

# Integrating Paper and Digital Documents

*Heather Brown, University of Kent at Canterbury*

*Peter Robinson, University of Cambridge*

## ABSTRACT

This paper outlines two very different approaches to using physical paper documents as natural input devices to computers. The DigitalDesk is a computer-enhanced desk that uses a video camera to detect and recognise paper documents, while Xerox's 'Intelligent Paper' is specially produced paper containing invisible marks that uniquely identify each sheet. Both allow the visible contents of the paper to be matched to a corresponding digital document. This, in turn, allows natural actions, such as pointing to an item on a page, to be used to initiate computer actions. The overall effect is to blur the conventional distinction between paper and digital documents.

The paper describes the basic mechanisms used in the two approaches and looks briefly at some existing and potential applications. It then looks in more detail at a specific application area: allowing printed documents to act as natural interfaces to the additional information found in digital versions encoded according to the Text Encoding Initiative (TEI) guidelines.

**Keywords:** Paper and Digital Documents, DigitalDesk, IntelligentPaper, Text Encoding Initiative.

## 1 Introduction

Many papers have been published about the relative merits of paper and digital documents, and many hypertext designers have attempted to make their systems easier to use and understand by incorporating familiar paperlike features into the interface [1, 2]. Hypermedia is widely recognised as a marvellous way of finding information and the usability and appearance of digital documents is continually improving [3], but users still prefer to print out documents of more than a few pages because they find reading from paper more convenient and pleasant than reading from a screen.

Several recent projects have attempted to incorporate the advantages and convenience of paper into portable computer systems. A report on one of these, the *The Virtual Book* project [4], gave the following succinct summary of the strengths of paper and digital documents when describing their goals for the design of a portable digital book to compete with paper

The task, then, was two-fold: come as close to paper as possible in the dimensions of legibility, ergonomics, robustness, and similar attributes; then, surpass it in other attributes such as searchability, flexibility, and interactivity.

The resulting prototype, *Lectrice*, has a high quality 10.4 inch diagonal display surrounded by a frame containing several physical buttons. These buttons and a pen are used to manipulate the documents displayed. The device can be attached to a computer network via a physical 'tether' or radio link, or it can be used as a stand alone device. Users found it comfortable to read from a *Lectrice* when sitting back on a couch and it was described as 'an agreeable bed companion: it provides its own light and turns pages almost silently'.

A rather different approach to mimicking the advantages of paper in a portable computer system is shown by the *XLibris Active Reading Machine* [5] and the *Dynomite* portable electronic notebook [6]. These concentrate more on merging 'the benefits of paper note-taking with the organisational capabilities of computers' and use 'digital ink' to mimic standard paper note-taking mechanisms instead of making users change their way of working to suit the computer.

In spite of this awareness of the relative merits of paper and digital documents, and the increasing recognition of the need to bring the advantages of paper into the digital world, little has been done to incorporate actual physical paper into computer systems. For most users, computer-based documents are almost invariably manipulated via a screen and mouse or keyboard and remain essentially separate from paper documents.

This paper reports on two novel ways in which physical paper can become an integral part of an interactive computer system, allowing familiar and convenient printed paper documents to provide access to the powerful facilities associated with digital documents. Section 2 outlines the basic mechanisms used by the DigitalDesk [7] and Xerox's 'Intelligent Paper' [8], explains how they allow paper to become an active part of a computer system, and describes some existing and potential applications.

Section 3 then looks in more detail at a potential application designed to exploit the detailed scholarly information that may be recorded in digital documents coded according to the Text Encoding Initiative (TEI) guidelines [9]. A TEI text may provide information about the state of old manuscripts, include details about the different versions of a work, and give linguistic information or notes about the work provided by scholars. Section 3 shows how a modern printed transcription of a work could act as the starting point for exploring the additional information in a TEI text.

## 2 New technology for integrating paper into the computer interface

### 2.1 *The DigitalDesk*

The Cambridge University Computer Laboratory and the Xerox Research Centre in Cambridge (formerly Xerox EuroPARC) have collaborated for several years on research into the use of video in user interfaces. The approach has been to add computational properties to the conventional office environment rather than replicating the office on a computer screen [10,11]. The original DigitalDesk was developed by Pierre Wellner, a student at Cambridge University sponsored by Xerox [12]. In addition, the collaboration has produced *BrightBoard* [13], a computer enhanced whiteboard, and a *Digital Drawing Board* [14], a version of the DigitalDesk for use in Computer Aided Design.

The DigitalDesk itself consists of an ordinary desk enhanced by

- a conventional computer system linked to a central document registry
- a video camera and projector mounted above the desk pointing down at the desktop.
- a LED-tipped pointing device

Figure 1 shows how the parts of the overall DigitalDesk system are related. Images from the video camera are processed to recognise documents (sheets of paper or the pages of open books) on the desktop and to match them up with corresponding digital documents stored in the central registry. Once this has been done, the computer system can associate positions on the paper pages with the appropriate digital information. User actions can then be initiated *from the page* using the pointing device, and the system can respond via the conventional computer screen or by projecting information down onto the desktop. In particular, it can project information directly onto the page — to highlight certain parts of the text, for example, or to superimpose additional information onto diagrams or graphs. Projecting coloured rectangles onto the page to highlight words or sections of text is an effective way of making the text 'live'. It is similar in many respects to the dynamic use of colour to highlight text on a screen. The overall effect is to involve the paper in the system in a completely new way.

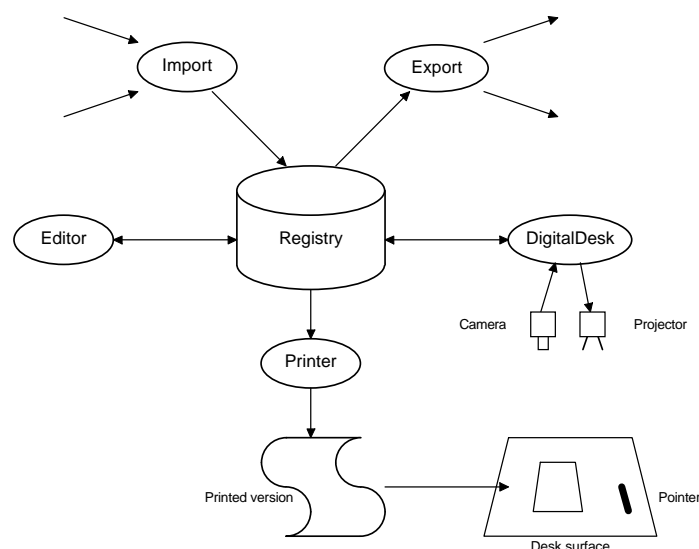


Figure 1. The DigitalDesk

The recognition of pages from the video images is crucial to the operation of the DigitalDesk. In some early applications an easily detected identifier was printed on each page to aid recognition. This worked well, but limited the use of the desk to specially printed documents. The recognition system was later improved so that normal printed documents could be used. The updated system works in four stages.

- i. The video images are processed to find the areas representing pages on the desktop.
- ii. These areas are then examined to find the positions of all lines of text.
- iii. Details of word spaces and ascenders/descenders are extracted for each line.
- iv. The pattern of lines on a page and the patterns of spaces, ascenders, and descenders for each line are then used to find the best match with a page in the central registry.



Figure 2: Text in a thresholded image

The current DigitalDesk technology is not capable of performing accurate optical character recognition. Figure 2 shows the resolution achieved for text after initial thresholding of the video image. It demonstrates that characters are unclear, that almost nothing in the way of font information is available to help with the recognition, and that the lines of text may not be horizontal in the image. For a book that does not lie completely flat on the desk, there is the added problem that the lines may be slightly curved. Figure 3 shows part of a page image after the base and x-height have been found for each line and further processing has been undertaken to identify the spaces between words and the ascenders and descenders within each word. The pattern of spaces/ascenders/descenders is then assembled into a 'line barcode' and the line barcodes for the page are checked against similar barcodes derived from the electronic text. For a relatively small document registry, this method gave an accurate page match in a reasonable time (a few seconds).

Once the corresponding digital page has been found in the registry, the position of each word in the electronic text can be mapped to a known position on the page.



Figure 3: Line image showing extracted features

The LED-tipped pen is easily recognised in the video image, so pointing to a position on the page with the pen can be interpreted in much the same way as moving a cursor to a position in a window on a computer screen. Similarly, clicking with the pen can be interpreted in much the same way as a mouse click and can therefore initiate an action associated with the visible text at that position.

The DigitalDesk thus involves physical paper in two novel ways. It allows the paper to act as an input device — each page acting rather like a fixed screen image. Its ability to project information onto the page also allows the paper to act in a limited but interesting way as an output device. (A further form of input, of course, is to draw on a piece of paper using a real pen or pencil and then capture the image.)

## 2.2 DigitalDesk applications — interactive mathematical notebook

An interactive mathematical notebook [15] was developed to demonstrate the potential of the DigitalDesk in teaching. This was a paper notebook that could be carried round and used in the normal way, but it took on additional features when placed on the DigitalDesk. Figure 4 shows a page from the notebook and (inset) a section of the video image showing it in use on the DigitalDesk. The page is part of a tutorial on the graphs of quadratic equations. The student can experiment with quadratics by writing coefficients into a general equation on the page and seeing the resulting graph projected onto a space on the page.

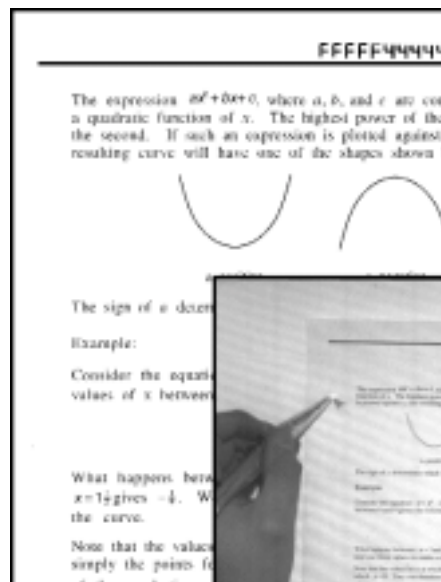


Figure 4: The interactive mathematical notebook

As this application used the earlier page recognition system, each page in the notebook contained a clear page identifier in an OCR font. This can be seen at the top right in Figure 4, together with a thick horizontal line that made it easier to recognise pages (and their orientation) in the video image.

### 2.3 DigitalDesk applications — the ‘Active Alice’ grammar lesson

The first major application using the new page recognition system to provide grammar lessons for children based on normal printed books with electronic versions available as part of the British National Corpus [16,17]. The application was called ‘Active Alice’ [18] because the book chosen as the first example was a simplified version of *Alice’s Adventures in Wonderland* [19].

The Alice text, in common with all texts in the BNC, is an SGML [20] text containing markup giving the part of speech of every word in the text. The BNC classifies words according to fifty-seven different *word class codes*. Some of the commonly used word classes are

AJ0	adjective (e.g. hot, strange)
AJC	comparative adjective (e.g. taller, worse)
AV0	adverb (e.g. suddenly, immediately)
CJC	coordinating conjunction (e.g. and, or)
DPS	possessive determiner (e.g. their, her)
NN1	singular common noun (e.g. book, sister)
NP0	proper noun (e.g. France, Alice)
PNP	personal pronoun (e.g. you, she)
VBD	past tense of verb to be (e.g. was, were)
VVG	-ing form of lexical verbs (e.g. living, thinking)
VVB	base form of lexical verbs (e.g. live, think)

The following (slightly simplified) version of the first two sentences of the electronic text shows that the word class codes are easy to extract from the text for use in the application.

```
<w NP0>Alice <w VBD>was <w VVG>beginning <w TO0>to
<w VVI>get <w AV0>very <w AJ0>bored.
<w PNP>She <w CJC>and <w DPS>her <w NN1>sister <w VBD>were
<w VVG>sitting <w PRP>under <w AT0>the <w NN2>trees.
```

When the book is placed on the DigitalDesk and identified, two windows of information are projected down onto the desk, one either side of the book. The area to the right is used for lesson information about word types, while the area to the left is used mainly for a quiz. These are both designed to involve the printed text where possible. Figure 5 shows the initial window that appears to the right of the book and the additional information that would appear below it if the user clicked the pen on the **Verbs** menu item

If the user subsequently chooses **More about tenses** in the ‘Verbs’ window, a third small window appears beneath the others giving information about different tenses and how certain forms of verbs may be derived by adding ‘ed’ or ‘ing’. There is always a **Show me . . .** item in the second level window. Selecting this involves the book in the lesson by highlighting all the words of the relevant type on the open pages (by projecting coloured rectangles down onto the book). About 30 different windows of information are directly available on the right of the book. The **Others** menu item in the initial window provides access to information on prepositions, conjunctions and some less common word types.

If the user clicks the pen on a word in the book, the word is highlighted by a coloured rectangle and information about the word is projected onto the area to the left of the book. Figure 6 shows the information that would be projected if the user clicked on the word ‘her’.



<p>To find out about words on a page, simply point to them with the pen</p> <p>To find out about different types of words, point to the word type in the list below</p>	<p>Verbs are words about doing (<i>singing, walk, played</i>), but they may also be about thinking, feeling or possessing</p> <p>Verbs have different tenses to express actions in the past, present, or future To <i>be</i>, to <i>have</i>, and to <i>do</i> are special verbs used in many different ways</p>
<p><b>Adjectives</b> (describing words) <b>Adverbs</b> (how/where/when/...) <b>Determiners</b> (a/the/my/some/that/...) <b>Nouns</b> (words for things) <b>Pronouns</b> (I/you/we/they/mine/...)</p>	<p><b>More about verbs</b> <b>More about tenses</b> <b>More about be/have/do</b> <b>Show me the verbs in the book</b></p>
<p><b>Others</b></p>	

Figure 5: Initial information (left) and information about verbs

<p style="text-align: center;">Determiner — possessive</p> <p>A determiner explains which noun we mean (a bull, your house, many books, that field)</p> <p>'her' is a possessive determiner — one that explains who owns the noun or nouns (my gloves, his pen, their books)</p>
--

Figure 6: Feedback on 'her'

A permanent quiz window invites the user to answer questions or find certain types of words. The problems always involve the currently open pages of the book. Some typical problems are

- Find 2 nouns
- Find a comparative adjective
- Find 3 different pronouns
- Find 2 words from the verb 'to have'
- Are the highlighted words adjectives or adverbs?

The interface is designed to involve the book as far as possible and to use a consistent colour coding for words of a given type, blue for adjectives and green for verbs, for example. Thus the words shown in italics in Figures 5 and 6 appear in colour in the projected information, and the same colour coding is used when highlighting words in the book.

## 2.4 Intelligent Paper

In 1998, Xerox announced a new technology known as Intelligent Paper. This was described in a paper in the Proceedings of Electronic Publishing 98 [8], and the information in this section is taken mainly from that paper.

The basic idea is that every sheet of Intelligent Paper should be uniquely identifiable to any computer attached to the World Wide Web. The only special equipment needed for this, apart from the paper itself, is a pen-like pointing device connected to the Web via a conventional or radio link. The pen can detect which sheet of Intelligent Paper it is pointing at and transmit this information to a Web-based computer.

Each sheet of Intelligent Paper is covered with special markings, invisible to the human eye, that uniquely identify both the sheet of paper and the position within the sheet. The pointer device can read these marks and, when the pen is clicked on the page, transmit the *page-id* and *position-within-page* over the Web to a network address derived from the

*page-id*. A program at this address interprets the page/position information and takes action. Typically this results in information being sent back over the Web to a screen or printer close to the user. Figure 7 illustrates a possible use. The user has an Intelligent Paper tourist map of Paris giving places of interest, such as the Eiffel Tower and the Louvre. To see information about the Louvre displayed on a screen, the user simply points to the Louvre on the map.

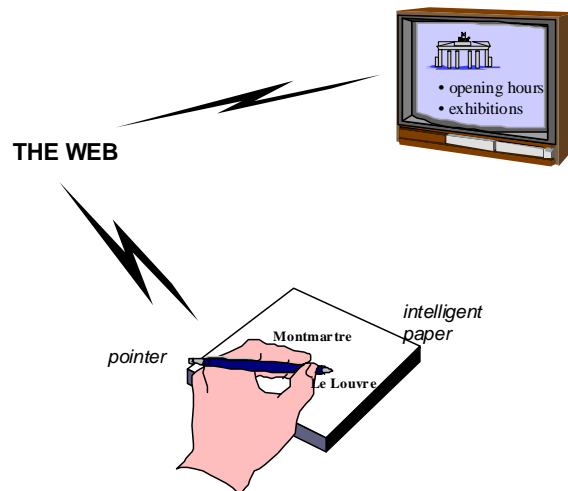


Figure 7: Intelligent Paper

The overall system is intended to work in the following way. Publishers buy supplies of Intelligent Paper from authorised producers. They print and distribute the paper in the normal way. They also develop Web-accessible software that associates specific actions with given positions on given sheets of paper. Intelligent Paper looks and feels just like any other paper, so users retain all the normal advantages of portability, feel, and flexibility. Thus maps, books, and catalogues can still be used in all the normal ways. As explained in [8], however, 'if the user owns a pointer and is near a peripheral device connected to the Web, the paper document assumes the behaviour of a touch-sensitive screen. It becomes enriched with arbitrarily complex information, provided by the document's publisher at its Web site'.




Figure 8: Map of Europe printed on Intelligent Paper

The underlying technology of Intelligent Paper uses two layers of ink. Figure 8 shows a map of Europe printed on Intelligent Paper. The top layer is normal visible ink. The lower layer provides the identification information and is known as the code layer. This is printed in invisible ink and uses a special technology known as Xerox DataGlyphs™. DataGlyphs may be used in a variety of ways, but in one of the simplest variants they allow cells approximately a quarter of an inch square to be laid out in a grid pattern. Each cell is printed so that it contains approximately 150 bits

of data. Pages up to 5 feet square only need 16 bits of data to represent the *position within-page*, so this allows 64 bits to represent a unique *page-id* and still leaves plenty of spare bits for security purposes, error correction, and future special applications.

The pen can be thought of as a camera that can view a complete cell. In practice it can see just enough of the DataGlyphs to determine not only the contents of a cell but also the position of the cell relative to its field of view. This allows it to determine its *position-within-page* to a high degree of accuracy.

A pen click causes the *page-id* and *position-within-page* to be routed over the Web to the publisher's Web site as a request for an action. Several routing mechanisms can be used for determining the Web address from the page identification. To keep the routing tables needed to associate *page-ids* with URLs down to a reasonable size, these may work in several stages. Once the request reaches the publisher's Web site, the responsibility for matching it up with the correct digital page and interpreting and acting on the request is entirely up to the publisher. Routing the resulting information back to a Web-addressable computer near the user is a problem that could be addressed in a number of ways. One solution could be the use of 'Intelligent Confetti' as described in the next section.

Intelligent Paper provides a new and very different way of using paper as an input device to a computer system. The 'hardware' needed (Intelligent Paper and a specialised pointer) is much less cumbersome than the DigitalDesk, but a large Web-based software and routing infrastructure is needed to realise the full potential of Intelligent Paper as an input medium for any Web-based computer. Once this infrastructure exists, however, the possibilities are enormous.

## 2.5 Intelligent Paper applications

The Xerox paper describes two applications in some detail and suggests some future possibilities. This section focuses on one application called Reading *Othello* and outlines a few of the future possibilities mentioned.

The Othello example is based on the actions a student might take when reading her Intelligent Paper version of Shakespeare's Othello. Ariane is a French student of English Literature. She is lying in bed, in front of her Web-addressable TV, reading her version of Othello. This looks like any other Othello except that some pages have boxes printed across the top looking like menu items on a screen. Some actions that Ariane might take are described below.

- i. She comes across a word she does not understand, so she clicks on the word and sees an explanatory note about the word displayed on the TV screen.
- ii. She wishes to hear the scene, so she clicks on the PLAY SCENE box at the top of the first page of the scene and hears the scene recited.
- iii. She wants to browse the Web site for the book, so she clicks on the MAIN WEB PAGE box and sees the initial Web page for that version of Othello displayed. The pen now starts to act like a mouse for the TV screen. Moving the pen on the page moves a cursor on the screen and clicking the pen acts like a mouse click. This allows her to browse the Web site.
- iv. Browsing the site leads her to a particularly relevant page. She wants to make a note of it for future use, so she pencils a note in the margin of the book, clicks the pen on the LINK box at the top of the page and then on the pencilled note. This sets up a new action associated with the position of the pencilled note on that page. Later on, if she clicks the pen on her pencilled note, the relevant Web page will be displayed.

The LINK facility outlined in (iv) above provides an interesting way of adding individually-tailored information to pre-printed Intelligent Paper documents.

Other potential Intelligent Paper publications might be

**maps:** providing access to tourist information and timetables for public transport

**pocket encyclopedias:** providing access to complete multimedia versions

**product catalogues:** providing access to order forms and multimedia information on products



A final variation is the possible use of Intelligent Confetti — small round stickers carrying a *page-id* only. These could be used to make additional facilities available when stuck on sheets of Intelligent Paper or, more importantly, used to identify Web devices to be used for output. This would allow a user to set up the return address for actions by clicking first on their Intelligent Paper and then on the Intelligent Confetti on the required output device.

## 3 Using the Text Encoding Initiative

### 3.1 Introduction to the TEI

The TEI is an application of SGML designed specifically for capturing a wide variety of information about literary texts. It covers complex literary and linguistic structures, different versions of texts, the state of old manuscripts, scholarly commentaries about texts and much more. The TEI guidelines are comprehensive; they occupy 1300 pages and cover the use of more than 400 different elements.

The British National Corpus (BNC) uses the TEI guidelines, so the ‘Active Alice’ project described in Section 2.3 was an example of accessing a TEI text from paper. Like all BNC texts, the Alice text used only a small subset of TEI facilities to record simple information about parts of speech. Several major TEI projects, however, have captured detailed scholarly information about old works (see [21] for a list of TEI projects). The Canterbury Tales Project [22], for example, created TEI versions of the fifty-four manuscripts of Chaucer's Wife of Bath's Prologue. A CD-ROM version of these [23] provides detailed comparisons of different versions and full collations of all the witnesses to the Prologue.

An interesting way to exploit the effort put into creating these rich TEI versions is to use a modern paper version of an old work to access the additional scholarly information in the TEI version. The rest of this section introduces a small selection of TEI facilities relevant to this application and Section 3.2 looks in more detail at a specific example.

Old manuscripts contain many deletions and insertions. Deletions may be made by crossing out or erasing. Insertions are typically written above the original line or in a margin. The two are frequently combined when text is replaced. Such a replacement might be indicated in the TEI text as follows

```
Descending by <del type=overstrike>the easy slope</del>
<add place=supralinear hand=DW>the beaten road</add>
that led
```

The `del` element indicates that the original text ‘the easy slope’ was deleted by drawing a line through the text (`type=overstrike`) and the `add` element indicates that ‘the beaten road’ was then added above the original line (`place=supralinear`). `hand=DW` specifies that the new text was written in a different handwriting identified by the name ‘DW’. Different types of deletion (e.g. `erasure`) and addition (e.g. `leftmargin`, `overleaf`) could have been specified and several other attributes are available to provide further details.

Elements such as `gap`, `damage` and `unclear` can be used to indicate that the original text was illegible or not clear enough to be transcribed with confidence. Where the text for a damaged section has been guessed or suggested by a scribe or scholar, it can be indicated by a `supplied` element. For example

```
<damage cause="rubbing">
<supplied reason="illegible" resp="PR3">
at this season </supplied> </damage>
```

might be used to specify that a scholar identified by the name ‘PR3’ was responsible (`resp="PR3"`) for supplying the text ‘at this season’ as a likely interpretation of text that was illegible due to rubbing of the manuscript.

Variations on these elements, including nested additions and deletions, may be sufficient to cover changes and problems within a single manuscript, but further elements — including `app`, `lem`, and `rdg` — are needed to encode information about different versions and copies of the same work. `app` is used to indicate the presence of some ‘critical apparatus’. `lem` and `rdg` distinguish base or ‘lemma’ versions of a text from different ‘readings’ given by other scribes or in other versions. The following simple example (taken from the TEI Guidelines) shows how variations of a single word from four different versions might be coded.

```
<app>
  <lem wit="E1 Hg">Experience</>
  <rdg wit="La" type="substantive">Experiment</>
  <rdg wit="Ra2" type="substantive">Eryment</>
</app>
```

The `wit` attribute identifies the different ‘witnesses’ or sources. At the start of the manuscript there will be `witlist` element containing information which relates the names used to identify the witnesses (E1, Hg, La, and Ra2 in the example) to a description of the relevant people or manuscripts. Where many variations exist, `rdgGrp` elements could be used to subdivide the appropriate `lem` and `rdg` elements into several different families or ‘reading groups’.

Two further examples should be sufficient to give an idea of the range of TEI information that might be available in a manuscript. The first of these is related to verse. TEI elements used for verse include `lg` for line groups, `l` for lines, and `seg` for segments within lines. Attributes for these elements can be used to give details of metrical analysis, rhyme patterns, and the type of verse. Thus

```
<lg type=stanza met="-+--+--+/-+--+-" rhyme="ababcdcd">
```

might be used to introduce a stanza and specify its metrical pattern and rhyme. If a line within the stanza did not conform to the overall metrical pattern, a `met` attribute could be used as shown below to override the pattern given by the line group element.

```
<l n=2 met="---+--">
```

The final, rather different, example concerns information about names and places. The TEI provides `name` and `rs` (referring string) elements for names and places. `name` is used for proper nouns and `rs` for general purpose names. For example

```
In June <name type=person key=SM3 reg="Smith, Dr John">
Smith</name> moved to <name type=county>Kent</name>. Later
in the year <rs type=person key=SM3>the Doctor</rs>
```

Here the `key` attribute provides a unique identifier for the person or place concerned, and `reg` gives a regularised or standardised version of the name. More detail may be encoded than shown above. For names, this could include forenames, surname, titles referring to rank or role (e.g. Lord, Colonel), nicknames, and embellishments like ‘Junior’, ‘de la’, or ‘IV’. For places, information can be given about the country, region, settlement (town), and relative position or distance from another place. If used consistently, these elements make it easy to find every reference to particular people and places in a work.

### 3.2 *The Wordsworth project*

The Wordsworth Trust [24] and staff at the Universities of Newcastle and Kent have held discussions about a project to create TEI versions of a selection of Wordsworth's manuscripts. The motivation for this is to allow scholars to access information over the Web instead of having to travel to Dove Cottage to inspect the manuscripts. If the project goes ahead, the resulting TEI texts could also be accessed via the DigitalDesk or Intelligent Paper. This section outlines some of the possibilities assuming the TEI versions contain the type of information outlined in Section 3.1 above.

Traditionally, humanities scholars have published weighty books containing transcriptions of manuscripts and comparisons of different versions, together with numerous footnotes describing the state of the manuscripts and giving information about the history or interpretation of the text. 'The Thirteen-Book Prelude by William Wordsworth' edited by Reed [25] is a good example. Wordsworth's 'Prelude' is a poem of approximately 8000 lines, but the Reed version consists of two substantial volumes containing a total of 2300 pages. They form part of a series known as *The Cornell Wordsworth*, which aims to present 'clean, continuous "reading texts" from which all layers of later revision have been stripped away' and also to provide 'a complete and accurate record of variant readings, from Wordsworth's earliest drafts down to the final lifetime (or first posthumous) publication'. They do this by providing three main versions

- the clean reading text
- transcriptions giving a printed approximation of all the insertions and deletions in the originals
- photographs of the manuscript pages

Extensive notes are provided explaining the history and state of the various versions and giving additional information about the meaning of the text. The following is a typical example taken from the transcriptions

[9<sup>r</sup>]  
 While yet  
 [ ? ~~Then never~~ ]  
 Ere I [ ? had ] travell'd from my native hills [XI, 224]  
 And afterwards when thought the [241 –242]  
 Gorgeous  
 Alp  
 I roam'd { <sup>my</sup> [ ? ] pleasures were entirely

The positioning and spacing of the text is a close representation of that used in the original manuscript. The items in square brackets within the text represent doubtful or illegible words which would be encoded using `gap` and `unclear` elements in a TEI version. The '[9<sup>r</sup>]' at the start identifies the sheet of the original manuscript. The numbers in brackets to the right of the text refer to the book and lines in the manuscript and, where appropriate, give the corresponding line numbers in the clean reading version. These numbers are used to link footnotes to the appropriate text.

Similar information concerning the manuscripts could be encoded using the TEI guidelines, although it would require the use of advanced versions of the features described in the previous section. For Web access or CD-ROM versions, this information could be made available via a standard hypertext interface. For the DigitalDesk or Intelligent Paper, a printed version containing the clean reading text could be used as the starting point.

On the DigitalDesk this would work as follows. Once a page of the work has been recognised, the user is presented with a menu of the various types of information available for the work. The user selects one or more of these. Areas of the open pages are then highlighted (using a colour code linked to the menu) to show where this type of information is relevant. To access the details, the user then points to a specific highlighted area. Further feedback may be given via information projected onto the paper/desk or via the computer screen. As a specific example, the user could ask to see all the places where a particular manuscript differs from the reading text. Highlighting all such areas on the page would provide a vivid illustration of the pattern and level of the differences. Full details of the other version would be projected onto a separate area of the desk or presented via a screen.

Should the reader wish to check up on references to particular names, the system might give information at two or three different levels. Initially, all names on the open pages could be highlighted. The user could then point to a particular highlighted name to ask to see the additional information. At this point the user could also have the option of asking for a list of all references to the same person or place (this could be given by displaying appropriate text extracts on the screen or by giving page references for the paper version).

The system would work in a rather similar way with Intelligent Paper, but the initial selection of the type of information required would probably be made by pointing to menu boxes printed on the page. Instead of highlighting

areas on the page to show where additional information is available, the system would display an appropriately highlighted version on a Web-connected computer screen. Further levels of information would then be requested by pointing to the appropriate text on the page or by using the pen as if it were a mouse for the screen.

## Summary

Much effort has been expended in replicating the qualities of paper on a computing device. Although great progress has been made, there is as yet no real alternative to the paper document that can be folded up and stuffed in a pocket. This paper has described two experimental systems for bridging the gap between the paper and digital worlds. Both approaches allow physical paper to act as a starting point for accessing and exploring the dynamic world of digital information.

Although the current DigitalDesk prototype has significant limitations in its recognition system, it has already shown that many applications are feasible. Its ability to project information onto the paper provides an eye-catching and attractive way of involving the paper at all stages, not just as a starting point for accessing the digital information.

Intelligent paper opens up a whole new world of Web-based applications and has potentially far-reaching implications for publishers. Although it does not provide an immediate equivalent of the DigitalDesk's ability to highlight information on the paper, it is an extremely flexible technology that fits well with the ubiquitous nature of the Web and advances in mobile computing. The LINK facility, referred to in Section 2.5, brings a whole new meaning to annotations on paper.

## Acknowledgements

The authors would like to thank: Robert Harding, Steven Lay, Dan Sheppard, and Richard Watts for their contributions to the DigitalDesk and the Active Alice project; Marc Dymetman and Max Copperman for the inspiration provided by their paper on Intelligent Paper and for their willingness to allow us to quote from it and use their Figures to illustrate the explanations given in Sections 2.4 and 2.5; and members of the Wordsworth Trust and staff at the Universities of Newcastle and Kent for ideas about the Wordsworth project.

## References

1. Victoria A. Burrill, 'VORTEXT: VictORias TEXT reading and authoring system', in *Text processing and Document Manipulation*, J. C. van Vliet (ed), pp. 43–57, Cambridge University Press, 1986.
2. Abigail Sellen and Richard Harper, 'Paper as an analytic resource for the design of new technologies', in *Proceedings of the ACM SigChi Conference on Human Factors in Computing Systems, CHI'97*, Steven Pemberton (ed), pp. 319–326, ACM/Addison-Wesley, 1997.
3. Andrew Birrell and Paul McJones, 'Virtual Paper', *Project Web Page*, Digital Equipment Corporation Systems Research Center. (<http://www.research.digital.com/SRC/virtualpaper/home.html>)
4. David Chaiken, Mark Hayter, Jay Kistler, and Dave Redell, *The Virtual Book*, SRC Research Report 157, Digital Equipment Corporation Systems Research Center, November 1998.
5. Morgan Price, Bill Schilit, Gene Golovchinsky, 'XLibris: The Active Reading Machine', in *CHI 98 Summary*, pp. 22–23, ACM Press, 1998. (<http://www.fxpal.xerox.com/papers/pri98.pdf>)
6. Bill Schilit, Lynn Wilcox, and Nitin "Nick" Sawhney, 'Merging the Benefits of Paper Notebooks with the Power of Computers in Dynamite', in *CHI 97 Extended Abstracts*, pp. 22–23, ACM press, 1997. (<http://www.fxpal.xerox.com/papers/sch97.pdf>)
7. William Newman and Pierre Wellner, 'A desk that supports computer-based interaction with paper documents', *Proceedings of the ACM Conference on Human Factors in Computing*, Monterey, 1992.

8. Marc Dymetman and Max Copperman, 'Intelligent Paper', in *Electronic Documents, Artistic Imaging, and Digital Typography*, Eds. Roger D. Hersch, Jacques Andre and Heather Brown, pp. 392–406, Springer-Verlag, LNCS, 1998.
9. C. M. Sperberg-McQueen and Lou Burnard (editors), *Guidelines for Electronic Text Encoding and Interchange (TEI P3)*, ACH/ACL/ALLC (Association for Computers and the Humanities, Association for Computational Linguistics, Association for Literary and Linguistic Computing), Chicago/Oxford, 1994.
10. Peter Robinson, 'Virtual Offices', *Proceedings of Royal Society Discussion Meeting on Virtual Reality in Society, Science and Engineering*, BT Publication SRD/R5/1, 1995.
11. Pierre Wellner, 'Interacting with Paper on the DigitalDesk', *Communications of the ACM*, 36(7), pp. 87–96, July 1993.
12. Peter Robinson, Dan Sheppard, Richard Watts, Robert Harding and Steve Lay, 'A framework for interacting with paper', *Proceedings of Eurographics '97*, D. Fellner and L. Szirmay-Kalos (Eds), Vol. 16, No. 3, 1997.
13. Quentin Stafford-Fraser and Peter Robinson, 'BrightBoard — a video augmented environment', *Proceedings of the ACM Conference on Human Factors in Computer Systems*, Vancouver, 1996.
14. Kathy Carter, 'Computer aided design — back to the drawing board', *Proceedings of the International Symposium on Creativity and Cognition*, Loughborough, 1993.
15. Peter Robinson, Dan Sheppard, Richard Watts, Robert Harding and Steve Lay, 'Animated Paper Documents', *Proceedings 7<sup>th</sup> International Conference on Human-Computer Interaction, HCI'97*, San Francisco, August 1997.
16. Gavin Burnage and Dominic Dunlop, 'Encoding the British National Corpus', in *English Language Corpora: Design, Analysis and Exploitation*, Jan Aarts, Pieter de Haan and Nelleke Oostdijk (eds), pp. 79–95, Amsterdam and Atlanta: Editions Rodopi, 1993.
17. G. Leech, '100 million words of English', *English Today*, 9(1), 1993.
18. Heather Brown, Robert Harding, Steve Lay, Peter Robinson, Dan Sheppard and Richard Watts, 'Active Alice: Using Real Paper to Interact with Electronic Text', in *Electronic Documents, Artistic Imaging, and Digital Typography*, Eds. Roger D. Hersch, Jacques Andre and Heather Brown, pp. 407–419, Springer-Verlag, LNCS, 1998.
19. Lewis Carroll (retold by Jennifer Bassett), *Alice's Adventures in Wonderland*, Oxford Bookworms, Oxford University Press, 1994.
20. Charles F. Goldfarb, *The SGML Handbook*, Oxford University Press, 1990.
21. TEI home page <http://www-tei.uic.edu/orgs/tei/index.html> and list of applications (<http://www-tei.uic.edu/orgs/tei/app/index.html>).
22. 'The Canterbury tales project', <http://www.shef.ac.uk/uni/projects/ctp/index.html>.
23. Peter Robinson, *The Wife of Bath's Prologue on CD-ROM*, Cambridge University Press, 1996.
24. Wordsworth Trust Home Page, <http://www.wordsworth.org.uk/>.
25. Mark L Reed (ed), 'The Thirteen-Book Prelude by William Wordsworth, Volumes I and II, Cornell University Press, 1991.