Web*, Toronto, Canada, May 2002.
16. M. C. Norrie, A. Steiner, A. Würigler, and M. Wunderli. A Model for Classification Structures with Evolution Control. In *Proceedings of ER'96, 15th International Conference on Conceptual Modelling*, Cottbus, Germany, October 1996.
17. The Open University, <http://www.open.ac.uk/>.
18. Paper⁺⁺, IST-2000-26130, <http://www.paperplusplus.net>.
19. A. J. Sellen and R. Harper. *The Myth of the Paperless Office*. MIT Press, November 2001.
20. B. Signer, A. Erni, and M. C. Norrie. A Personal Assistant for Web Database Caching. In *Proceedings of CAiSE'2000, 12th International Conference on Advanced Information Systems Engineering*, Stockholm, Sweden, June 2000.
21. I. Sio, T. Masui, and K. Fukuchi. Real-world Interaction using the FieldMouse. In *Proceedings of UIST'99, 12th Annual ACM Symposium on User Interface Software and Technology*, November 1999.
22. NeoMedia Technologies. IDOCs Linking the Worlds of Print and Electronic Media. Technical report, NeoMedia Technologies, September 1998.
23. S. Whittaker and J. Hirschberg. The Character, Value, and Management of Personal Paper Archives. *ACM Transactions on Computer-Human Interaction*, June 2001.
24. A. Würigler. *OMS Development Framework : Rapid Prototyping for Object-Oriented Databases*. PhD thesis, ETH Zurich, 2000.