# Switching over to Paper: A New Web Channel

Moira C. Norrie
Institute for Information Systems
ETH Zurich
CH-8092 Zurich, Switzerland
norrie@inf.ethz.ch

Beat Signer
Institute for Information Systems
ETH Zurich
CH-8092 Zurich, Switzerland
signer@inf.ethz.ch

## Abstract

*We present a general web-based information infrastructure capable of supporting the rapid development of highly-interactive information environments that cater for widely varying requirements across application domains and all forms of fixed and mobile client devices. In particular, we describe how this infrastructure has been extended to support digitally augmented paper through a special transformation component that can map active areas of document pages to information objects so that user- and context-dependent interaction can be supported. Our infrastructure is sufficiently general and flexible to adapt to, not only emerging and even unanticipated technologies in the area of interactive paper, but also the rapidly expanding interaction sphere of hypermedia.*

## 1 Introduction

Many researchers have argued for the retention of paper as a key information medium and its integration into cross-media environments as opposed to its replacement [15, 32]. This has resulted in a wide variety of projects and technological developments over the past decade. Some of these such as the Logitech io Personal Digital Pen [14] have now reached the marketplace. However, it is important to note that not all of these technologies have the same goal in mind and hence the same functionality. Many focus on the issue of information capture rather than interaction by enabling either pre-printed or hand-written information to be transferred to the digital world.

As part of the European project Paper$^{++}$ [29], under the Disappearing Computer Programme (IST-2000-26130), we are investigating the augmentation of physical paper. Paper has a number of advantages over digital media in terms of navigation within and across documents. It is much simpler to scan a book by rapidly flicking through pages than to browse through a digital document [27]. Furthermore, various forms of paper-based collaboration and interaction are very difficult to support in today's digital environments [19]. On the other hand, the digital medium allows for flexible updates of content, indexing and search facilities.

Our goal therefore is not to replace paper by digital gadgets, trying to mimic some properties of physical paper such as, for example, Microsoft's OneNote in combination with Tablet PCs [28]. Instead, we are developing a cross-media information management infrastructure to *fully integrate paper and digital media*, thereby gaining the best of the physical and the digital worlds. We aim not only to provide links between printed and digital information, but also to build highly-interactive systems enabling users to easily switch back and forth between paper and digital channels [25].

With the rapid development of technologies in the field of input and output devices, it becomes important to have information infrastructures that are sufficiently general and flexible to adapt to, not only emerging (and perhaps even unanticipated) technologies, but also the rapidly expanding interaction sphere of hypermedia. It is no longer sufficient to think only of desktop browsers as the output channel. A modern information infrastructure has to be able to cater for all forms of mobile devices, user preferences and context-dependent delivery. It is a matter of *delivering the right information, to the right person, at the right time*.

Our overall goal is to develop such an information infrastructure capable of supporting highly-interactive information environments. As part of this, we believe that it is important to support paper as one form of client device. This includes, therefore, not only the printing of information on demand, but also support for its linking to digital information. In this paper, we present our web-based link infrastructure and, specifically, its support for interactive paper. Further, we present our authoring tool for digitally augmented paper allowing us to define active areas on paper and link them to digital or physical objects. In addition, the Paper$^{++}$ authoring tool provides functionality for user management to handle link visibility and a component to manage multi-layering and overlapping links.

We begin in Section 2, with an overview of state-of-the-art information infrastructures that can support hypermedia systems. Section 3 then goes on to discuss the specific requirements of interactive paper in terms of information infrastructure. In Section 4 we describe a general model for interactive paper based on active areas and layers. Some applications developed based on our infrastructure for digitally augmented paper are presented in Section 5. Section 6 deals with the issues of cross-media link generation and the required authoring tools. Concluding remarks are given in Section 7.

## 2 Information infrastructures

Global and mobile access to information is of increasing importance. As a result, modern database systems are no longer simply static repositories of application data, but rather central system components that manage information about users, roles, contexts and presentations as well as the application domain.

Within our group, we are investigating technologies and methodologies to support the rapid development of information environments that offer access to shared information from various forms of client devices and employing different modes of interaction. In particular, we are exploiting object-oriented, XML and web technologies to achieve general and flexible solutions that can adapt to a constantly evolving world of new technologies and requirements.

We have extended the OMS Java data management framework [17] with a general access layer to support universal client access through a clear separation of the notions of information, structure and presentation [18, 34]. The access layer is based around an XML server and the use of XSLT for document generation from default and customised templates. Included in the current set of devices are desktop browsers, PDAs, mobile phones (WAP/WML and iMode), and regular voice phones through the use of VoiceXML [35, 36]. In addition, as described in the next section, we have further extended the framework with support for access from digitally augmented paper.

The resulting eXtensible Information Management Architecture (XIMA) shown in Figure 1 can provide multi-channel content delivery. However, for the generation of hypermedia and support for multi-lingual, multi-format, personalised and context-dependent delivery of information, it is necessary to also integrate the key concepts of content management [8].

A wide variety of software solutions classified generally as content management systems have been developed to support both the development and operation of complex websites. These systems emerge from the publishing com-

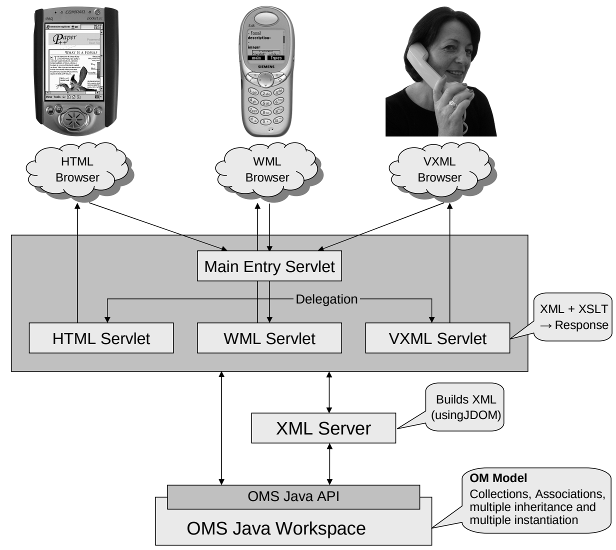


**Figure 1. XIMA architecture**

munity rather than the information systems community and they focus on supporting a team of individuals — authors, editors, graphical designers, programmers and administrators — working together to develop and maintain a website. Typically, the concepts in which they deal at the information level are primarily document concepts such as text, image and link rather than with semantic concepts of a particular application domain. Thus they deal with an image of a given format, but have no information as to what it is an image of. The same is true of the many commercial solutions available to support both the development and operation of complex web sites e.g. Coldfusion [9], Interwoven [13] and Vignette [39]. This lack of semantic information about content makes it difficult to ensure consistency and to link pieces of content together that represent the same logical entity. For example, if information about a person appears on several pages of a website, the content editor has to manually keep track of the fact that the associated content actually belongs together as a semantic unit. This means that when information about that person is updated, or even the entire person is to be replaced, it is difficult to track where to make the necessary changes and tedious to make them. If the information is stored together in one place and the different instances that appear on the various web pages dynamically generated from that information, then it is easy to make changes and avoid inconsistencies. Note that this solution applies even if the actual instances differ greatly, for example, a link to a person based on their name, or a full description of a person inclusive of an image.

The situation becomes much more critical when one starts to consider multi-format and multi-lingual content. In this case, we do not simply want to have the concept of a person and their associated properties, but also different content for these property values so that the appropriate content is delivered depending on factors such as access device and user preference.

We therefore integrate both perspectives into our information infrastructure and manage, not only content, but also representations of application entities. This enables the hypermedia designer to work in terms of high-level concepts such as persons and departments, specifying which concepts appear together in a page and the navigation between pages, rather than being concerned with the details of content and its format as well as layout and presentation. In addition, to support personalised and context-dependent interaction and ensure security of access, it is important to integrate support for the modelling of users, their roles and context.

This requires a very flexible data management framework that can dynamically construct information presentations from entity representations, associations between entities, pieces of content, document structures and templates at access time. Such flexibility comes at a price in terms of response times, and therefore multi-level caching schemes are vital to ensure acceptable performance [33].

To enable semantic entities to be represented and manipulated directly, while retaining maximum flexibility, the framework is based on an object-oriented information model with a separate association construct for representing bidirectional links between objects and uniform handling of both data and metadata. The model includes a full operational model that enables objects, collections of objects and associations to be manipulated directly.

It is important to emphasise that it is also possible to generate entire static websites off-line with such an infrastructure. In particular, and relevant to our work on paper interfaces, it is possible to generate printed documents through the use of appropriate XSLT templates.

## 3 Architecture for interactive paper

The 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, sending a request to a server. The server then processes the request and delivers the required information to the appropriate display device. A number of technical solutions have been proposed and it is beyond the scope of this paper to discuss these in detail. We would recommend the following references for discussions of related technologies [2, 3, 6, 10, 12, 16, 20, 38].

Given these technologies, the key to full integration of paper into hypermedia systems will depend on the flexibility of the information server which is responsible for mapping links embedded in paper to information objects and de-

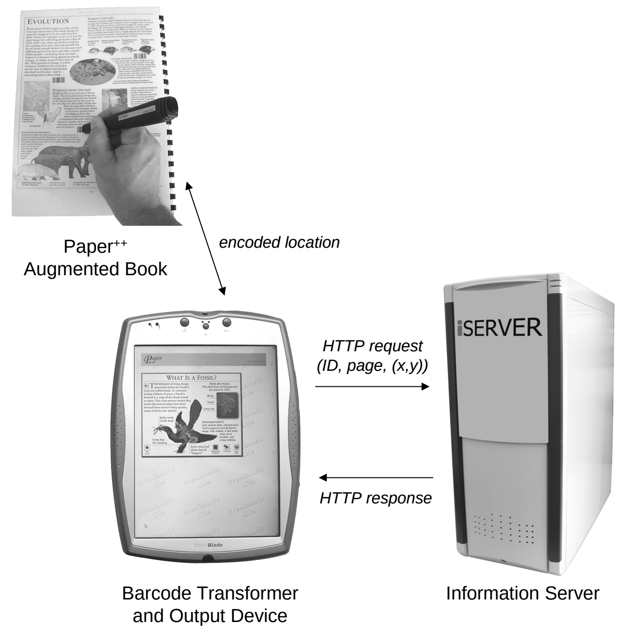


**Figure 2. Paper$^{++}$ architecture**

livering that information in the format required by the client device.

Our Paper$^{++}$ information infrastructure is based on a classic client server architecture as shown in Figure 2. Every page is covered by a grid of invisible conductive barcodes encoding the page number and the (x,y) position. A special barcode reader enables us to detect the (x,y) position within a paper document. At the moment, these paper documents have to be produced in an additional preprocessing step, but the vision is that we will be able to print the conductive grid and the corresponding artwork in a single step (e.g. using an additional ink jet cartridge containing the conductive ink). The document identifiers are handled separately and not encoded within each document page. Instead, we can use radio frequency identification (RFID) tags to identify different documents as was done in [3]. The inductive barcode reader immediately transmits the encoded location and page number information to a barcode transformer component on the client device. The barcode transformer decodes relevant information such as the page number and the (x,y) position from the received stream of ASCII characters. In the next step, an HTTP request containing this information about the pointing device's current position on a specific paper document is sent to the information server which will process the request and return the result in the format requested by the digital presentation device (e.g. HTML for a regular web browser or WML for WAP-enabled mobile phones).

The information server contains the three main components presented in Figure 3. In the remaining part of this

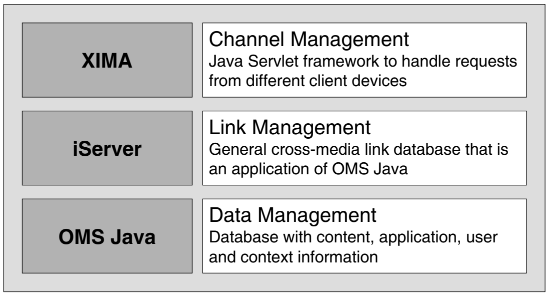


**Figure 3. Information server components**

section, we describe the functionality of these three components and show how they are related.

The OMS Java data management component is used by the iServer as well as by the XIMA component. First, it manages the specific application databases containing all sorts of information about the domain of discourse which then can be used as link resources by the iServer component. All of this application data is already highly interlinked due to the fact that OMS Java is an implementation of the OM model [24] which supports the powerful concept of bidirectional associations to link single information objects. Further, the OMS Java component stores information about user preferences and context information used by the XIMA component for customised output generation. Last, but not least, the iServer architecture is an application of OMS Java.

Based on work in the hypertext community we have developed iServer, a general cross-media link management component enabling links to various kinds of resources. Every bidirectional associative or referential link is bound to one or more anchors and points to one or more target components (*multi-headed links*). Furthermore, the iServer does not distinguish between components which can be used as link anchors and those available as target objects. Every object that can be used as an anchor also is a valid target component and vice versa. We can either link from/to *entire resources* or use the concept of *selectors* to control the link granularity by addressing specific parts of a resource. The sharing of link knowledge is handled by the iServers's user management component allowing access rights to be defined on the link level.

The iServer architecture is designed as a platform which can be extended to support new kinds of media types. As will be shown in Section 4, we have extended the selectors and resources — resulting in the concepts of *shapes* and *pages* — to support interactive paper as a new hypermedia resource.

Finally, the XIMA component is responsible for the deployment of all information managed by OMS Java and the iServer component on different output channels. As described in Section 2, XIMA uses XML as an intermediate data representation and applies XSLT stylesheets for the final output rendering. XIMA not only supports a variety of client devices, but also enables customised output generation based on users' preferences managed by OMS Java.

Various Open Hypermedia projects [11, 30, 40, 41] provide concepts similar to those of the selectors, external links etc. introduced by our iServer link management component. However, most of these systems focus on the integration of different digital media types and do not pay much attention to augmentation of physical objects by supplemental digital content.

A number of previous projects have used some form of barcode as active links. The barcode is a simple encoding of a unique ID and a simple mapping table on the server could be used to map this ID to a multimedia file (e.g. a video), which is then displayed. Such a solution has many limitations. It supports only unidirectional, single level linking. There is no possibility to browse further in the digital media and there is no way of linking from, say a digital version of an image, back to all the documents linking to this image.

An improved solution is to activate areas of a page and map these to information objects, rather than directly to media files. These objects can store metadata about files and, through associations to other information objects, can support deeper levels of linking. For example, assume we want to link the occurrence of the word antelope within a page of a nature book to an image of an antelope. Instead of mapping directly to the *antelope.jpg* file, we map to an instance of objects of type animal that specifies information such as the species name, average size, description and habitat. Not only can we display all (or some subset) of this information in addition to the image, but also links to associated information objects such as predator species or locations where they can be found. In other words, we can generate full hyperlinked web documents from the database as described in Section 2. It is important to note that with this general solution, we can support all forms of linking between active areas and digital resources — whether these are simple media files or hypermedia documents as shown in Figure 4.

By linking paper into a much richer information environment with information about application concepts and not simply media content, we have the potential to provide readers of printed material with a wealth of related digital materials, inclusive of dynamically generated hypermedia documents designed to meet the needs of that reader at that point in time and space. However, it does come at a price since it requires the development of application-specific solutions in terms of a database with information about the current domain of discourse. This requires domain expertise and large-scale investment. However, it is important to point out the potential long-term rewards of such an investment when one considers reusability.

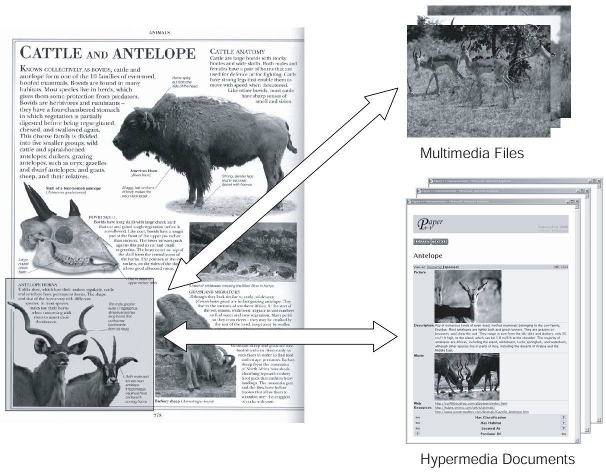


**Figure 4. Active areas**

An example of a cross-media publisher is the BBC (British Broadcasting Corporation). They are well known for their nature documentary series and not only sell these to other broadcasters world-wide, but also some material is used in various educational establishments as a basis for long-distance learning (e.g. the Open University in the UK has an oceanography course based on the Blue Planet materials). A nature database with information on species and their classifications, habitats, geographical locations etc. 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 addition, they have repositories of multimedia that they have built up over the years. By linking these together, and using the sort of infrastructure that we present in the next section, they could support their cross-media publishing for many future series [26].

As another example, consider a system to support researchers in their annotations, recommendations and cross-referencing of different articles. What is the domain of discourse? There are two possible answers — the specific research domain, e.g. digitally augmented paper, or the general research activity of literature search and survey. Clearly, the former is much more specific and a system to support it would require a database about digitally augmented paper that represents concepts such as for example DataGlyph [16]. The second is much more general and only requires a database that knows about concepts such as citations, references, annotations etc. which enable linking between different articles and information about authors of those articles.

Generally, the requirements of any system will depend, not only on the application domain, but also on the envisaged sphere of interaction and scope of information. It is therefore essential that the underlying information structure be able to support all forms of system — including the basic linking of printed documents to multimedia files as well as advanced specialist information environments that freely link back and forth between the digital and paper worlds.

## 4 Active areas and layers

As mentioned already, we map active areas of a page to digital resources, which represent concepts of the domain of discourse. An active area can be defined as one of a number of simple geometrical shapes or as a complex shape composed from other shapes.

It is possible that we may want overlapping active areas (similar to the overlapping links described in [23]). For example, a piece of text may overlap an image. Complex images may contain sub-images that we also want to be able to activate. Similarly, it is often the case that we may want to link individual words or phrases within a piece of text as well as the entire piece of text. To enable this, we have introduced a virtual page layering scheme as shown in Figure 5.

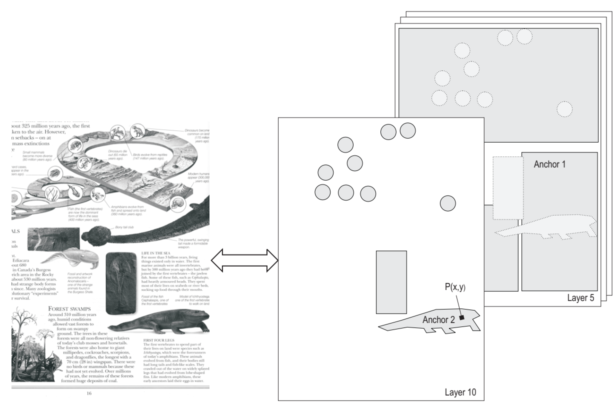


**Figure 5. Virtual page layers**

Each active area is associated with a specific logical page layer and there can be no overlapping areas within a layer. We assume the reader device delivers a pair of (x,y) coordinates that indicate a position within the page along with the page ID. From the (x,y) coordinates, we can identify all the active areas to which the point belongs. If it belongs to more than one active area, then these areas must be assigned to different layers and the one from 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 information objects depending on the current active layer set. This could, for example, be used to provide a "zoom-in" effect on images if the areas and layers were defined in such a way that repeated selection of a position on a page caused the uppermost layers to be deactivated in turn,

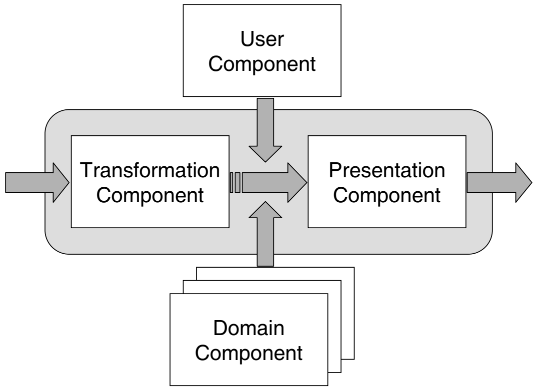


**Figure 6. Underlying database components**

moving down to smaller image areas defined on lower layers. Virtual page layers and the potential to support zooming provides a whole new set of interaction possibilities for paper interfaces, where digital media can be used to give information about the current context.

In addition to the layering concept, link activation may depend on the user and even the user role. A special user component manages information about the different users of the system. A user can define their preferences and also be classified into various user groups and roles. This enables an author of an application to, not only define different active areas for different users, but also to map the same active area to different information objects based on the user profile. Such forms of user-dependent information delivery could be used in an educational environment where an instructor obtains different information from the students. For example, the instructor may receive the solutions for exercises included in printed course material, whereas the students only receive hints on how to solve the exercise or links to other course materials. The information infrastructure is based around four main database components — the transformation, user, domain and presentation components as shown in Figure 6.

The transformation component deals with documents, pages, active areas and layers and is responsible for mapping positions within a printed document set to information objects. The presentation component is responsible for information delivery to various devices based on document structure definitions and presentation templates. These two components are specific to a particular application.

As discussed above, the user component manages information about user preferences and profiles and can be used to both control access and deliver user-dependent information. This component may be shared by more than one application.

The domain component manages information about the domain of discourse and may also be shared by different applications as discussed in Section 3 for the example of a nature database to support various BBC cross-media publishing activities. Further, more than one domain component could be used in a given application. In the case of the researcher's application also discussed in Section 3, we could use a general domain component dealing with information about publications, citations and authors and a further domain component dealing with concepts about a specific research area.

## 5 Applications

Within the Paper$^{++}$ project, we have developed several demonstrator applications to illustrate the potential of tightly integrating paper and hypermedia infrastructures and carry out some preliminary user studies. At the same time, these demonstrators helped us validate some of the new concepts introduced by our information server for digitally augmented paper.

Some mechanism was required that would enable us to implement these demonstrators in parallel to the design and testing of the paper, printing and reader technologies also being developed within the Paper$^{++}$ project. Our method was to use off-the-shelf barcode readers and conventional visible barcode technologies, but to use them in an unconventional way to simulate the technologies under development. We did this by using barcodes to encode position information rather than simple unique identifiers. For each active area, we created a barcode that encoded a position within that area and positioned it in the document alongside that area. The information delivered by swiping a barcode was therefore equivalent to that to be delivered by the new technologies when any position within an active area is selected.

Based on a children's nature encyclopaedia published by Dorling Kindersley and the corresponding software application delivered on a CD-ROM, we have designed a new Paper$^{++}$ version of the encyclopaedia linking parts of the book to pieces of digital information. Three different versions of the encyclopaedia (book only, book and CD-ROM and the Paper$^{++}$ version) have been compared in a user study where groups of children were given specific tasks to solve using one of the encyclopaedias. Initial results of this evaluation are presented in [7].

A second demonstrator has been built supporting researchers in cross-referencing and annotating publications as described in Section 3. The application domain we have chosen for this demonstrator is the Paper $^{++}$ internal project database providing information about related technologies, publications and contact persons. Some of the publications stored in the database have been augmented in the printed version with links into the database. An example of such a link from a publication to a specific technology is shown in

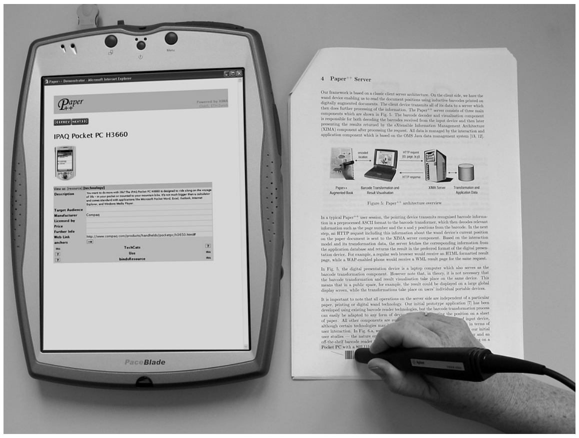


**Figure 7. Augmenting research papers**

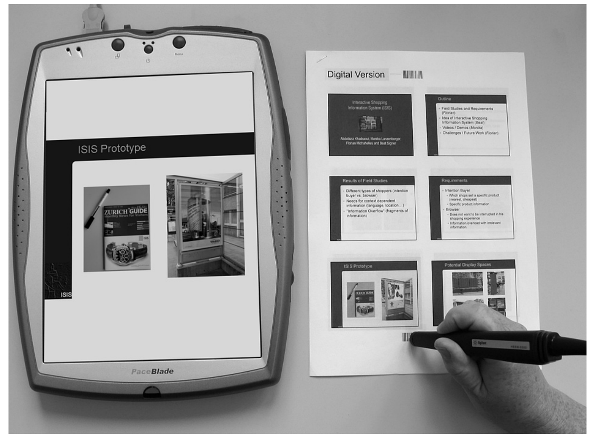


**Figure 8. Paper-driven presentations**

Figure 7. After swiping the highlighted word *Pocket PC* the user immediately gets detailed information about the *Pocket PC* concept introduced in this part of the document.

As mentioned before, a major advantage of our cross-media information management architecture compared to similar projects is the possibility to have links back from the digital medium to paper without any additional authoring effort. Based on the metainformation stored together with every bidirectional link (document ID, page number and shape) the system can automatically generate links from a digital information concept to all the documents referencing it. In the case of the application shown in Figure 7, we defined a method named *anchors* for every digital information unit. If a user executes this method, the system will return a collection of all documents linking to the selected piece of digital information. In the current implementation, those links back to paper are text-based only which means that, for every anchor, the system provides information on which page and within which document it has been defined (we do not use the shape information in the text-only feedback). As part of our ongoing work, we are extending the text-only feedback by using the shape location information to generate a combination of textual and graphical feedback (showing the exact location of an active area within a page).

The links back to paper become even more powerful as soon as multiple users start to use the system to augment papers they are reading. Since the system allows for sharing of link information, a user can profit from other users' work and knowledge by having access to their link spaces. To come back to the scenario presented in Figure 7: a user can get information about other publications in which the *Pocket PC* technology has been mentioned by just using the community's link knowledge to get all links from the digital concept of the *Pocket PC* to related research papers.

Even without any additional digital information used in augmenting parts of a document, the sharing of links is a potent tool somehow extending the very static concept of today's citations used in research papers. For such *paper to paper links* we have extended the user interface by applying XIMA's facility to generate voice output. Instead of returning a text string containing the information about the the citation's title and related information, XIMA transforms this string to spoken output using a text-to-speech engine (TTS). This allows a user to use multiple channels at the same time to absorb information. After swiping a link referencing a paper, the user can go on reading while simultaneously listening to the aural information about the referenced publication.

Another use of paper as an input device is shown in Figure 8. We have digitally enhanced some of the handouts of presentations managed by the project database. The handouts shown on the right hand side provide a user with access to the digital version of the presentation. First of all by swiping the link on the top (title) we get access to some metainformation associated with the presentation such as the location where the presentation was held, the date when the slides were presented and also a digital version of the document in the form of a PDF file. Furthermore, the handouts can be used to control the presentation. The user can specify the slide to be shown next by just swiping the corresponding slide's image in the handouts. The idea of giving flexible paper-driven presentations is similar to the approach described in [22].

## 6 Authoring tools

A general information infrastructure for interactive paper implies a general authoring framework that can support the different forms of authoring that may be required. First of all, we can distinguish between content authoring and link authoring. If we want to link together existing printed and digital materials, then the focus is purely on link authoring.

If however, we want to support a publisher of new cross-media materials, then we need to support the authoring of content as well as links. We will focus first on the former case and then go on to consider the second case where an entire publishing framework needs to be provided.

A link authoring tool requires some means of linking positions within a document to objects in the information server. Two key issues 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 such as images, headings and words?

In a static link authoring system, links between content are fixed and defined by the publisher of the hypermedia system. Users can freely navigate back and forth between the printed and digital worlds, according to the links provided by the developer of the system, but they cannot add their own links. This is, of course, similar to the situation in current standard web browsers. Conventions such as the use of colour to indicate linked pieces of text have been adopted and these could also be used to indicate links in printed documents. Further, users have come to expect where links might be — such as headings and images. If the consistency of links is to be maintained under editing changes to documents, then the authoring tool requires access to the document rendering system so that the link positions will be updated in the new version of the printed document.

The ability for users to create their own links leads to a concept of dynamic link authoring systems. In the case of web browsers, dynamic link authoring is being introduced through the adoption of standards such as XLink [5] and research prototype browsers such as Amaya [1]. The dynamic authoring of links from existing printed documents to digital resources is relatively simple in that the documents are static and no special consideration has to be given to maintaining link consistency under changes. The authoring tool must simply provide a means for the user to specify an active area of a page and the object to which it should be linked. The main issue in this case is how to make the users aware of the existence of such links. One possible solution would be to integrate a highlighter pen with the reader device as described in [2].

The explicit authoring of links either by the publisher or the user is one possible way to build up a rich hypermedia infrastructure. Analogous to ongoing work on adaptable links by the hypertext community [4, 21, 31], we have introduced the concept of derived links. Let us assume that we have two paper documents, document A and document B, as shown in Figure 9. Further assume that from both of them we have authored links to the same application domain database. Just by analysing different users' browsing behaviour, we may conclude that most of the users starting from the rectangular active area in document A, end up following some links corresponding to associations in the domain DB, but finally come back to physical paper, namely to the circular shape defined in document B. This implies that there must be a strong correlation between the concept represented by the *rectangular active area* in document A and the concept represented by the *circular active area* in document B. In such cases, it makes sense to generate a direct paper-to-paper link between the two different concepts based on the user access statistics. These derived links are a first step towards an adaptable and flexible hypermedia infrastructure which automatically evolves over time.

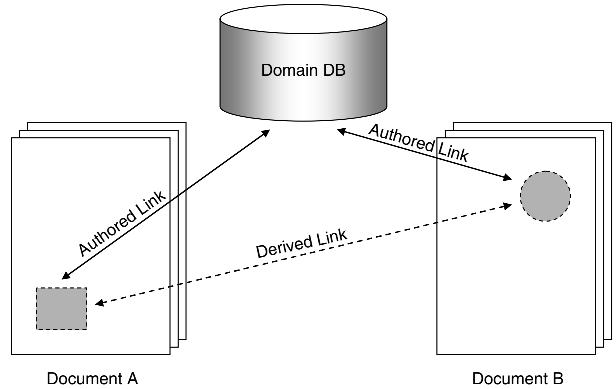


**Figure 9. Derived links**

Authoring tools proposed in other projects have tended to be restrictive, allowing only rectangular active areas to be specified and not supporting context-dependent links and personalisation (as for example in [37]). In contrast, the authoring tool that we are developing supports a number of simple shapes (rectangles, circles, ellipses and polygons) and complex shapes which are a composition of multiple shapes. Further, it introduces the concept of virtual page layers to control the granularity of active areas. Our authoring tool for digitally augmented paper is presented in Figure 10. The main window shows the digital version of a document's page together with all the active areas (shapes) defined for this page. A user can manipulate existing shapes or add new ones and link them to digital information. On the left hand side, there is a navigator component providing easy access to single document pages or to specific shapes, a component to control the layering and a user management tool.

In addition to the authoring tool based on a digital document version, we are developing a paper-based link authoring component. The idea is to use the "reader" device to define the boundaries of active areas directly on paper without the need for a digital document version.

In the case where both printed and digital materials are to be developed from scratch, then a quite different approach could be adopted. In this case, both content and links need


**Figure 10. Paper$^{++}$ authoring tool**

to be authored and an integrated authoring tool could be developed based on recently developed web content management solutions. Instead of using traditional word processing tools to generate printed documents and then linking them to digital information, all content can be managed in a database and a general XML/XSLT framework can be used to generate not only digital documents such as HTML, but also printed documents with embedded active links.

Since we have already integrated concepts for content management into our data management framework as described in Section 2, this approach can be supported with minimal effort. The main issue is how to automate at least part of the task of link generation to make this process less tedious. In the sample applications considered so far, we do this by generating links in one of two ways — either based on explicit associations stored in the database or derived from metainformation such as concept naming so that keywords contained in text are linked automatically. Clearly, such techniques for automatic link generation are also used in the generation of websites and general hypermedia systems.

## 7 Conclusions

We have shown how the eXtensible Information Management Architecture and the iServer link management framework have been extended to cope with links from paper to web-based information and vice-versa, thereby introducing paper as a new form of channel in multi-channel information systems.

The information infrastructure that we have developed has great potential as an experimental platform for the investigation of emerging technologies for interactive paper and its potential uses. By remaining independent of particular hardware solutions and modes of interaction, we are able to easily adapt to both new technologies and applications.

Early demonstrator applications were tedious to develop due to a lack of authoring tools and they tended to exploit only simple linking schemes. However, they were sufficient to carry out early user studies based on printed and digital materials for a children's nature encyclopaedia.

With the development of an interactive authoring tool and the implementation of a fully general link model, we are now in a position to experiment with much more complex demonstrator applications and a range of new technologies and devices. However, the generality and flexibility that we offer has raised even more issues in terms of link visualisation, authoring schemes and the sharing of links among user communities. Further, consideration of various potential application domains has shown that a wide variety of solutions are required to support different kinds of activities and environments.

## References


[1] W3C's Editor/Browser, World Wide Web Consortium (W3C), http://www.w3.org/Amaya/.

[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] M. Back, J. Cohen, R. Gold, S. Harrison, and S. Minneman. Listen Reader: An Electronically Augmented Paper-Based Book. In *Proceedings of ACM CHI 2001, Conference on Human Factors in Computing Systems*, Seattle, USA, March 2001.

[4] C. Bailey, W. Hall, D. E. Millard, and M. J. Weal. Towards Open Adaptive Hypermedia. In *Proceedings of AH 2002, 2nd International Conference on Adaptive Hypermedia and Adaptive Web Based Systems*, May 2002.

[5] S. J. DeRose. XML Linking. *ACM Computing Surveys*, 31(4), December 1999.

[6] M. Dymetman and M. Copperman. Intelligent Paper. In *Proceedings of EP'98, 7th International Conference on Electronic Publishing, Document Manipulation, and Typography*, April 1998.

[7] 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.

[8] 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.

[9] A. Group. ColdFusion MX: Raising the Return on Investment of Internet Application Development. White Paper, April 2002.
[10] F. Guimbretière. Paper Augmented Digital Documents. In *Proceedings of UIST'2003, 16th Annual ACM Symposium on User Interface Software and Technology*, Vancouver, Canada, November 2003.
[11] W. Hall, H. C. Davis, and G. Hutchings. *Rethinking Hypermedia: The Microcosm Approach*. Kluwer Academic Publishers, May 1996.
[12] D. L. Hecht. Printed Embedded Data Graphical User Interfaces. *IEEE Computer*, March 2001.
[13] Interwoven. Application Development using Interwoven : Version 1.1. White Paper, December 2000.
[14] Logitech io Personal Digital Pen, Logitech Inc., http://www.logitech.com.
[15] W. Johnson, H. Jellinek, L. Klotz, Jr., R. Rao, and S. Card. Bridging the Paper and Electronic Worlds: The Paper User Interface. In *Proceedings of ACM INTERCHI'93, Conference on Human Factors in Computing Systems*, April 1993.
[16] C. Kafka. DataGlyphs Bridge Paper and Digital Worlds. *Docu World*, 1998.
[17] A. Kobler and M. C. Norrie. OMS Java: Lessons Learned from Building a Multi-Tier Object Management Framework. In *Proceedings of OOPSLA'99, Workshop on Java and Databases: Persistence Options*, Denver, USA, January 1999.
[18] 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.
[19] D. M. Levy. *Scrolling Forward: Making Sense of Documents in the Digital Age*. Arcade Publishing, October 2001.
[20] W. E. Mackay and A. L. Fayard. Designing Interactive Paper: Lessons from three Augmented Reality Projects. In *Proceedings of IWAR'98, 1st IEEE International Workshop on Augmented Reality*, November 1998.
[21] D. T. Michaelides, D. E. Millard, M. J. Weal, and D. D. Roure. Auld Leaky: A Contextual Open Hypermedia Link Server. In *Proceedings of OHS'7, 7th International Workshop on Open Hypermedia Systems*, August 2001.
[22] L. Nelson, S. Ichimura, E. R. Pedersen, and L. Adams. Palette: A Paper Interface for Giving Presentations. In *Proceedings of ACM CHI'99, Conference on Human Factors in Computing Systems*, May 1999.
[23] T. Nelson. *Literary Machines*. Mindful Press, 1982.
[24] 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.
[25] M. C. Norrie and B. Signer. Web-Based Integration of Printed and Digital Information. In *Proceedings of DIWeb'02, 2nd International Workshop on Data Integration over the Web*, Toronto, Canada, May 2002.
[26] M. C. Norrie and B. Signer. Issues of Information Semantics and Granularity in Cross-Media Publishing. In *Proceedings of CAiSE'2003, 15th International Conference on Advanced Information Systems Engineering*, Klagenfurt/Velden, Austria, June 2003.
[27] K. O'Hara and A. Sellen. A Comparison of Reading Paper and On-Line Documents. In *Proceedings of ACM CHI'97, Conference on Human Factors in Computing Systems*, March 1997.
[28] Microsoft OneNote, Microsoft Corporation, Redmond, USA, http://www.microsoft.com/office/onenote/.
[29] Paper$^{++}$ Project, IST-2000-26130, Disappearing Computer Initiative, http://www.paperplusplus.net.
[30] A. Pearl. Sun's Link Service: A Protocol for Open Linking. In *Proceedings of ACM Hypertext'89, Conference on Hypertext and Hypermedia*, November 1989.
[31] S. Reich, L. Carr, D. D. Roure, and W. Hall. Where have you been from here? Trails in Hypertext Systems. *ACM Computing Surveys*, 31(4), December 1999.
[32] A. J. Sellen and R. Harper. *The Myth of the Paperless Office*. MIT Press, November 2001.
[33] 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.
[34] B. Signer, M. Grossniklaus, and M. C. Norrie. Java Framework for Database-Centric Web Engineering. In *Proceedings of WebE'2001, 4th Workshop on Web Engineering (in conjunction with WWW10, 10th International World Wide Web Conference)*, Hong Kong, May 2001.
[35] B. Signer, M. C. Norrie, P. Geissbuehler, and D. Heiniger. Aural Interfaces to Databases based on VoiceXML. In *Proceedings of VDB6, 6th IFIP Workshop on Visual Database Systems*, Brisbane, Australia, May 2002.
[36] B. Signer, M. C. Norrie, P. Geissbuehler, and D. Heiniger. Telephone Interface for Avalanche Warnings based on Information Server for Adaptable Content Delivery. In *Proceedings of Pervasive 2002, International Conference on Pervasive Computing*, Zurich, Switzerland, August 2002.
[37] I. Siio, 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.
[38] L. J. Stifelman. Augmenting Real-World Objects: A Paper-Based Audio Notebook. In *Proceedings of ACM CHI'96, Conference on Human Factors in Computing Systems*, April 1996.
[39] Vignette. Vignette Content Management Server. White Paper, February 2001.
[40] E. J. Whitehead. An Architectural Model for Application Integration in Open Hypermedia Environments. In *Proceedings of ACM Hypertext'97, 8th Conference on Hypertext and Hypermedia*, Southampton, UK, April 1997.
[41] U. K. Wiil and J. J. Leggett. The HyperDisco Approach to Open Hypermedia Systems. In *Proceedings of ACM Hypertext'96, 7th Conference on Hypertext and Hypermedia*, Washington DC, USA, March 1996.