Supporting Tangible User Interaction with Integrated Paper and Electronic Document Management Systems

A thesis
submitted in fulfillment
of the requirements for the degree
of
Doctor of Philosophy
in
Computer Science

at

The University of Waikato

by

Matthew Jervis

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Department of Computer Science
Hamilton, New Zealand
2011

Abstract

Although electronic technology has had a significant impact on the way that offices manage documents, in most cases electronic documents have not completely replaced paper documents. As a result, many present-day offices use a combination of paper and electronic documents in their normal work-flow. The problem with this, however, is that it results in information and document management becoming fragmented between the paper and electronic forms. There is, therefore, a need to provide better integration of the management of paper and electronic documents in order to reduce this fragmentation and, where possible, bring the advantages of electronic document management to paper documents.

Previous research has investigated methods of incorporating management and tracking of paper documents into electronic document management systems. However, better integration between paper and electronic document management is still needed, and could potentially be achieved by augmenting elements of the physical document management system with electronic circuitry so they can support tangible user interaction with the integrated document management system. Therefore, the aim of this thesis has been to investigate this.

The approach that was taken began by identifying the requirements of such integrated systems by studying the document management needs of a number of real-world offices. This was followed by the development of a series of prototype systems designed to function as tangible user interfaces to the integrated document management system. These prototypes were then evaluated against the identified requirements, and a user study was conducted in order to evaluate their usability. The results of these evaluations demonstrate that it is possible to develop systems systems that can utilise tangible user interaction techniques to enhance the integration of paper and electronic document management, and thus better bridge the divide between the physical and virtual worlds of documents.

Acknowledgements

I would never have started, let alone completed, this thesis without the help and support of many people.

Firstly, I owe a great deal of gratitude to my chief supervisor, Masood Masoodian, without whose guidance and encouragement thesis would never have been possible. I am also grateful to Mark Apperley, my second supervisor, for his assistance. Also, I would like to thank Doris Jung for proof reading and providing me with valuable feedback.

A big thank you to all the university staff who have helped with all the important things along the way, like scholarship applications, ethics approval, keeping the tea room stocked.

It would not have been possible for me to complete my research without funding. After all, there's only so long a person can survive on two-minute noodles. So I am thankful for the funding that I have received from the Tertiary Education Commission, the University of Waikato, Waikato Link Limited, BuildIT, and the New Zealand Vice-Chancellors' Committee.

Thanks to all those who participated in the two studies that I conducted, and without whose participation the studies would not have been possible.

I would also like to thank the WAND network research group for their support during my research.

Thanks to my friends who have supported me in so many ways, and made my world a more enjoyable place: you rock.

These acknowledgements would not be complete without expressing my sincere gratitude to my family, especially Mum and Dad, for their love and encouragement.

Contents

List of Figures

List of Tables

List of Acronyms

ADC	Analog-to-Digital Converter
API	Application Programming Interface
AR	Augmented Reality
CRC	Cyclic Redundancy Check
DMS	Document Management System
ECM	Enterprise Content Management
ECRM	Enterprise Content and Records Management
ERM	Electronic Records Management
FPGA	Field Programmable Gate Array
GPIO	General Purpose Input/Output
GUI	Graphical User Interface
HMD	Head-Mounted Display
IC	Integrated Circuit
ID	Identifier
ITO	Industry Training Organisation
LAN	Local Area Network
LCD	Liquid Crystal Display
LED	Light Emitting Diode
MSB	Most Significant Bit
NZQA	New Zealand Qualifications Authority
PA	Personal Assistant

PCB Printed Circuit Board

PDA Personal Digital Assistant

PDM Personal Document Management

PIM Personal Information Management

RDF Resource Description Framework

RFID Radio Frequency Identification

RFIG Radio Frequency Identity and Geometry

RM Records Management

SSH Secure Shell

TCP/IP Transport Control Protocol over Internet Protocol

TUI Tangible User Interface

UI User Interface

USB Universal Serial Bus

Publications

The following publications were produced during the course of this research.

1. Seifried, T., Jervis, M., Haller, M., Masoodian, M., Villar, N. (2008). Integration of virtual and real document organization. In *TEI '08: Proceedings of the 2nd international conference on Tangible and embedded interaction*, pp. 81–88. New York, NY, USA: ACM.

2. Jervis, M., Masoodian, M. (2009). Digital management and retrieval of physical documents. In *TEI '09: Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, pp. 47–54. New York, NY, USA: ACM.

3. Jervis, M. G., Masoodian, M. (2010b). SOPHYA: a system for digital management of ordered physical document collections. In *TEI '10: Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*, pp. 33–40. New York, NY, USA: ACM.

4. Jervis, M., Masoodian, M. (2010a). Design of a Modular Architecture for Integrating Paper and Digital Document Management Systems. In *Workshop Proceedings of PaperComp 2010, The 1st International Workshop on Paper Computing, Part of the 12th ACM International Conference on Ubiquitous Computing*, pp. 44-47. New York, NY, USA: ACM.

5. Jervis, M., Masoodian, M. (2011). Evaluation of an integrated paper and digital document document management system. In *Proceedings of INTERACT 2011, The 13th IFIP TC13 Conference on Human-Computer Interaction*, Springer-Verlag, in print.

Part I

Background

Chapter 1

Introduction

There can be no doubt that paper has had a significant role in the development of the society that we live in today. It has served as a relatively low-cost and easily-produced medium to support information storage and communication. From the early days of its use paper has played a role in business. For example, during the Tang dynasty in China (618–907 AD) half a million sheets of paper were used annually by the Board of Finance for tax assessments (Wood, 1986). By the turn of the twentieth century paper had become an integral part of business and society in the West.

However, the rapid development of computer technology towards the end of the twentieth century threatened to render paper obsolete. As this technology has evolved, there have been predictions of a future *paperless* office (e.g., "The Office of the Future", 1975). This desire for a paperless office is motivated by a range of factors, such as the cost and difficulty of storing and managing paper documents, the interactional limitations of paper documents (e.g., lack of support for dynamic content), and the commonly held notion that paper is an outdated relic, a symbol of out-dated business practice. In comparison, electronic documents have numerous advantages, such as being easier to index and retrieve, taking less physical space, and allowing for faster delivery.

Despite the predictions, the paperless office has largely remained a myth (Sellen and Harper, 2003). While a relatively small number of offices[1] have made a complete transition to a paperless environment, most offices continue to use paper in some way, whether as a central part of their work-flow or in a lesser role. Significant reasons for this include the fact that paper documents have

[1]The term office is used in this thesis to describe any workplace where documents are used.

certain advantages that electronic documents are not yet able to fully match (e.g., fast and flexible navigation, ease of annotation, etc.) and that many offices still have large quantities of historical paper documents.

1.1 Motivation

The continuing existence of paper documents, and their use alongside electronic documents, raises two key sets of problems: those due to the fragmentation of information and document management between the paper and electronic forms, and those due to the limitations of paper (e.g., the costs and difficulties of managing paper documents in comparison to electronic documents). These problems can be approached from two different directions. One direction is to bring the advantages of paper documents to electronic systems so that they can completely replace paper. The other direction is to bring the benefits of electronic documents to paper documents and provide seamless integration between them.

Although these two approaches are complementary, and not mutually exclusive, the research presented in this thesis has focused more on the second approach. The main reason for deciding to focus on making the integration of paper and electronic document management more seamless is that, despite the progress that has been made in enhancing electronic documents, paper still plays a crucial role in present-day offices (as will be discussed in Chapter 2), and this is likely to remain so for the foreseeable future.

However, one of the major barriers to providing integration of paper and electronic document management, which would need to be overcome by any proposed system, is the divide that exists between the physical world (in which paper documents exist, and their users reside) and the virtual world (in which electronic documents exist, and computer software operates). Current research (reviewed in Chapter 3) has shown that two main problems need to be addressed in order to bridge this divide:

- providing links between the electronic document management systems in the virtual world, and paper documents in the physical world; and
- providing users with means of interacting with virtual electronic document management systems that go beyond the limitations of

conventional Graphical User Interfaces (GUIs).

Most existing approaches to integrating paper and electronic document management focus on the first of these two problems. For example, they aim to link paper and electronic documents by manually giving paper documents unique identifiers and using this as a key to their electronic metadata (e.g., Wilson, 2001); or they use some form of machine readable tag, such as a barcode or Radio Frequency Identification (RFID) tag, that can provide—depending on the implementation—varying degrees of tracking of physical documents and triggering of events in software (e.g., Smith et al., 2006; Bodhuin et al., 2007).

What is lacking, to a large extent, is effective systems that provide users with means of interacting with the virtual components of the document management system that go beyond the GUI. Most software applications, including electronic document management systems, require users to interact with the virtual world using a GUI. However, such an interface becomes limiting when managing paper as well as electronic documents. For example, even when some form of electronic document tracking system is provided to integrate management of paper documents, the user interaction is generally limited to displaying instructions for finding a given paper document using a conventional computer software interface, without directly presenting such information in the physical world in which that paper document exists.

In contrast to such an approach, this thesis proposes that a more seamless user experience across the divide between the physical world of paper documents and the virtual world of electronic documents can be provided by utilising tangible user interaction techniques (see Chapter 3). Especially given that one of the most enduring reasons for wide-spread use of paper documents is related to the affordances of paper as a tangible physical entity, which most people enjoy interacting with (see Chapter 2). Therefore, the question that needs to be investigated is: how can such a tangible form of interaction be facilitated?

1.2 Objectives

The primary goal of the research described in this thesis has been to investigate the following question:

To what extent can tangible user interaction facilitate a more seam-

less integration of paper and electronic document management?

In order to answer this broad question, the research presented in this thesis addresses the following secondary questions:

1. *What roles do paper and electronic documents serve, how are they managed, and what problems are encountered with their management in present-day offices?*
2. *What are the requirements of an integrated paper and electronic document management system?*
3. *How could a tangible user interface be designed to provide a more seamless integration of paper and electronic document management?*
4. *To what extent would the proposed tangible user interface fulfil the requirements of an integrated paper and electronic document management system?*

1.3 Approach

The research methodology that has been followed in order to answer the questions proposed in the previous section comprises four main stages:

1. literature review;
2. identification of requirements;
3. design and implementation; and
4. evaluation.

The research described in this thesis begins with a review of relevant literature, which is split into two chapters. Chapter 2 reviews document management literature, beginning with a discussion of the relative merits of paper and electronic documents, and why, as a result, paper documents are likely to have a role in many offices for the foreseeable future. This is followed by a review of previous studies of document and information management. Chapter 3 follows with a review of previous research towards integrating paper and electronic document management. This chapter begins with an overview of the research towards bridging the divide between the physical and virtual worlds, which is followed by a more in-depth look at research towards integrating paper and electronic document management.

In order to gain a better understanding of the problem, and confirm that the findings of previous studies of paper and electronic documents still hold true in present-day offices, a study of paper and electronic document use in offices was conducted (Chapter 4). This study investigated the roles paper and electronic documents play in modern offices of today, as well as how they are managed, to specify a set of requirements for integrated paper and electronic document management systems.

The design aspects of an integrated paper and electronic document management system that were taken into consideration are discussed in Chapter 5, and these guided the development of an initial proof-of-concept prototype system, which is described in Chapter 6. This system was developed as a tangible component of the Aroo integrated document management system, which is also described in Seifried et al. (2008).

In order to improve on the first prototype and create a system that could be used for user evaluation, a second prototype system was developed, the design and implementation of which is discussed in Chapter 7. Details of this system, called SOPHYA[2], have also been published in Jervis and Masoodian (2009). Because this system is designed for management of document collections in which strict ordering is not necessary, it is referred to as the *unordered* SOPHYA prototype.

Another prototype, which is based on the concepts of the unordered SOPHYA prototype, but also provides a means of managing ordered collections of documents was subsequently developed. This system, referred to as the *ordered* SOPHYA prototype, is described in Chapter 8, and reported in Jervis and Masoodian (2010b). A description of the architecture and complementary nature of these two SOPHYA systems, particularly in terms of their middleware, and their potential applications is described in Chapter 9, and introduced in Jervis and Masoodian (2010a).

The two SOPHYA systems were then evaluated in terms of fulfilling the specified requirements (described in Chapter 4), and this evaluation is discussed in Chapter 9. Furthermore, in order to gain feedback from potential users

[2]SOPHYA is short for 'Smart Organisation of PHYsical Artefacts'.

of such an integrated document management technology, an evaluation of the SOPHYA prototypes was conducted, and is described in Chapter 10. Results of this study are also published in Jervis and Masoodian (2011).

1.4 Contributions

The research in this thesis makes the following original contributions:

- A critical review of literature concerning document management (Chapter 2) and existing technology for bridging the divide between paper and electronic document management (Chapter 3).
- Investigation of the use of paper and electronic documents in present-day offices to identify methods of their management. This study (Chapter 4) contributes a review of the roles that paper and electronic documents serve, and the ways that these are managed in a range of present-day offices.
- Identification of a set of requirements for integrated paper and electronic document management systems (Chapter 4), and a design based on this, as proposed in Chapter 5.
- Design and implementation of a series of three different prototype systems to support integrated paper and electronic document management (Chapters 6–8). The details of these prototypes have been published in several papers (Seifried et al., 2008; Jervis and Masoodian, 2009, 2010a,b).
- Evaluation of the two SOPHYA prototype systems. This has included an evaluation of the systems in terms of fulfilling the specified requirements (Chapter 9), as well as a study of their effectiveness in providing an integrated document management system (Chapter 10). This has also been published in Jervis and Masoodian (2011).
- Identification of the benefits and potential barriers to deployment of integrated document management systems as viewed by their potential users (Chapter 10). Limitations of the prototype systems have also been identified and possible solutions to these have been proposed.

1.5 Thesis structure

The thesis is structured into five parts. The first four parts correspond to the four main stages of the research that were conducted, and the fifth part contains references and supporting material in several appendices.

Part I Background

Chapter 1	Introduction
Chapter 2	Background
Chapter 3	Related Work

Part II Identifying Requirements

Chapter 4	Study of Document Use in Offices
Chapter 5	Design of an Integrated System

Part III Development

Chapter 6	Prototype of an Integrated System
Chapter 7	Architecture and Implementation of a Revised Prototype System
Chapter 8	Prototype System for Managing Ordered Document Collections

Part IV Evaluation and Conclusions

Chapter 9	Evaluating the Functional Requirements of SOPHYA
Chapter 10	Evaluation Study
Chapter 11	Conclusions and Future Work

Part V References and Appendices

References	
Appendix A	Ethics Approval for Observational Study
Appendix B	SOPHYA Container Network Protocol
Appendix C	SOPHYA Circuitry
Appendix D	Ethics Approval for Evaluation Study

Appendix E Evaluation Study Handbook

Appendix F Video Demonstration

Chapter 2

Background

The focus of the research presented in this thesis is the integration of the management of documents in paper and electronic form. This chapter describes the background for this research. It begins with a discussion in Section 2.1 of the problems and advantages of paper, and why it is still important to account for paper documents in present-day offices where documents are increasingly being produced, stored, and transmitted in electronic form. This is followed by a discussion of document management in Section 2.2, that defines a number of important terms, outlines several related fields, and describes a framework for the discussion of document management. Section 2.3 provides a review of the literature relating to document management. This is followed by a discussion (Section 2.4) and summary (Section 2.5).

2.1 Paper Documents in the Present-Day Office

Frohlich and Perry (1994) describe the continued use of paper in the office, despite the availability of electronic alternatives, as the *paperful office paradox*. In order to gain a better understanding of why this paradox exists, they conducted a study comprising interviews and observation of 15 people from a variety of backgrounds working at a single organisation. They found that paper was mainly used by their participants for interpersonal communication, as well as idea generation and thought structuring, external memory store, planning, long-term storage, problem solving and note taking. The advantages of paper that they identified were mostly due to its physical and tangible nature. These advantages included transportability, ease of annotation, and pleasantness to read. Conversely the advantages of electronic documents that they identified were mostly related to their organisation, rather than the documents themselves.

The advantages of paper described by Frohlich and Perry are mostly derived from its affordances. Affordances are "*the perceived and actual properties of the thing, primarily those fundamental properties that determine just how the thing could possibly be used*" (Norman, 1990). Paper has numerous affordances that have not generally been duplicated with electronic documents to date (at least without expensive, uncommon hardware).

Some of the affordances of paper identified by Sellen and Harper (2003) include: grasping, carrying, manipulating, folding, writing on and reading. Paper affords fast, flexible navigation and the ability to cross reference multiple documents simultaneously. It also affords annotation and the interweaving of reading and writing.

Furthermore, when individual sheets of paper are combined into books these have a new set of affordances. Books afford flicking through, reading from and place holding (e.g., by modifying the book by dog-earing, folding, or even ripping pages, or through adding elements such as page tabs and bookmarks (Rosner et al., 2008)).

Paper documents are also useful for long-term storage of information. This was demonstrated by Whittaker and Hirschberg (2001), who conducted an investigation of office paper usage from a processing and archiving perspective. The study comprised an online survey of 50 people, followed up by interviews with 14 of the original group. The participants were office workers from a range of roles at a single company, whose office was undergoing a move. Whittaker and Hirschberg noted that whereas previous research had identified *reading* and *collaboration* as tasks where paper was commonly utilised, it was shown by their study that paper that paper documents were also useful for *long-term memory and reference*. They argue that information in paper form is in fact valuable in the long-term, and is not just ephemeral.

Recently, advances in technology have led to the falling cost and increasing availability of mobile devices such as e-book readers, tablet computers and smart phones. These devices come closer to replicating affordances of paper documents such as portability and, in the case of e-book readers, pleasantness of reading. However, there are a number of requirements that need to be met before they will be able to replace paper completely. Golovchinsky (2008) outlines several requirements of knowledge workers (e.g., reading, annotating and comparing between books) that need to be met by such devices before they will be important to that group. Until devices are able to meet these requirements it appears that paper will continue to play a role in many offices.

2.1.1 Problems with paper

Despite the benefits of paper, as described above, there are numerous problems associated with paper documents. These problems generally lead to paper documents being regarded as inferior to electronic alternatives. Sellen and Harper (2003) divide these problems into three classes, which are described below.

Symbolic. This class of problems is not directly due to paper itself, but rather people's perception of paper. To many people, paper is seen as old-fashioned and out-dated, while paperless offices are considered modern and futuristic.

Cost. Compared to electronic documents, paper has higher costs for delivery, storage and retrieval. Cost problems are a significant factor in causing paper to be viewed as inferior to electronic alternatives. Cost can be considered not only in monetary terms but also in terms of time and effort required to carry out a given task.

Interactional. These problems are due to limitations in terms of interaction with paper. For example, paper occupies physical space, requires physical delivery, and cannot be easily revised, reformatted or replicated. It is important to note, however, that these limitations can also be viewed as benefits depending on what the purpose of a given paper document is (e.g., the difficulty of revising paper documents makes them valuable where a permanent record is required, such as for legal or administrative reasons). A significant body of research has been dedicated to investigating ways of overcoming the interactional limitations of paper, and several of these are discussed in the next chapter.

Possibly the biggest problem with paper documents is the difficulty of managing them in comparison to electronic documents. These management difficulties are discussed in more detail in Section 2.3.

2.2 Document Management

This section presents an overview of document management. It begins by defining several terms used within this thesis, followed by an overview of the differences between personal and shared documents, which are then discussed in more detail.

2.2.1 Information, documents, content, and records

As previously stated, the focus of this thesis is *document* management. However, there are significant bodies of literature that discuss *information* management, *content* management, and *records* management. As there is some inconsistency in the definitions of these used in the literature, this section, therefore, provides definitions of the terms, and the relationships between them, as used in this thesis.

Information. Boardman (2004) quotes Feather and Sturges's (2003) definition of information as "*an assemblage of data in a comprehensible form capable of communication and use*", and based on this gives the following definition for information: "*any assembly of data which carries some meaning for one or more people*". This is the definition of information that is used in this thesis. As this definition makes no mention of physical or electronic form, information can exist in either form.

Documents. The definition of document used in this thesis is derived from the definition given by Henderson (2009), who used it in describing personal document management. Her definition comes originally from the International Telecommunication Union's Open Document Architecture standard,which she quotes as "*a structured amount of information intended for human perception, that may be interchanged as a unit between users and/or systems*" (Henderson, 2009, p. 8). Henderson notes that, since there is no mention of physical format, this definition covers both physical and electronic documents. Furthermore, a document need not be textual, and it may be static or dynamic. There is some variation in the literature as to whether documents are structured or unstructured, and therefore in this thesis documents may be structured or unstructured amounts of information.

Content. This term refers to any recorded information, whether it is stored for computer access (e.g., in a database) or for human perception (i.e., documents). It is a term that is commonly used by commercial vendors when describing their software (see Section 2.2.4 on Enterprise Content Management below). Although content can be physical or electronic in form, the term is mostly used to refer to electronic content.

Records. Records are defined by Strong (2008) as "*a subset of all content—structured and unstructured—that are generally governed by legal or tax record-keeping requirements*". This term is also commonly used in describing commercial software.

In summary, a document, as referred to in this thesis, is an amount of information that is in a form for human perception and can be interchanged as a unit. Thus, paper documents, books, and electronic files are all examples of documents covered by this definition. Content refers to any recorded information, whether or not it is primarily for human use, while records are content that is kept for legal or administrative reasons.

2.2.2 Personal vs. shared documents

Having defined what is meant by the term document in the previous section, it is now possible to discuss management of documents. Research in the field of document management, and the broader super set of information management, can be divided into two main categories: that which concerns the management of documents and information exclusively for personal use, and that which deals with management of documents and information in shared environments.

2.2.3 Managing personal documents

Personal Document Management (PDM) is described by Henderson (2009) as "*the activity of managing a collection of digital documents performed by the owner of the documents*". Although Henderson's definition of PDM is constrained to electronic documents, it could also be applied to paper documents.

It is worthwhile to consider how individuals manage personal documents, even in cases where organisations have shared libraries of information. This is because even when there are shared libraries of information, individuals will still often have their own personal information library. For example, in their diary study of document usage in engineering, Wild et al. (2010) observed that "*a significant portion of documents are still being sought within personal file stores*".

Personal Information Management (PIM), has been described as "*the practise and study of the activities a person performs in order to acquire or create, store, organise, maintain, retrieve, use, and distribute the information needed to complete tasks (work-related or not) and fulfil various roles and responsibilities*" (Jones, 2007). In this context, the meaning of *personal* information is "*information owned by an individual, and under their direct control*" (Boardman, 2004). As such, personal information is not necessarily private, but rather information that is kept for personal use (Lansdale, 1988).

Barreau (1995) discusses PIM systems in terms of five functions: *acquiring* information, *organising and storing* information, *maintaining* the system, *retrieving* information from the system, and *producing output*. Boardman (2004) has adopted and refined this list to *acquisition*, *organisation*, *maintenance*, and *retrieval*. Further adapting this for document management, Henderson (2009) identifies the following seven activities as the basic activities that document management systems need to support: *creating/acquiring*, *deleting*, *organising*, *finding*, *reminding*, *activity logging*, and *versioning/synchronising*.

2.2.4 Managing shared documents

While it is important to consider the management of personal information by individuals, this is only a subset of information management. There is also a need to manage information within organisations where information may be shared among groups of people, or even across the whole organisation. In this context the definition of information management given by Wilson (2003) is appropriate: "*the application of management principles to the acquisition, organisation, control, dissemination and use of information relevant to the effective operation of organizations of all kinds*".

Similarly, Boyette et al. (2005) state that document management "*refers to the process of capturing, storing, sorting, codifying, integrating, updating and* *protecting any and all information*". Likewise, Strong (2008) describes electronic document management tools as being "*focused on the capture, indexing, storage, and management of electronic documents*".

There are several other topics that relate to the management of shared documents and are worth discussing, among these are: knowledge management, enterprise content and records management, and information retrieval, which are discussed in the following sections.

Knowledge management

Knowledge management comprises document management and some method of recording tacit knowledge (see Henderson, 2009). As such, while the document management element component of knowledge management is discussed in this thesis, knowledge management as whole is not explicitly considered.

Enterprise content and records management

A common set of terms used to describe information in a commercial sense are based around the term enterprise. These terms include *Enterprise Content Management (ECM)*, *Electronic Records Management (ERM)*, which when combined lead to *Enterprise Content and Records Management (ECRM)*. The evolution of these terms mirrors the evolution of the systems themselves as they have continuously grown to deal with information in a wider range of electronic forms (Strong, 2008).

ECM tools, as their name suggests, support the management of content. *Records Management (RM)*, and its subset *ERM*, is concerned with the management of records, which must be managed carefully due to legal or administrative requirements. Finally, *ECRM* is the combination of ECM with support for records. There is some flexibility in how the terms are used, and ECM is often used interchangeably with ECRM.

Porter-Roth (2006) identifies five basic stages in the life-cycle of a document: capture, manage, store, deliver, disposition. Within the management stage, Porter-Roth lists four sub-tasks: metadata input, secure access, indexing, and versioning. These five stages are similar to the major components of ECRM identified by Kampffmeyer (2006), which are: *capture*, *management*, *storage*, *delivery*, and *preservation*.

Information retrieval

Information retrieval (IR) "*deals with the representation, storage, retrieval, organization of, and access to information items*" (Baeza-Yates and Ribeiro-Neto, 1999). Boardman (2004) argues that PIM is a higher level activity than IR, and makes use of IR in two of its sub-activities (acquisition and retrieval, previously mentioned in Section 2.2.3). For this reason, information retrieval is not directly considered in this thesis.

2.2.5 Framework for discussion of document management

The previous sections gave an overview of the major areas of interest with regard to document management. Based on literature from these areas, five topics have been identified which form a framework that is used for the discussion of document management in this thesis.

Organisation concerns the methods people employ to organise documents. This includes such sub-topics as categorisation of documents.

Maintenance comprises the tasks that need to be performed in order to keep document collections organised and relevant, such as purging, version control, and maintenance of indexes.

Control includes the tasks that involve monitoring and controlling access to documents, such as access control, logging, and tracking.

Storage involves the methods of storing documents. For physical documents this includes how the documents are physically stored, for example in filing cabinets, folders, or archival boxes. For electronic documents storage methods may include the file system, a document management system, or email. This also includes factors such as storage duration.

Retrieval regards the methods used for retrieving documents from the collection. This includes, for example, locating and delivering the document.

2.3 Review of Document Management

The previous section introduced and defined document management. This section provides a review of the literature surrounding document management as it pertains to the five topics identified above: organisation, maintenance, control, storage and retrieval.

2.3.1 Organisation

Organisation is an important facet of document management. It has a direct impact on how easy it will be to find and retrieve documents in future. This section reviews a number of studies that have investigated the ways in which people manage their documents, both in paper and electronic form. It begins with a discussion of two different organisational strategies: filing and piling. This is followed by a discussion of the difficulties that are commonly encountered when trying to classify documents, and how people deal with situations where they need to classify a document into multiple categories or wish to delay classification. The section ends with a discussion of the different levels of information people work with.

Filing and piling

In his well known and frequently cited article, “How do people organize their desks?” Malone (1983) describes a series of ten interviews that he conducted in order to gain a better understanding of how people organise information on their desks and around their offices. Each interview included a tour of the participant’s office, after which they answered a set of questions about their desk organisation and how they felt about it. In addition, six of the participants were afterwards asked to find several documents (chosen by co-workers) in their office, so that the processes they used to find them could be observed.

Based on this study, Malone made two primary claims: firstly, that the organisation of documents within a person’s workspace serves not only to assist finding these documents, but also to act as a reminder; secondly, that the cognitive difficulty of classifying things is a significant contributor to how people organise their desks. Additionally, Malone observed two main strategies for how people organise documents in their workspace. He defined *filing* as an organisation where the elements, such as folders, are explicitly titled and arranged in some systematic order, while *piling* is where the elements are not necessarily titled or arranged in a systematic order.

A similar study was conducted by Case (1986), which identified that people keep things in piles on their desks rather than filing them because "*things tend to get lost in filing cabinets*" (Case, 1986). Case also points out that "*perhaps certain materials would be better in a form that is more easily stored and retrieved than their current format*", suggesting, for example, that printing principal authors' names on the spines of journals or colour coding by date would allow them to be easily viewed at a glance.

These findings were mirrored by Cole (1982), who studied filing systems in thirty offices in order to investigate how people organise information in paper-based offices and "*gain insight into their concepts of office information and its representation in their memories*". He found that "*filing systems were not rigidly maintained; rather organisation was idiosyncratic, dependent on the demands on the users and their motivations*", and that "*users prefer to spend as little time as possible actually filing*".

Clearly organising documents is not a simple task. One of the major difficulties when organising documents is classifying them, which is discussed in more detail in the next section.

Classification difficulties

Dumais and Landauer (1983) describe two classes of problems that make classification difficult. The first class of problems concerns category names: these may be inaccurate, hard to understand, or interpreted differently by their creator and other people who use them. The second class of problems is due to the nature of categories in the physical world: it is not usually possible to create clearly defined, non-overlapping categories to which items can be unambiguously assigned.

Such difficulties with classification were observed by Malone (1983), who found that the significant cognitive difficulty involved in filing documents is largely due to the difficulty of assigning them to categories. As one of his participants stated: "*the hardest problem for me organizationally is deciding what the categories are and what category something is in*".

The difficulty of classifying documents is exacerbated by the fallibility of human memory. Whittaker and Sidner (1996) note that users who organise their email into a large number of folders have to be careful in order to remember what the classifications meant in order to avoid duplication. Similarly, one of the participants in Gordon's (1997) study, who was a secretary, noted frustration at trying to file documents for her manager, who was not consistent with the way he categorised documents: "*he picks a different name every time [for the same topic] . . . There may be 10 files in there on the same topic under different file headings*".

This difficulty was also observed by Whittaker and Hirschberg (2001) in their study of personal paper archives. They found that "*filing does not always guarantee easy access to information*", and that as the complexity of their archives increases people forget categories they have previously created, leading to duplicates and incomplete retrieval later on. Moreover, as more documents are filed in a category there may be a shift in the meaning of that category.

Multiple and delayed classification

Given the above problems with categories, it is therefore not uncommon that an item will belong to multiple categories. For example, one of the participants in a study conducted by Case (1991) stated: "*So often a piece of material falls into two or three categories. You can't constantly take it out of one and put it in other files, but you can make duplicates of it and the computer is going to be of value in doing that kind of thing*". This participant's view on this matter was not unique as, for instance, another participant would sometimes file duplicates of a given document in additional folders because that document could belong in any one of them.

Another difficulty when organising documents is that it is not always clear at the time of filing how a document should be classified. This was observed by Case (1991), one of whose study participants stated that at the beginning of a project he would not necessarily be sure of what categories to use, so instead he would file in alphabetical order initially and create categories as he went along.

This desire to delay classification is not unique to paper documents. In a study of email management conducted by Whittaker and Sidner (1996), it was found that one reason users would not sort their email into folders was that they wanted to postpone judgement because the value of the message was not immediately known. As a participant in this study stated: "*if information seems like it eventually will be overcome by events, I'd be very loath to move it into a [folder]*".

In fact, in some regards paper documents support delayed classification better than electronic documents because they can be piled. Leading Malone (1983) to suggest that electronic filing systems should support deferred classification and allow items to be piled without needing to be explicitly given a name or categorisation. However, the major limitation of piling is that as the size and number of piles increases their usefulness decreases, because it becomes harder to find things in them.

Levels of information

Cole (1982) identified three levels of information that his participants worked with:

action information, information that is currently being used or is soon to be used;

personal work files, closely kept, relevant material that is not currently being used; and

archive storage, information stored long-term.

Nardi et al. (1994) identified a similar division of information in their investigation of how fifteen computer users filed and found information:

Ephemeral information is information that has value for a limited time and is rarely filed away.

Working information has a longer life than ephemeral information, but is also relevant to the current work needs of users. Parts of this information may be archived or kept for longer.

Archive information is stored information that is not directly relevant to the current work of the user. This is generally stored in a structured archive and infrequently accessed.

A similar three tier model is presented by Sellen and Harper (2003, p. 132), who describe how the files used in an office they studied could be divided into three types:

hot files for active relationships with suppliers;

warm files for recently dealt with relationships, that may need to be dealt with again; and

cold files for those long dealt with, and largely kept "*just in case*".

It is important to take these levels of information into consideration as interaction with documents can vary greatly depending on which level they belong to. For example, Sellen and Harper (2003) describe how hot files required portability and needed to be quick to hand; warm files had a physical presence that served to remind, and they needed to be able to quickly move from storage (e.g., in a desk drawer) to active (hot) use; finally cold files have been "*put to bed*" and their meaning is gradually lost as knowledge of them declines.

2.3.2 Maintenance

The tasks involved in maintaining a document collection are ongoing, and necessary in order to keep the collection organised and relevant. Organisation, as discussed above, is itself partly a maintenance task (e.g., in terms of maintaining organisational categories and hierarchy) but was discussed on its own due to its significance.

An important part of maintaining a document collection is purging out-of-date and unnecessary material. This reduces storage space requirements and makes it easier to find useful documents by removing documents that are no longer relevant. The policy surrounding document destruction is important as the destruction of documents can have serious ramifications, for instance in the case of legal documents, as was evidenced in the case of McCabe vs. British American Tobacco Australia Limited (see Newman, 2004).

Another element of maintenance is the management of revisions of the same document. A single document may go through a number of draft stages before being completed, and it may be necessary to keep track of these for legal or administrative reasons. Therefore, there needs to be a mechanism in place for maintaining these various revisions. Although electronically this can be achieved easily through version control software, in the case of paper documents it can be more complex and costly (e.g., in terms of keeping track of and storing revisions).

Maintenance difficulties

A common thread of the studies into document management, especially those in which the participants were managing paper documents, is that the overhead of time required to manage their collections of documents is a significant problem. For example, a study of personal indexes conducted by Jahoda et al. (1966) showed that 42% of their participants found the time taken to prepare their index was too long.

Similarly, Whittaker and Hirschberg (2001) note that due to the large overheads associated with the construction and maintenance of paper-based filing systems, some people require secretarial assistance to maintain such systems. Part of the problem is the mechanical difficulty involved when, for example, creating new categories. Malone (1983) found that the mechanical effort involved in creating new categories is often non-trivial, and can act as a major barrier to people keeping their filing systems up-to-date.

As well as being time consuming, maintaining document collections is often considered mundane, boring, or unenjoyable. For example, Cole (1982) noted "*a general lack of motivation towards the upkeep of elaborate filing systems—in fact, the less time spent filing the better*".

The time consuming and generally boring nature of maintaining document collections can lead to them being maintained sporadically, with action only being taken when triggered by some event, for example moving office. Such a lack of maintenance can lead to the collection becoming out-of-date or full of irrelevant material, such that its usefulness is decreased.

2.3.3 Control

Another facet of document management is monitoring and controlling access to documents. While this is undoubtedly of importance when dealing with documents of a sensitive nature, such as patient records in a hospital, it can also be valuable in preventing loss or misplacement of more mundane documents. For example, in their observation of a patent office, O'Neill et al. (2006) observed that the manager of invention disclosure (ID) documents would keep "*a hawk-like watch on these files*" and would try "*not to let anyone into the ID room because otherwise files go missing or get disorganised*".

2.3.4 Storage

Storage is an implicit part of having a document collection. Documents must be kept somewhere, whether it is a physical or virtual location. Storage is another aspect of document management where paper documents suffer in relation to electronic documents. Sellen and Harper (2003) identify storage as one of three key costs for paper document management systems. They give as an example the cost of filing cabinets for two million documents in paper form, which runs into tens of thousands of dollars, while the same documents could be stored electronically on fewer than ten CD-ROMs (see Sellen and Harper, 2003, pp. 27–28).

The problem of storage can be exacerbated by inefficient organisation and maintenance. For example, filing multiple copies of the same document under different categories in order to simplify the process of finding them results in an increase in the amount of storage space required. Alternatively, it may be possible to store only a single copy of the same document if a better means of locating it were available. Similarly, if the collection is not purged of irrelevant documents, this may, over time, increase the required storage space unnecessarily.

2.3.5 Retrieval

There is, of course, little use in storing documents if they are unable to be retrieved when required. Therefore, one of the primary reasons for organising and maintaining document collections is to facilitate document retrieval. As such, problems with organisation and maintenance often result in greater difficulty retrieving documents. This was found, for example, by Whittaker and Hirschberg (2001), who noted that inconsistent classification of documents within an archive can result in incomplete retrieval later on.

Four step process

Jones (2007) describes finding personal information that one has previously stored as a four step process. While Jones was referring to finding personal information, these four steps are equally applicable to finding shared documents. The four steps, listed below, may be performed once or repeated multiple times in order to find one or more documents.

Remembering to look. Before one can find a document, one has to remember to look for it. At this stage the organisation of documents to remind, as described by Malone (1983), may play a part.

Recalling information about the document (e.g., which pile of papers it is in, which folder it is in, or some key words) in order to narrow the search. The more information a person can recall about the document they want to retrieve the easier it will be find it. This was reflected in the study of Cole (1982), who noted that the more familiar someone is with an archive, the easier it is for them to find things within it.

Recognising the required document.

Repeating the process until all appropriate documents are collected. However, it is possible that this process will not be repeated a sufficient number of times. Gordon (1997) notes that searches may be concluded prematurely and action taken on the mistaken assumption that all the relevant information has been found.

Orienteering and teleporting

Two contrasting search strategies, referred to as orienteering and teleporting, have been identified by Teevan et al. (2004). These strategies were identified based on a study they conducted which investigated how people search in their email, electronic files, and on the Web. This study comprised 151 semi-structured interviews with 15 participants, discussing recent search activity of the participants.

Searches following the orienteering strategy are conducted in a series of small steps. Each of these steps, is in effect an iteration through the four steps listed above, and these may be repeated multiple times to find a single document. For example, one of the participants 'knew' the document she was looking for and the folder it was stored in, but could not describe them, therefore she performed a number of iterations through the four steps with the local context aiding the remembering step at each iteration. Browsing through folders to find a document is an example of orienteering.

Alternatively, searches following the teleporting strategy are conducted in one large step. However, Teevan et al. note that a teleporting strategy was rarely taken by their participants, and keyword-based searching was usually performed as a step in an orienteering search. Similarly, Malone (1983) notes that when retrieving documents, the description of a document to be retrieved will often be specified in multiple dimensions, for example "*a message from M.A. Smith, last week, about a meeting in Palo Alto*".

The ability to perform teleport-type searches can have an influence on the organisational strategy used. For example, searching using metadata elements may enable the use of a flat hierarchy, thereby removing the need to arrange documents into folders.

2.4 Discussion

Despite numerous predictions, the paperless office has not yet become a reality, and paper documents continue to be used in most modern offices. One reason for this is the fact that paper has certain affordances due to its physical, tangible nature, and these have not yet be replicated satisfactorily by electronic documents.

However, while paper documents have some advantages in comparison to electronic documents, they also suffer from several limitations, especially with respect to their management. For example, electronic documents generally cost less to store, retrieve, and transmit than paper documents. They also take less space to store, can be grouped and sorted faster and easier than paper documents. Electronic documents can have metadata stored with them (which can be automatically generated, e.g., document creation time), whereas paper documents can only be sorted along a single dimension (e.g., sorted by last name in alphabetical order) unless additional information such as an index is kept, which requires extra effort to create and maintain. Therefore, as Frohlich and Perry (1994) note, advantages of electronic documents over paper are more to do with organisation of documents than the use of the documents themselves.

It seems, then, that for the foreseeable future offices will continue to host an ecosystem that includes both paper and electronic documents. This in turn leads to fragmentation of information between paper and electronic forms. A pertinent description of this type of information fragmentation is given by Jones (2004):

> *But this diversity has become part of the problem leading to in-*

> *formation fragmentation. A person may maintain several separate, roughly comparable but inevitably inconsistent, organizational schemes for electronic documents, paper documents, e-mail messages and Web references. The number of organizational schemes may increase if a person has several e-mail accounts, uses separate computers for home and work, uses a PDA or a smart phone or uses any of a bewildering number of special-purpose PIM tools.*

Problems due to information fragmentation include the need to maintain separate organisational systems for managing paper and electronic documents, which can be more costly in terms of requiring more time, and can be less than satisfactory, for instance by leading to future problems when retrieving documents, such as difficulty of finding required documents and incomplete retrieval of all relevant documents.

One method of overcoming this problems is to bring the affordances of paper documents to electronic documents, such that paper documents can be phased out. Examples of research towards this end include Rosner et al. (2008) who investigated paper use in order to guide design of electronic technology, and Liesaputra et al. (2009) who aim to bring the look and feel of paper books to electronic books. However, while such research works towards improved user interaction with electronic documents, the complete removal of paper from the office—if at all possible—is still a long way off. Furthermore, such developments do not provide a solution for existing legacy paper documents that may be too expensive or time consuming to digitise.

An alternative approach to solving this problem is to bring the benefits of electronic document management to paper documents, and investigate systems that better integrate the management of paper and electronic documents. This is the approach that is taken in this thesis.

2.5 Summary

This chapter began with a discussion of paper documents in present day offices (Section 2.1), which identified that, although they have a number of shortcomings, it appears likely that paper documents will play a role in offices for the foreseeable future due to their affordances. Section 2.2 provided definitions for several terms that are used in this thesis, and identified five topics that provide a framework for discussing document management: organisation, maintenance, control, storage, and retrieval. A review of literature relating to document management was provided in Section 2.3. This section discussed a number of previous studies of document management, and from these it is clear that this can be a challenging task. One of the issues that is encountered in offices work with paper and electronic documents is fragmentation of information between the paper and electronic forms (as discussed in Section 2.4). There is, therefore, a need to investigate systems that integrate paper and electronic electronic document management. The next chapter provides a review of existing technology and previous research towards integration of paper and electronic document management.

Chapter 3

Related Work

The previous chapter discussed some of the significant problems encountered with document management, especially with regard to paper documents and fragmentation of information. The discussion at the end of the chapter noted that the problems of paper document management and fragmentation of paper and electronic document management could be overcome or reduced by providing better integration of paper and electronic document management.

This chapter begins with a discussion of the divide between the physical and virtual worlds, and research towards bridging this divide in Section 3.1. This is followed by a review of existing technology and systems that work towards integrating the management of paper and electronic documents (Section 3.2). The chapter concludes with a discussion (Section 3.3) and summary (Section 3.4).

3.1 Bridging the Physical and Electronic Worlds

Information fragmentation is a problem that results from information being kept in different forms and managed with different tools. As such, information can be fragmented across a range of electronic tools as well as between paper and electronic forms. However, the fragmentation of information between paper and electronic forms provides an additional challenge in comparison to fragmentation of information in purely electronic form, since in this case the fragmented information is split over the divide between the physical and virtual worlds. This division between the physical and virtual worlds is an important issue in computer science research, as noted in the following quote from the renowned interaction designer and researcher Bill Buxton (2002):

> *One of the most significant issues confronting computer users is the problem of bridging the gap between the physical and virtual worlds. For most activities, most current systems make it too difficult to move the artifacts back and forth between these two worlds, the physical and virtual. Hence, the relevant documents, designs, and so on are isolated in one or the other, or split between the two.*

The divide results from the fact that people and physical objects exist in the physical world, whereas electronic content exists in the virtual world. These are two separate worlds, and the divide between them must be crossed in order for interaction to occur between them. The divide can be crossed in two directions: from the physical to the virtual (input) and from the virtual to the physical (output). Since users exist in the physical world, some form of interface between the worlds is required in order for them to interact with the virtual world. Figure 3.1 shows how the divide is crossed for a typical GUI. In this case the crossing of the divide is provided by, for instance, the keyboard and mouse (in the input direction) and monitor (in the output direction).

The problem with crossing the divide between physical and virtual in this manner is that the virtual world effectively exists 'in the box'. That is, from the perspective of the users, the virtual world exists inside the computer, because the only way of interacting with the virtual world is through the computer.

Overcoming this problem is not a simple matter and there are several fields of research that approach the problem from different angles. Significant among these are:

- augmented reality;
- ubiquitous computing; and
- graspable, tangible and embodied user interfaces.

These fields are introduced in more detail below.

3.1.1 Augmented reality

Research into Augmented Reality (AR) seeks to bridge the divide between the physical and virtual worlds by overlaying electronic content on the physical world. Azuma (1997) defines AR systems as having three characteristics:

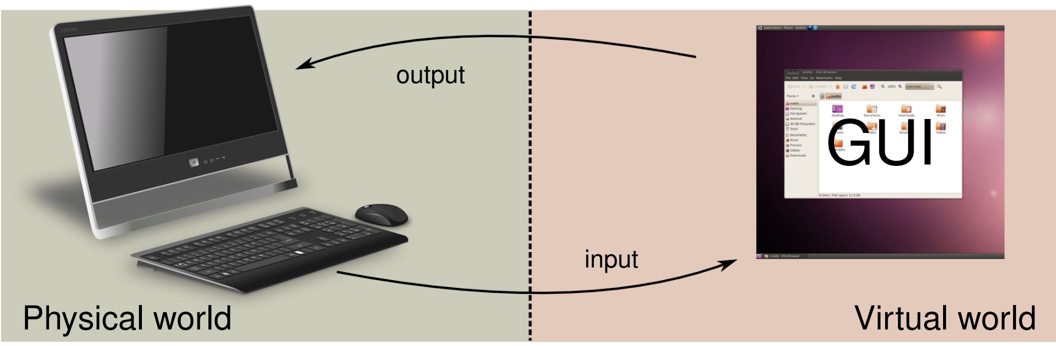

Figure 3.1: Crossing the divide between the physical and virtual worlds with a typical GUI.

- they combine physical and virtual;
- they are interactive in real-time;
- and they are registered in 3D.

Augmented reality requires crossing the divide between the physical and virtual worlds in both directions: physical-world objects are registered in the virtual world, and virtual objects are displayed in the physical world. This is shown in Figure 3.2. Numerous approaches to registering physical-world objects can be taken, but a method commonly used is visual tracking using a camera and printed tags (e.g., ARTags Fiala, 2005)—visual tracking is described in more detail in Section 3.2.2. The display of virtual objects, overlaid on the physical world, is usually achieved through some form of wearable or carried device, such as a Head-Mounted Displays (HMDs), or by projecting electronic content into the physical-world environment (Silva et al., 2003).

There are several significant limitations to augmented reality systems. In order to display virtual content, either a wearable display or a projector is typically used. Wearable displays require that everyone who is to interact with the system wears such a device, because those who do not wear a device will not be able to see the virtual elements. Systems that use a projector for display have limitations in terms of the coverage area of the projector. Once virtual objects move out of the coverage area of the projector they will no longer be displayed. Additionally, the beam of the projector can be blocked by the users interacting with the system, thereby hiding the virtual objects.

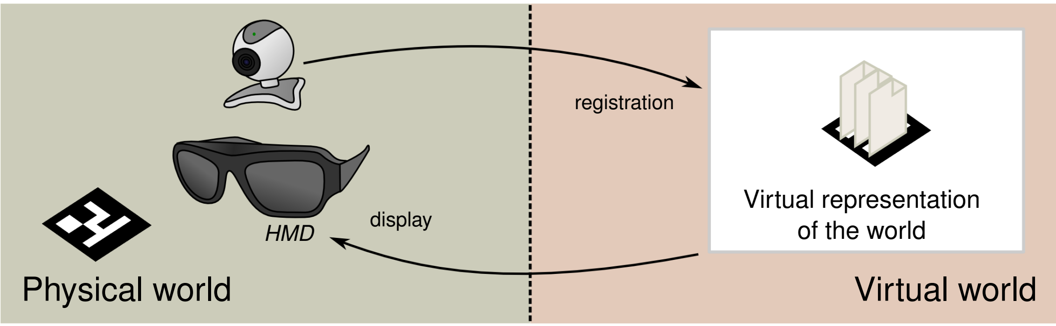

Figure 3.2: Crossing the divide between the physical and virtual worlds with Augmented Reality (AR). Physical-world objects are registered in the virtual world, and virtual objects are overlaid on the physical world.

3.1.2 Ubiquitous and pervasive computing

The term ubiquitous computing was introduced by Weiser who states that "*the most profound technologies are those that disappear*" (Weiser, 1991). He presents the example of writing, which has become so ubiquitous that we are rarely consciously aware of it. The field of ubiquitous computing aims to make computers, like writing, so integrated into our lives that they "disappear".

Ubiquitous computing goes beyond simply making computers more mobile so that they can be taken everywhere. It involves splitting the computer up, so that it becomes a number of simpler task specific devices rather than a single monolithic device. Buxton (2002) describes a number of limitations suffered by monolithic *superappliances* as compared to separate, task-specific appliances. These limitations of superappliances include: greater complexity (largely due to the extra layer needed in order to switch between modalities); being usable only in a single location at a time, whereas a number of task specific devices can be distributed; typically being usable by only a single person at a time, while a number of task-specific tools can be used simultaneously by different people; functions being time-multiplexed, whereas a number of functions can be performed simultaneously when using tasks specific tools.

Interesting examples of ubiquitous computing technology include the *active badge*, which wirelessly broadcasts the identity of its wearer (Weiser, 1991), and *context-aware kitchen utilities*, a knife and cutting board augmented with sensors that are able to differentiate between different foods being cut (Kranz et al., 2007).

This ubiquitous computing approach to bridging physical and virtual can be applied to the management of documents. In this case, the physical document management infrastructure would be augmented with electronic functionality that provides input and output. This would allow a system to, for example, track the movement of documents, determine their current locations, and present information back to the user in a way that is integrated with the physical document management system.

3.1.3 Graspable, tangible and embodied user interfaces

As computing technology evolves away from superappliances towards more task-specific devices, often the ways in which users interact with these devices also evolve. Just as GUIs provided a more natural interaction paradigm than command line interfaces, so too do graspable and tangible user interfaces provide a more natural interface than GUIs.

Graspable user interfaces (e.g., Fitzmaurice et al., 1995) provide interaction through physical objects that are tightly coupled with virtual objects. Manipulation of these physical objects (such as rotation or movement) results in corresponding manipulation of the corresponding virtual objects. Graspable interfaces allow more direct interaction with virtual objects than otherwise available with a graphical user interface. Additionally, interaction with graspable interfaces is space-multiplexed rather than time-multiplexed. In other words, each virtual object has its own input device in physical space, as opposed to graphical interfaces where interaction is typically carried out using a common set of devices. This, for instance, affords two handed interaction, or allows multiple users to interact simultaneously with the virtual world.

The related concept of Tangible User Interfaces (TUIs) was introduced by Ishii and Ullmer (1997) as a way of achieving more seamless interaction between the physical and electronic worlds. Similar to graspable interfaces, TUIs extend interaction beyond conventional input devices such as the keyboard and mouse. However, TUIs take this even further, incorporating physical interactions occurring in the foreground with peripheral background interactions, so that the physical world itself becomes an interface to the virtual world.

Examples of TUIs include mediaBlocks (Ullmer et al., 1998), wooden blocks with an embedded electronic identifier that serve as physical icons of electronic content; and VoodooIO (Villar and Gellersen, 2007) a “malleable control surface" that can be dynamically altered to suit the needs of the user (discussed in more detail in Chapter 6).

Embodied user interfaces are similar to tangible interfaces in that interaction is through physical manipulation of objects. However, the point of difference between the two is that where the tangible interaction is reflected on the computer system, with an embodied interface the virtual representation is part of the object that is manipulated. Fishkin et al. (1998) believe that embodied interfaces are the next evolutionary step beyond tangible interfaces towards the invisible user interface. For example, Fishkin et al. (2000) describe an embodied interface for the task of reading a book that allows the user to turn pages with physical flicking motions.

In recent years the line between tangible and embodied interfaces has become blurred, and therefore in this thesis tangible and embodied interfaces are considered to be the same.

3.2 Integrating Management of Paper and Electronic Documents

This section describes several of the ways in which paper and electronic documents and their related infrastructure can be linked, and provides examples of a number of systems that take these different approaches. It is divided up into several sections based on the approaches taken to linking paper documents (and other artefacts) to electronic information. It begins with a review of systems that manually maintain a link between paper documents and their corresponding electronic content. This is followed by a discussion of systems that automate the linking between content in the physical and virtual worlds, first looking at those that use visual tags and then those that use wireless tags to achieve this, then other approaches that can be taken.

3.2.1 Manual systems

The simplest approach in terms of technology is to manually maintain the link between the physical and virtual worlds. In this case the user is responsible for keeping the electronic system up-to-date with information from the physical world, and making sure that the physical world matches the view held be the electronic system.

An example of a filing system using a manual link between paper documents and an electronic index is described by Wilson (2001). In this system, paper documents were given a unique serial number and stored in numerical order in a central repository. The same serial number would then be used as an index to electronic metadata and related electronic documents.

By integrating the management of paper documents with an electronic index this system reduces some of the difficulty of document organisation. This system mitigates the limitation of storage along a single axis that paper documents suffer from. Instead, the electronic metadata allows the identifier of a paper document to be located using the electronic index which, once known, allows the user to retrieve the document. This also simplifies the organisation and maintenance requirements of the paper-based collection since only a single copy of each document needs to be kept.

Another example of a system that manually maintains the link between paper documents and electronic metadata is the *self-organizing file cabinet* (Lawrie, 1997; Lawrie and Rus, 1999). This system comprises a physical world filing cabinet and software that assists users in organising the information in it. When documents are to be added to the filing cabinet they are first scanned and entered into the system's database. By comparing information in the scanned document to information in the database, the software is then able to suggest which drawer of the filing cabinet the document should be placed in, so that it is stored with documents of the same or similar topic. The database can be searched or browsed to locate documents, and is able to indicate to users through an on-screen visualisation the position of the documents within drawers of the filing cabinet. As with the system described above, this system reduces the effort required for storing and maintaining documents in a paper-based collection.

However, the problem with manually maintaining the link between the physical and virtual worlds is that there is no automation and the user is responsible for keeping the paper documents and electronic information associated with them synchronised. For example, with the self-organized cabinet, when documents are placed in the cabinet, they must be placed at the front of the drawer so that the software knows where they are. This lack of automation means that it would be difficult to keep the self-organized filing cabinet synchronised in practise, as there is no way of enforcing actions or detecting when incorrect actions have been taken. For example, if the user were to retrieve or return documents without using the software, or to misplace a document, the physical and electronic systems would lose synchronisation—a problem that would compound as more such mistakes are made. Therefore, while the manual approach is the simplest to implement, the lack of enforced synchronisation between the physical and electronic entities is not optimal. In order for such a system to work, discipline is required on the part of the users so that consistency is maintained.

3.2.2 Automated document tracking

A link between the paper documents and electronic content can be automated by attaching some form of machine readable tag to paper documents. The main types of tags that may be used are visual tags (e.g., barcodes) and electronic tags (e.g., Radio Frequency Identification or RFID tags). Using an appropriate tag reader it is possible to track the location of physical world documents and incorporate this information with an electronic document management system. Several approaches of this type are discussed below.

Document tracking with visual tags

In their patent, Meunier and Grasso (2002) describe a system in which paper documents are marked with an identifying barcode along their edge. This barcode can be used to identify documents for retrieval purposes, by scanning the document collection with a scanning device until the desired document is reached. Such a system would speed up the retrieval of documents as the reader would be able to identify documents faster than a human, especially when documents are stored so that only their barcoded edge is visible. Additionally, the faster identification of documents reduces dependency on the ordering of documents on the shelf for retrieval purposes, as a result reducing the effort required for organisation and maintenance. However, the primary limitation of this system is that scanning of documents is not automated and must be performed using a barcode scanner.

An example of a more automated solution using visual tags is PaperSpace (Smith et al., 2006). Rather than using 1D barcodes along the edge of documents, PaperSpace makes use of 2D barcodes which are embedded in the margins of documents as they are printed. A webcam monitors the users' desk, and the computer vision component of the software is able to identify visible documents on the desk by parsing their 2D barcodes. This information is passed to a tracking component that is able to determine the location and orientation of the documents on the desktop, and tracks their position as they are moved around the desk, even when placed in piles. This information is stored in a database that can be queried by application software, which enables software to assist the user in retrieving documents.

A similar system called TRIP is described by de Ipiña et al. (2002). TRIPtags are printed tags with an encoded identifier that can be used to track the location of the physical artefacts to which they are attached by processing the recored video stream from a camera. The authors describe an example application in which TRIPtags are attached to library books and shelves. A librarian with a video camera can walk around recording video of the books and shelves, and when the video is processed, the system is able update its database of books with the current locations of each book relative to the books around it. This reduces the effort required for maintaining such collections because the detection of misplaced books can be partially automated.

While vision-based systems have a number of benefits (e.g., low cost and ability to print tags directly onto documents), they also have several limitations. Chief among these is that line-of-sight is required between the tag and tag-reader in order for the tag to be read. For instance, if the document is visually obscured, rotated out of the plane of the reader, or if there is insufficient light to read its tag, it may effectively disappear from the system tracking it. This may happen if, for example, documents are stacked and then moved, or if a user moves a document while standing in between the document and the camera tracking it.

Document tracking with electronic tags

An alternative to using visual tags that overcomes the line-of-sight limitations that these impose is to use electronic tags. Electronic tags include some form of circuitry that allows them to communicate with a tag reader. Typically this communication is wireless, and at radio frequency. Tags that communicate at radio-frequency have the advantage of not requiring clear line-of-sight between the tag and the tag reader. RFID tags (Want, 2006) are a common example of such tags. RFID tags can be passive, in which case they do not require a power-supply and instead draw power inductively from the tag reader, or they may include their own power-supply (e.g., a battery), in which case they may be referred to as active, semi-passive, or battery-assisted passive depending on their design. RFID tags can be used to track the movement of artefacts such as documents.

MagicTouch (Pederson, 2001) is an RFID-based system that tracks the movement of artefacts based on the assumption that they only move when explicitly moved by the user. RFID tags are attached to artefacts and the user wears a device on their hand that can both read RFID tags and determine its current physical-world location. Thus, as the user picks up and moves RFID-tagged objects their location can be tracked. Information about artefacts (e.g., their location) is stored in a database that can be queried by higher level application software. This increases the control that the users have over their physical collections, allowing documents to be tracked and document access to be logged. It also assists with document retrieval by storing the current location of each artefact. However, the main limitations of MagicTouch are due to its assumption that documents will be moved one at a time, and will only move when moved by a user wearing an appropriate reader.

RFInD (Saxena et al., 2007) takes a different approach, using a single RFID tag reader with rotatable antennae and variable gain control. Tags are placed on fixed locations in the area in which artefacts are being tracked, and these act as reference points. Artefacts are tagged with RFID tags, and the system is able to determine for a given artefact tag which reference tag it is closest to. Based on knowledge of the positions of the reference tags it can then display approximate location information to users through a GUI. Similar to RFInD, Realfind (Câmara et al., 2008) also tracks artefacts. However, Realfind uses RFID readers mounted at various locations around the space in which artefacts are being tracked.

Both RFInD and Realfind provide similar advantages to MagicTouch in terms of allowing tracking of artefacts and making locating physical artefacts easier during retrieval. However, they have limitations in terms of the accuracy with which they can locate artefacts. While they may be able to determine that an artefact is, for example, on a certain desk or in a given pile, they are not able identify to its location with better accuracy.

Existing passive RFID location determination systems are limited in their accuracy. For example, the RFID-based location determination systems described by Heo et al. (2007) and Zhao et al. (2010) are limited to accuracy of one metre and twenty centimetres respectively, which is not enough to determine, for example, the ordering of books on a shelf. Although finer grained accuracy can been achieved using moving antennae (Hinske and Langheinrich, 2008), or by using active tags (Hekimian-Williams et al., 2010), these approaches are less than ideal for document management, since moving antennae require moving parts, which can be expensive and require maintenance, and active tags have a limited battery life or require additional power supply.

While the approaches discussed above were aimed more generally at tracking artefacts, there has been some research more specifically related to managing documents using RFID tags. Arregui et al. (2003) describe a system that involves tagging documents with RFID tags to link the paper and electronic representations of a document. The link between the paper and electronic instances of a document is created when the document is printed. Blank sheets of paper have RFID tags attached to them and these are read by a printer augmented with an RFID reader in order to associate the physical manifestation of the document with its electronic manifestation. Places in the environment where the documents are stored or used (e.g., filing cabinets, brief cases and desks) are augmented with RFID readers, enabling the tracking of documents. Additionally, the activity of individual users can also be tracked through wearable RFID badges.

Using RFID in this manner simplifies the process of locating a paper document for retrieval, because documents are continually tracked and their location is constantly updated in the document database when they pass within the range of RFID readers. Additionally, the system allows the entry of rules that can cause alerts to be generated; such alerts may be displayed on a computer or sent to a mobile phone or pager. Another use of this system described by Arregui et al. is detecting association between folders by identifying those that are often kept together.

An example of the practical application of this technology is described by O'Neill et al. (2006). They discuss how the system could be deployed in the patent department of a law firm. However, this application is described purely as a scenario of use, and no actual implementation or evaluation of the system is presented.

One of the limitations of RFID-based systems is the cost associated with manually tagging documents. However, a system that could potentially reduce the cost and difficulty of tagging documents by printing RFID tags directly onto paper using e-Ink technology is described by Hark et al. (2008). They also describe a database that could be used to store information associated with RFID-tagged documents, such as document name, version, and location.

A more detailed model for incorporating paper documents into an electronic document management system is described by AbuSafiya and Mazumdar (2004). Their model relies on using RFID technology, such as that described by Arregui et al. (2003), to provide a link between electronic and paper documents. This model takes into account the paper documents being managed, the physical location where are these are stored, and the organisational hierarchy of the enterprise where it is to be used. Using this model it is possible to perform queries to determine, for example, which documents are in a given location, what the current location of a document is, or which documents belong to a given person.

Similarly, Bodhuin et al. (2007) describe how RFID technology can be deployed in a business office to improve management of paper documents. Their system also integrates elements of work-flow management, for example by alerting users to documents with urgent priority, and the system aids maintenance of document collections, for instance by alerting users to documents which have expired and can thus be destroyed. This system also improves control over documents, since it can detect when a document is taken out of range and generate alarm.

Limitations

With the approaches described above, both those using visual and those using RFID technology, the complexity is in the tag readers and associated systems rather than the tags themselves. This has the advantage of relatively low tag cost. However, this means that the divide between the physical and virtual worlds is only crossed in one direction, the input direction. That is to say, while the electronic system is able to determine the location of tagged documents there is no complementary way of outputting information back into the physical world, and therefore this interaction is, in most cases, limited to traditional user interfaces (see Figure 3.3). Therefore, such document tracking systems alone do not provide a seamless integration between the physical and virtual worlds.

3.2.3 Other approaches

An alternative to the approaches described above is to provide more functionality within the tags that are attached to documents. An example of a system taking this approach is DigiClip (Decker et al., 2004), an augmented bulldog clip that can be attached to paper documents. The DigiClip has internal memory that allows it to store information about the attached document (e.g., its author, theme, keywords, and revision history). As DigiClip is able to communicate wirelessly, such information can be read by portable devices such as Personal Digital Assistants (PDAs). It is also able possible to inform the user when the paper and electronic documents lose synchronisation by lighting an LED on the DigiClip.

The limitations of DigiClip stem from the fact that it is an active device. As such, it requires its own power supply, which is provided by a battery, and

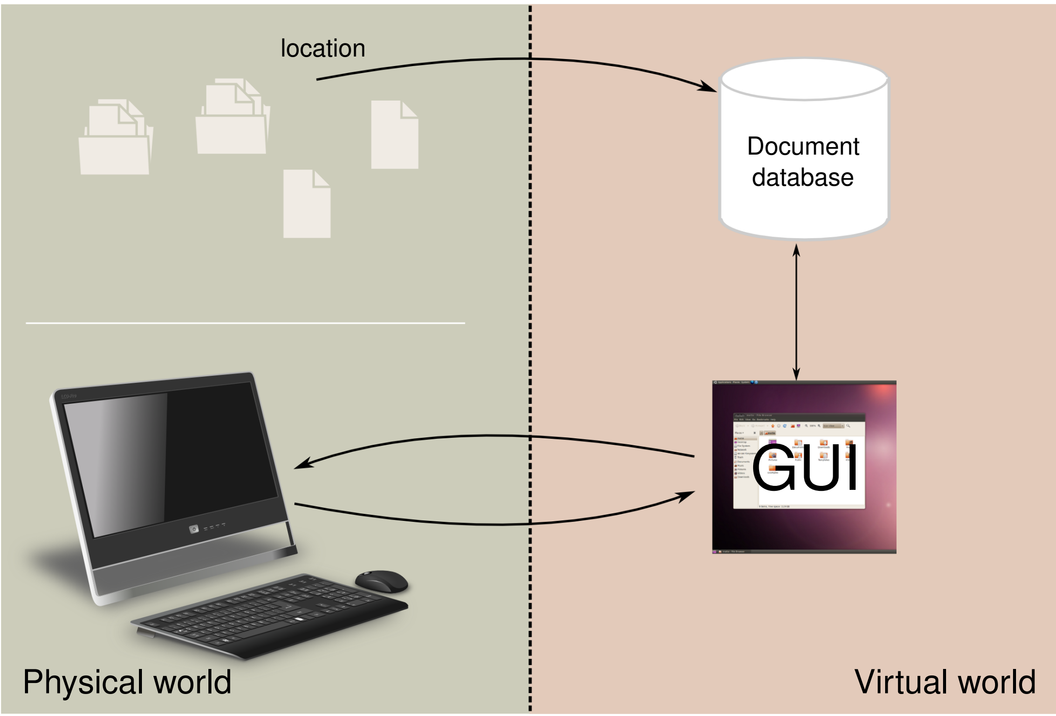

Figure 3.3: Tracking documents using visual or RFID tags allows location information to be inputted to the virtual world, but this information is still communicated back to users in the physical world via a traditional user interface.

this would need charging or replacing—the authors estimate that a lifetime of one year could be achieved. Also, the size, weight, and cost of the clip is large in proportion to the documents it is clipped on to. For these reasons it would be impractical to attach a DigiClip to every document in an archive. The DigiClip would be more suitable for hot or warm documents (as discussed in Section 2.3.1).

Raskar et al. (2004) describe a system that uses RFID tags with an inbuilt photo-sensor (which they refer to as RFIG, or Radio Frequency Identity and Geometry tags) in combination with a handheld projector. Because the beam of the projector can be sensed by the RFIG tags, the system is able to determine tag position by projecting a series of specially designed patterns and reading back from the RFIG tags what was sensed. Once the positions of the tags are known, an augmented reality style display can be projected over top of them. Figure 3.4 shows how the RFIG system crosses the divide between the physical and virtual worlds.

An example application of RFIG technology is to facilitate management of books in a library (Raskar et al., 2005). By tagging books with RFIG tags it is possible to determine the ordering of books on the library shelves, and project an overlay that can, for example, indicate when a book is filed out of place and show where it should be put.

However, the major limitation of the RFIG system is that it requires users to carry a handheld projector in order to interact with it. Currently such devices are still bulky and not practical for general use. Also, the number of people able to independently use such a system simultaneously is limited by the number of projectors available. Similarly, the system would likely be

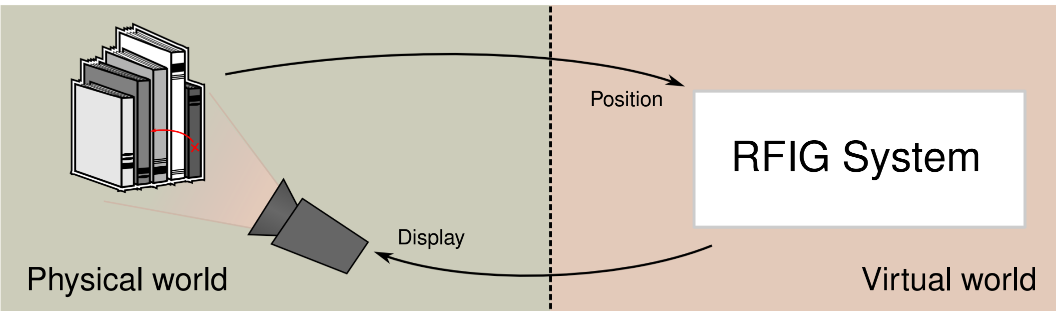

Figure 3.4: Crossing the divide between the physical and virtual worlds with the RFIG system.

unable to support multiple people using the system in close proximity, where the projected images would overlap and thus interfere with each other.

A system with a similar library application is described by Choi et al. (2006). This system uses RFID tags on library books and shelves in combination with a hand-held RFID tag reader to determine the position of library books. The system requires users to periodically scan the library shelves with the hand-held RFID reader so that it can build up a database of associations between book and shelf tags. When a book is to be retrieved, library patrons can be guided to the appropriate shelf by a PDA. Each shelf tag has an LED display associated with it, and this is used to show each patron the approximate location of the book they are looking for.

While this system provides integration between the physical books and the electronic catalogue, it suffers from two main limitations. Firstly, it requires the shelves to be periodically scanned with an RFID reader in order to build the shelf-book tag associations. Not only would this be a time consuming process, but also movements of books would not be detected until the next scan. In a library setting where books are constantly moved, this is a serious drawback to this system. The second limitation of this system is its coarse granularity, which only provides an approximate position. Since there would be a number of books associated with any given shelf tag, the user would still need to find the appropriate book among these.

3.3 Discussion

A number of approaches to linking between paper and electronic documents were described in the previous section. All of the approaches helped overcome the limitations of paper documents to some degree, by allowing electronic metadata to be associated with these documents. This assists with organisation and maintenance, and as a consequence retrieval of paper documents. Linking paper documents to electronic metadata removes, for example, the need to store multiple copies of a document, thus reducing the maintenance and storage overheads of the document collection.

However, the systems described above mostly focus on crossing the divide from physical to virtual (i.e., tracking physical documents and sorting information about them in an electronic document management system), and as such they only provide support for accessing this information through a traditional GUI. Even those that do provide user interaction beyond a GUI, such as DigiClip and RFIG, have limitations that would likely make them impractical for general office use.

In order to achieve an interface that is closer to Weiser's vision of an 'invisible' computer, the integration between paper and electronic document management can be taken further, by embodying the interface in the paper document management system itself. Including not only input (e.g., document tracking), but output as well (e.g., display of document location through visual means). This would provide an even stronger link between the paper documents and their electronic metadata. The rest of this thesis describes the development and evaluation of a new system for document management that focuses on crossing the divide between the physical and virtual worlds in both directions.

3.4 Summary

Providing better integration of paper and electronic document management requires crossing of the divide between the physical and virtual worlds. Traditionally this divide is crossed using conventional input and output devices, for instance a monitor for output and keyboard and mouse for input. However, the fields of augmented reality, ubiquitous computing, and graspable, tangible, and embodied interfaces investigate alternative ways to interact with electronic content.

Such approaches can be applied to document management, in order to provide better integration between paper and electronic document management as was discussed in the Chapter 2. A number of approaches were discussed in this chapter but, as pointed out, each of these approaches has their limitations. Most of the systems described provide only input in terms of document location, but do not allow for any form of output beyond that of traditional GUIs. Therefore, it is important to investigate development of a system that better integrates paper and electronic document management by embodying the document management interface in the physical document management system to provide seamless interaction across the divide between the physical and virtual worlds.

Part II

Identifying Requirements

Chapter 4

Study of Document Use in Offices

The previous chapters have identified the need for better integration between paper and electronic document management systems, and discussed several existing approaches that have been taken to achieve this. In order to design such systems, an understanding of how paper and electronic documents are used, and the purposes they serve, is required.

This chapter discusses the design, execution, and findings of a study conducted to better understand the use of documents in a range of modern office environments. It begins with an introduction of the study in Section 4.1, describing in more detail its purpose and the questions it sought to answer. Section 4.2 discusses the study design, while Section 4.3 gives an overview of the results and observations of the study. Section 4.4 provides discussion of these results, and defines a set of requirements for systems that aim to integrate paper and electronic document management. The chapter concludes with a summary in Section 4.5.

The study described in this chapter was conducted with the approval of the ethics committee of the Faculty of Computing and Mathematical Sciences, University of Waikato. A copy of the approval letter can be found in Appendix A.

4.1 Purpose of the Study

The design of an effective system for improving integration of paper and electronic document management needs to be guided by an understanding of the roles of paper and electronic documents, and the current systems used to manage them. Although, as discussed in Chapter 2, a number of previous studies have investigated paper document use and management in offices (e.g., Case, 1986, 1991; Cole, 1982; Frohlich and Perry, 1994; Malone, 1983; Whittaker and Hirschberg, 2001) at least a decade has elapsed since most of these studies were conducted. The intervening time has seen continued developments in electronic technology. The cost of computer systems has decreased, while processing ability and storage capacity have risen markedly. Electronic communication via intranets and the Internet has become ubiquitous, and is constantly increasing in throughput capacity. Interaction with computers has improved dramatically from the textual interfaces that would have been used in the days of studies such as Cole's (1982). Additionally, widespread availability and use of mobile devices, such as smart phones, netbooks, and laptops, mean that electronic documents are no longer limited to the desktop.

In light of the developments in electronic technology in recent years, it cannot be simply assumed that office work-flows and document usage have remained the same as they were in those earlier studies. It was therefore necessary to conduct a new study to investigate the use of documents, with special attention given to documents in paper form, in present offices.

Although there is a more recent study (O'Neill et al., 2006) of this type, which provides a description of paper document use and work-flow in a patent office, the scope of this study is limited to a single office.

The purpose of the study described in this chapter, therefore, was to investigate the roles documents play in present-day offices and to gain a better understanding of the motivations for using paper documents. The study also investigated the systems and methods for organising and managing documents in each of the offices studied, and the problems that were encountered managing documents. The main questions that this study sought to answer are:

1. What roles do paper and electronic documents serve, and why?
2. How are documents managed, both in paper and electronic form?
3. What problems are encountered when working with and managing paper and electronic documents?
4. How do people deal with the fragmentation between paper and electronic forms?

4.2 Study Design

In order to answer the above questions, an observational study of paper and electronic document usage in offices was designed and conducted. The design of the study is described in this section, which begins with a discussion of the study methodology, followed by an overview of the offices that participated in the study, and concludes with a discussion of the methods of data collection and analysis.

4.2.1 Methodology

A qualitative approach has been taken in many previous studies of document use in offices, such as those referenced in Section 4.1. Data collection methods used in these studies include interviews with participants, and observation of the participants' workspaces (e.g., the "tour of their office" conducted by Malone, 1983).These two methods of data collection (observation and interviews) are commonly used in qualitative research (Stainback and Stainback, 1988, p. 48).

This study follows in a similar vein, and comprises two components: a tour of the office given by the participants from that office, and an interview with the participants. During the office tour, the participants were asked to:

- describe and demonstrate the work that they normally perform in their office;
- show the types of documents they use to carry out this work;
- explain how these documents fit into their work-flow; and
- describe how these documents are managed in their office.

During this tour, participants would also be asked questions in order to clarify their comments, to concentrate on points of interest, and to keep from going too far off topic.

After the office tour an unstructured, open-ended interview (see Gorman and Clayton, 2005) was conducted. A number of pre-defined interview questions were used to guide the interviews. These questions, which are listed in Table 4.1, were not used verbatim. Some questions were omitted or modified depending on responses to previous questions during the office tour or the interview. The questions are split into three groups. These correspond to the four questions listed above (with the exception being questions 2 and 3, which are combined into a single group because they are closely related).

4.2.2 Overview of participating offices

The study was conducted at eight different organisations, which covered a diverse range of organisation and office types. In most cases, only one office was visited per organisation. However, in two cases, multiple separate tours and interviews were conducted at a single organisation. Even though the offices were from the same organisation, they are treated separately here because they worked independently. For the rest of this chapter, the term *office* is used to refer to the individual offices in the organisations, whereas *organisation* is used to refer to the organisation as a whole.

While the types of offices studied were diverse, they were mostly of a smaller size. One reason for this was that even in the larger organisations that were visited, the offices that participated in the study were generally smaller offices, operated by a few key administration staff. Another reason for this was the difficulty of finding organisations with larger offices dealing with both paper and electronic documents that were willing to participate in the study.

The organisations that took part in the study are described below. A total of eleven offices, and fourteen individuals participated in the study. Each office has been given a short ID that is used to refer to it in this chapter, these are listed along with a brief description of each office and the number of participants, in Table 4.2.

Organisation 1 — Event/exhibition

The first organisation to be visited was an event and exhibition organisation and consultation business. The office, when visited, comprised three staff members. The study was conducted with the administrative assistant who performed the majority of filing tasks, for both paper and electronic documents. The most common types of documents worked with included records of communication, contracts, and invoices.

Table 4.1: Questions used as a guide when conducting the unstructured, open-ended interviews.

Roles of paper and electronic documents
What does a typical task involve in terms of documents? Do you use paper documents when performing typical tasks? When do you use paper vs. digital documents when performing typical tasks, and why? Is your office moving towards using digital documents? Do you see paper documents being phased out in your office in the near future?
Document management

How do you manage your digital documents (e.g. Document Management System (DMS) software)?
How do you manage your paper documents?
 	
Document Storage
How do you store paper documents?
Do you have different levels of storage (e.g. long term, short term, etc.)?
    If so, how do these levels differ in terms of storing and retrieving documents?
How do you store a digital document?
How long is a given document stored for?
Who can access documents in storage?
 	
Document Retrieval
How would you locate and retrieve a given digital document?
How long does this typically take?
How would you locate and retrieve a given paper document?
How long does this typically take?
How do you group your documents (e.g. are they kept together in per client folders, etc.)?
 	
Access Control
Is there control over access to documents?
How is access to digital documents controlled?
How is access to paper documents controlled?
 	
Lost Documents
Are documents ever lost?
    If so, is it frequent or occasional?
Do you have a procedure in place for finding lost document?
Do you have a set document tracking procedure?
 	
Sharing Documents
Are there cases where multiple people need to share a paper/digital document?
How do multiple people share a physical/digital document?
    E.g., are multiple copies of the document created?
    How long are these multiple copies kept?
Are there cases where multiple copies of the same physical/digital document are kept long-term?

Fragmentation
How do you find the integration of the paper and digital document management?

Table 4.2: Summary of the organisations that participated in the study.

ID	Short description	No. participants
1	Event and exhibition organisers and consultants	1
2	Production department of a newspaper company	1
3	Company operating forestry training courses	3
4	High school	
4a	Principal's personal assistant	1
4b	Student records	1
5	Power generation company	
5a	Drawing management	1
5b	Library	1
5c	Generator project	1
6	Emergency operations centre	1
7	Medical centre	2
8	Legal office for a multinational company	1

Organisation 2 — Production department of a newspaper company

Organisation 2 was a newspaper company that publishes numerous newspapers and magazines, and the office that was visited was the production department of this company. The main function of the production department is creating advertisements for the publications from specifications provided by the sales representatives who sell the advertising space. The department had around 30 staff working in it, and the tour of the office and interview were conducted with the production manager.

The typical work-flow of the production department begins with sales representatives collecting specifications and material from the client for a given advertisement. This information is entered into an in-house advertisement tracking system. Supporting physical materials are either brought into the office and placed into job bags (A4 envelopes with a job information sheet attached to the front) in the case of local jobs or, in the case of jobs from remote offices, they are scanned at the remote office and added as electronic documents to a shared storage space on the company intranet. The jobs are picked up by operators who produce the advertisements electronically. The sales representative gives a proof of the advertisement to the client, either in paper or electronic form, depending on the client's preference. If alterations are requested these are transmitted back to the production department (generally by fax) and an iterative process of updating and proofing ensues.

Organisation 3 — Forestry training organisation

The third organisation participating in the study was a company operating forestry training courses. This organisation was of a small size, with 4 people working in the office. The office tour and interview were mostly conducted with the managing director of the company. The secretary and project manager for the company also providing input during the tour of the office.

Tasks performed in the office included: preparing training materials; organising, scheduling, and operating courses that are run off-site; assessing participants in the courses; and generating invoices. Documents used in performing these tasks includes training materials, reference material, records (for example of assessments), correspondence, and invoices.

Organisation 4 — High school

Organisation 4 was a New Zealand high school, with approximately 1000 students and around 100 staff members. At this organisation the study was conducted in two parts.

4A — Principal's Personal Assistant (PA). The first part of the study comprised an office tour and interview with the principal's PA. Document-based tasks that she performed included: managing personnel files, typing documents, and processing incoming mail. Documents used in doing this included correspondence, both in email and hard-copy form, records, and documents used as part of the authoring process.

4B — Student records. The second part of the study was conducted with the student records administrator and centred around the student records that the school kept in paper form. The records contained hard-copy documents, such as enrolment forms, birth certificates and other such legal documents.

Organisation 5 — Power generation

Organisation 5 was an electricity generation company that has a number of offices around New Zealand. The study was conducted at a single site and comprised three separate parts.

5A — Drawing management. The first part of the visit to this organisation covered the technical drawing management system, a library of over 50,000 technical drawings stored both electronically in a drawing database and in paper form. The tour of this office and interview were conducted with a staff member whose job included administration of the system. In addition to the drawings kept in the central drawing library that was visited, each power station also stores copies of the drawings that pertain to it.

5B — Library. The second part of the study was conducted at the on-site library, with the staff member who managed it. The library included books, reports, and other documents in both paper and electronic form. The library catalogue was maintained in an electronic database that was accessible on the company intranet through a web-based interface.

5C — Generator project. The final part of the study was conducted with the document systems manager for the generation department. This part of the study looked at the document system for a generator development project that was underway. The document management system for this project was mostly paperless, however while paper was generally not used in the office, there were still elements of paper use, for example at the construction site.

Organisation 6 — Emergency operations centre

This organisation was the emergency operations centre of a city council, used for communications and co-ordination of operations when responding to civil emergencies. The study was conducted with one of the four staff members who operate the centre normally. During an exercise or emergency a large number of people use the centre as a base from which to operate.

In normal work (i.e., outside of an emergency), a mix of paper and electronic documents are used. During an emergency paper plays an important role as it is used to carry and log communications (except for email). It is also used for reference material, such as operations manuals and the contacts database.

Organisation 7 — Medical centre

The seventh organisation participating in the study was a medical centre, with approximately twelve staff members, including administration staff, doctors and nurses. The tour and interviews were conducted with two members of the administration staff. Documents used in this office mainly comprised patient records in electronic form, as well as legacy patient records, forms, and correspondence in paper form.

Organisation 8 — Law office

The eighth organisation was the New Zealand branch of a large multi-national company in the fast moving consumer goods industry. The tour and interview was conducted in their legal office, which was relatively small, comprising two staff members. The study was conducted with one of these staff members, who is a legal assistant. This office dealt with a number of documents, organised into folders, and stored in either paper or electronic form, or both.

4.2.3 Data collection and analysis

While conducting the study, the main method of data collection was video recording. Although some hand-written notes were made, most of the important information was spoken or observed, and thus recorded as video. Table 4.3 gives the length of video recorded at each organisation. As video recording was stopped during interruptions, these times account only for active tour or interview time, and not for other interactions such as greetings. The length of the video varied depending on the time taken for the tour of the office and the interviews.

Analysis of the recorded data began with processing of the video. This involved watching the video, transcribing the conversations, and recording any noteworthy observations. At this stage, some initial filtering was performed, removing or summarising extraneous parts of the conversation that had little or no importance to the study, rather than transcribing them literally. The length of the transcription for each interview is also given in Table 4.3.

The analysis of the study data followed an iterative process, that passed several times through the five phases of qualitative data analysis described by Dey (1993): reading and annotation, categorising, linking and connecting, corroborating, and producing an account. Four high-level categories corresponding

Table 4.3: Summary of video recording for each study.

Organisation	Video Length (mm:ss)	Transcription Word Count
1	37:26	7137
2	37:33	5148
3	83:05	2065
4a	37:54	4142
4b	8:29	598
5a	60:55	2859
5b	18:35	2010
5c	33:33	3019
6	36:58	2424
7	84:41	1493
8	34:13	3723

to the questions previously posed in Section 4.1 were used:

1. the roles paper and electronic documents serve;
2. how these paper and electronic documents are managed;
3. problems associated with their management; and
4. problems due to fragmentation and ways that these were dealt with.

Within these categories, sub-categories were identified from the data, except in the case of Category 2, in which case the framework described in Section 2.2.5 provided the sub-categories.

4.3 Findings

The results of the processing and analysis from the study data are presented in this section. These findings are structured to correspond to the categories described in Section 4.2.3, beginning with roles of documents (Section 4.3.1), followed by a discussion of management of documents (Section 4.3.2), problems with managing documents (Section 4.3.3) and fragmentation of documents (Section 4.3.4).

4.3.1 Roles of documents

This section discusses the major roles of paper and electronic documents that were observed in the offices that were studied. The roles that were identified were: storage of records, communication, authoring and proofing, reference material, and data collection. Each of these is discussed in more detail below.

Records

One of the most prevalent roles served by paper documents was the keeping of records. The term records refers to content that is kept for legal or administrative purposes (see Section 2.2.1 for further discussion of this term). Common types of records observed included legal documents (e.g., filled in and signed forms, signed contracts, birth certificates) and records of correspondence (e.g., copies of letters sent and received, both by email and post).

All of the offices studied had records of some form. There was some variety as to what was kept in the records, and how much information was kept in paper form. At one extreme there was Office 1, where a paper copy of nearly every document that was produced or received was filed, while at the other extreme Office 5C kept most records electronically, with only signed contracts stored in paper form. The other offices fell in between these two extremes.

One of the major reasons for using paper rather than electronic documents that was observed during the study was that it was required either by law (or another external authority), or by internal policy. An example of records being kept due to the requirements of an external authority was observed at the forestry training organisation (Office 3). This organisation was required by the New Zealand Qualifications Authority (NZQA) to keep records of assessments for review and moderation.

Records of assessments were also kept as part of internal policy so they could be referred to in cases of disputes or to identify where data entry errors had been made. Office 1 provided an example of records being kept in paper form due to internal policy. The company director required paper copies to be kept for most documents because having such copies was perceived to increase accountability. Similarly, in the newspaper production department (Office 2), where paper records were kept so that they could be referred to in case of queries, since they provided proof of the information and instructions received from clients, forming a paper trail that could be followed in case of any problems. Likewise, records of communications were also kept in paper form by Office 4A in the event of a dispute. In this case, although letters were drafted electronically, and may have been revised multiple times, a paper copy of the final version was kept as proof of what was actually sent, since it included the school letterhead and principal's signature, which would act as proof that this was indeed the final copy of the document. The student records in Office 4B are another example of records as they were required to be kept in paper form since they were legal records such as signed forms and copies of birth certificates.

Another reason for keeping records in paper form that was observed was that some documents originated in paper form and it made more sense to keep them in paper form rather than digitising them. An example of this could be seen in the newspaper production department (Office 2). While material coming into the production department from off-site offices was scanned and transmitted electronically, material received locally was kept in paper form. Likewise, the law office (Office 8) had a large number of legacy documents that were in paper form. The size of the paper-based archives meant that the cost of digitising these would be large, and the benefits gained would be unlikely to outweigh the cost.

Similarly, the medical centre (Office 7) maintained an archive of patient records in paper form since it would be too time consuming to digitise these. Most of the records in this office (e.g., patient notes, medication, transactions, accidents, and immunisations) were kept electronically using electronic document management software designed specifically for medical records. However, records that were received from outside typically came in paper form, and had too many pages of varied sizes for the benefits of digitising them to outweigh the cost in terms of the time it would take to digitise them. In this case storing them in paper form was a better trade-off than scanning them.

Communication

Another role commonly served by documents in the offices studied was that of communication, both internally within an office or organisation, and externally with other entities. This confirms the findings of others, such as Frohlich and Perry (1994) who found in their study that interpersonal communication was the primary reason for the creation of paper documents.

The emergency operations centre (Office 6) relied heavily on paper for internal communications during an emergency. All incoming messages (except email) were recorded on a specialised paper form with carbon copy triplication. The original copy, coloured pink, acted as a physical token of that message, and whoever had that copy was responsible for acting on the message or assigning it to someone else. The second copy was given to the planning and intelligence team so they were kept up-to-date, while the third copy (known as the 'file copy') would be married up to the pink copy when the message had been dealt with. The advantage of paper in this setting was that it was easy to quickly make notes on and would always work, regardless of whether or not power was available.

Communication, including correspondence and face-to-face, with individuals and entities outside the organisation often involved paper. For example, some clients of the production office (Office 2) would prefer to view proofs of their advertisements in paper form while others were happy to receive them through email. Similarly, although the office of the generator project (Office 5C) was paperless, paper was often used at the actual construction site due to its portability.

Even in the medical centre (Office 7), which was mostly paperless internally, communication with external entities was typically conducted using paper documents. For example, when transferring patients, all notes kept electronically on the system would be printed and given to the patient or their new healthcare provider in paper form. While this could in theory be done electronically they choose not to, due to it being more difficult with their current software, and because the patient records may include notes in paper form that would also need to be transferred.

Electronic means of communication, such as email, were also commonly used for communication in most of the offices studied. On multiple occasions participants remarked on the increase in email-based communication, and corresponding decline in the use of telephone and fax for communication.

In Office 2, jobs collected at remote offices traditionally had been faxed into the main office. However, as the study participant from this office pointed out "*there's all sorts of problems with the fax, it's just a bottleneck, it's archaic*"[1]. For this reason, a move had been made away from faxing such documents to

[1]Quotations in this chapter are from the study participants.

scanning documents them at the remote office and storing them electronically on a shared network filesystem.

Authoring and proofing

Both paper and electronic documents were commonly used as part of the document creation process. It was observed on multiple occasions that the creation of a document, regardless of whether the final goal was an electronic or paper document, would involve using a combination of paper and electronic means. For example, in the cases of Offices 1 and 4A, the process would begin with the manager/principal giving the PA a draft of a document (either handwritten or on dictaphone), and the PA would type the document and return a printed copy of it to be annotated; thus, an iterative process of annotation on paper, entry of changes electronically, and printing would continue until the document was completed. The final document would either remain electronic (e.g., an email) or be printed, signed, and mailed. This confirms the findings of Frohlich and Perry (1994) who describe a similar iterative process, using both paper and electronic documents, for document authoring.

In the newspaper production department (Office 2) paper was used in the later stages of the authoring process. While the design work was done electronically, paper was used for proofing. Paper copies of the electronically created documents would be printed, since this enabled them to be proofed on the medium on which they would eventually appear (i.e., paper). This also had a psychological benefit, as shown by the participant from this office stating that the proofs acted as a "*security blanket*". It was also noted in this office, that as they moved to an electronic system for tracking jobs and job information, that it lacked some of the benefits of physical, tangible documents. The physical arrangement of job bags awaiting attention on the shelves provided an indication of how much work was left to do, however the electronic job tracking system "*doesn't scare you as much*".

Reference material

Another observed use of documents was as reference material. Such documents included drawings, books, and reports—mostly in paper form—that were kept for reference. Whittaker and Hirschberg (2001) similarly found that participants in their study kept paper documents to serve as long-term memory store and reference.

The drawing database (Office 5A) and the reference library (Office 5B) are both examples of libraries specifically designed to facilitate the storage and management of reference material. In these cases the whole purpose of the given office was to serve reference material to the organisation at large.

Other offices kept reference material at a smaller scale. For example the emergency operations centre (Office 6) had a contacts database and operations manual that were kept in paper form. The paper copy of this database was beneficial as it could still be used under adverse conditions where electronic systems may not be operational.

Data collection

Paper documents were commonly used when information needed to be collected. Examples of paper forms used included the collection of job information from clients of the production department (Office 2), assessment forms used by the forestry training provider (Office 3) to gather feedback from their participants, and forms filled in by first-time patients of the medical centre (Office 7).

In some cases these forms were ephemeral, serving to gather data that was entered electronically before being destroyed (in Office 7), or being stored as records (in Office 3). In other cases the filled-in forms had an ongoing use, such as the ‘job sheets’ that were used to gather job information in Office 2. These job sheets would be attached to the front of the job bags to provide information about the job.

4.3.2 Management of documents

This section discusses the management of paper and electronic documents in the offices visited. It is divided into five sub-sections based on the framework described in Section 2.2.5: organisation, maintenance, control, storage, and retrieval.

Organisation

A variety of methods of organising documents were employed in the offices studied. The organisation of documents for most of the offices can be represented by the structure shown in Figure 4.1. An *individual document* was rarely managed as an individual unit, but rather was grouped with other related documents into units commonly known as files (or job bags in the case of Office 2). However, the term file is not used in this discussion to prevent confusion with electronic files, which are usually individual documents. Therefore, these will be referred to as *organisational units*.

In some cases, such as when the number of documents outgrows the storage capacity of a *physical organisational unit*, several of these will be grouped together, forming a single *logical organisational unit*. For example, as shown in Figure 4.1, a logical unit could be all the documents relating to a given client, but all such documents may not fit in a single physical folder, and would therefore be split across several folders (i.e., the physical organisational units). In cases where such division is not necessary a physical unit will also be a logical unit. In some cases, such as the reference library (Office 5B), an artefact (e.g., a book) could be considered to be an individual document, as well as an organisational unit, both physically and logically.

The next layer of the hierarchy is the *category/grouping* layer. This layer accounts for any number of levels of grouping of the logical units (e.g., the client files in the Office 8 example were further grouped by department).

The top layer of the hierarchy is the storage level. The concept of levels of

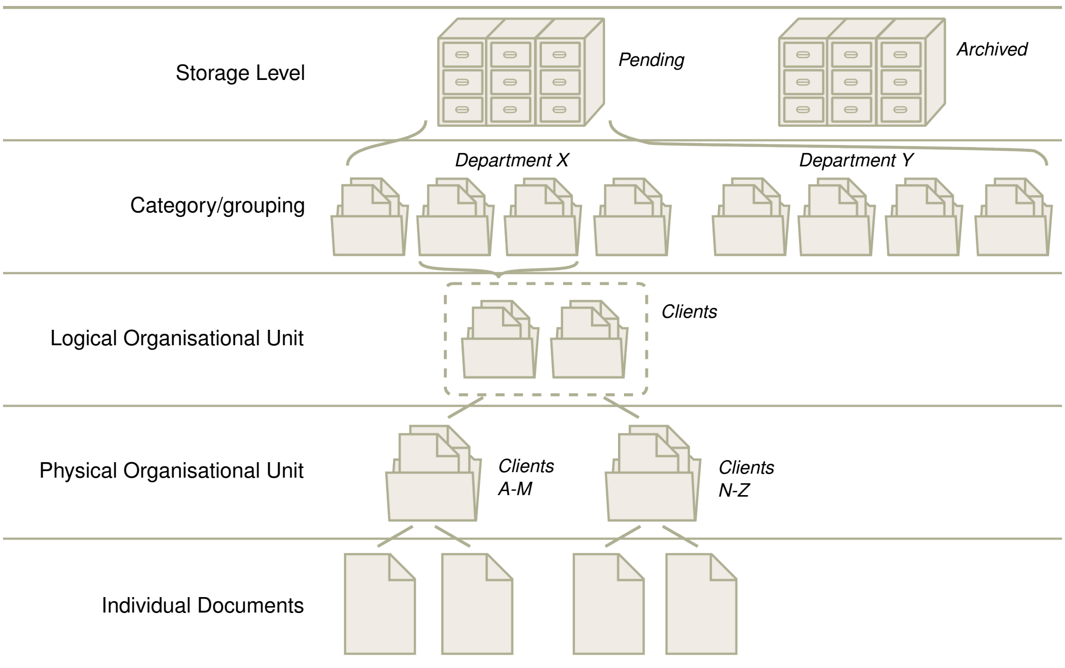

Figure 4.1: Document organisation hierarchy.

information was previously discussed in Section 2.3.1, where three levels of information that people commonly work with are identified. The observations of previous studies were confirmed in this study, since multiple levels of information were used in most of the offices that were visited.

For example, in Office 1, documents for the current year were kept on specific shelves. At a given point in time they would be moved to the lower shelves, while the documents from the lower shelves would be moved into an archival box. The box from the previous year would then be moved into cold storage. Aside from Offices 5A and 5B, all of the offices worked with multiple levels of information (at minimum, working and archived).

In most of the offices, information in electronic form was fragmented across a number of tools. In numerous cases (Offices 1, 2, 3, 4A, 6, and 8) these tools included a database, document files, email, and possibly other programs such as accounting software. In the cases of Offices 4B, 5A, and 5B, the offices were responsible for specific tasks (Office 4B student records, Office 5A drawing management, Office 5B library management), so did not need to work with information in the same range of forms as may be the case in other offices.

In Office 5C organisation of documents was in DMS software, and was flat. Instead of using folder hierarchies to organise documents, metadata was used since it was believed that this would prevent the problem of people filing material in different locations because they think differently. Instead, depending on the type of document, different fields of metadata would need to be mandatory.

Maintenance

One of the maintenance tasks was keeping an up-to-date index of the paper documents. In Office 4A two separate indexes were maintained. One for the archives, which was maintained by the school's archivist, and one that indexed the organisational units used to manage the rest of the documents. Both these indexes were kept in paper form. An electronic copy of the index was also kept which could be edited when updates needed to be made. However, when the indexes were referred to, a paper copy was used. A similar system was used to refer to the folders of drawings kept in Office 5A.

Also, in some of the offices, some form of electronic metadata would assist with the management of paper documents. An example of this was the online advertisement-tracking system used in Office 2, which contained metadata about each job. Similarly, in Office 5A information about each drawing was kept in a drawing database. Likewise, Office 5B made use of a database in order for people to discover documents in the library, and a spreadsheet was used in Office 8 to store metadata about each organisational unit.

Another maintenance task was moving documents from one level of storage to another. For example, in Office 1 there were multiple levels of information storage (as discussed in the previous section), and each year documents would have to be moved down a level.

Different levels of storage for electronic documents were also observed. For example, in both Offices 6 and 8 separate folder hierarchies were maintained for active and archived documents. In Office 4A documents were transferred to optical media and deleted from the system at the end of each year. Also related to moving documents between storage levels is the task of purging documents that are no longer required (discussed in the Storage section below).

Control

For all of the offices visited, access control for paper documents was enforced by controlling physical access to the office where the documents were stored. In most cases access to the physical documents was open to any person who was in the office. In Offices 3 and 4A, access to personnel documents had tighter control as they held sensitive information that was not to be viewed by all employees. Access to these documents was still limited by locking them in an office or cupboard.

Similarly, in Office 8 access to documents was controlled by limiting access to the room in which they were stored, which was known as "*the vault*". This control prevented files from being lost by people taking them away: "*no one's allowed in, else I'd lose everything I think . . . they just come through me and then I control it*".

In such offices where people were able to take documents away, different methods of tracking were applied depending on the office. In Office 4B, when someone needed to take documents from the student records archive the records manager would manually make a note of what was taken and who it was taken by. Email also played a role in keeping a record of who had documents: "*these days you usually get an email so you've got the trail there*". Likewise, in Office 5B, when people borrowed books from the library they would fill in a card from the back of the book with their name and the date they borrowed the book, and would leave the card with the librarian.

Access control to electronic documents was, in most cases, more advanced than with paper documents. In most cases users were required to login to computer systems and would access documents on shared network filesystems. Access to these shared filesystems could then be limited to appropriate users, (e.g., in Office 6, where only people within the office were able to access the shared filesystem). In the case of the drawing management system in Office 5A, although read-only access was available to anyone with access to the system, checking in updates would have to go through the administrator.

Storage

The storage duration and policies surrounding document destruction varied between the offices. For some offices the storage duration was a set time, such as "*a couple of months*" for proofs in Office 2, 18 months for assessments in Office 3 and seven years for student records in Office 4B. In several of the offices there was no set time, and documents were kept indefinitely. In the case of Office 6, "*at some stage we'll have to work out how long we need to keep them for*".

In all of the offices except Office 1, electronic documents were kept on a shared network filesystem, or in some form of electronic records management system. In some cases the electronic documents would be shared in this way, for example the electronic documents relating to a job in Office 2 would be stored on a network share, grouped by job. Similarly, in Office 4A documents were shared between two people in this way as well. Alternatively, as in the case of Office 6, a central network share was used but individuals kept their own personal filing systems on it.

Retrieval

The frequency of retrieval of paper documents varied depending on the office and types of documents. For example, in Office 2, due to the number of advertisements being put together (1200-1500 advertisements a week), there was therefore a small number of queries regularly coming through. These queries were mainly for more recent advertisements, and therefore recently filed information would be referred back to more often than older information.

Access to stored information in Office 4A was infrequent: "*once it's in there we generally don't refer back to it*". However, for some documents, although they may be retrieved infrequently, when they were needed they had to be found: "*tends to be that it is filed just in case it is required (which isn't often), but when it comes time to find something, it has to be found*" (Office 3).

In some of the offices (e.g., Offices 1 and 3), if the location of an electronic document was not immediately known, then a filesystem search would be used to find it. However, especially when the documents were stored on a network share, this could be time consuming.

As discussed above, in some other cases there would be some form of metadata stored that could aid retrieval of documents. For example, the job tracking system of Office 2, in which document files for each job were stored on a shared filesystem in a directory corresponding to the ID of the job in the job tracking system.

4.3.3 Problems encountered when managing documents

Numerous problems with paper documents were observed or described by the participants during the study. The major problems that were identified are discussed below.

Labour intensive and time consuming

One of the main problems with the management of paper documents that was observed during the study was, unsurprisingly, the time and effort that it required. Filing documents can be time consuming, while being a low priority task, and it can therefore be difficult to find time to do it within a busy schedule. This was evidenced by one of the participants from Office 3, who said "*I put it into here, and when [the secretary's] got time she files it . . . I don't have time to do that. So [the secretary] does it for me*". Similarly, the participant from Office 4A said of her pile of filing: "*this is my pile of filing here . . . and then when I've got a spare 3 days then I—unfortunately filing is low priority*".

Time was not only consumed by filing documents, but also by other manual tasks such as opening and closing files, archiving, and purging. For example, quoting the participant from Office 8, "*to open and close a file is just messy, you've got to stick all these stickers on, fill in all the forms, the Excel* *spreadsheet, all that*". This problem was partly due to the fragmentation of information, and the repetition of tasks that this caused: "*to open a file and to close a file you have to double up a lot. Filling out the matter form, and then you fill out that Excel spreadsheet and all that, then print the label and that's all the same information. I find it quite frustrating*".

Similarly, the participant from Office 2 remarked: "*we try to do as much as possible [electronically] because it's labour intensive [otherwise], you've got people carrying these things around, physically moving stuff around and following up; matching—particularly the fax in this instance—we're matching the fax sheets up together, . . . so it's why we want to get as much digital work-flow as possible, to speed things up basically*".

These findings are in accord with those of Gordon (1997) who notes that it is estimated that a quarter of the work time of information workers is spent distributing, filing, and retrieving documents.

Finding and retrieving documents and folders

Retrieving filed documents was also time consuming, and was not guaranteed to give a successful result. For example, the participant from Office 8 pointed out that it could take "*anything from around two minutes to maybe never*" to find a given folder of documents. Likewise, in the case of Office 1, although documents would be filed in paper form, often it would be quicker to print them again rather than retrieve the filed copy.

Similarly, talking about retrieving a specific paper document, the participant from Office 3 remarked: "*I couldn't find the one that they referred to. It's somewhere in my system, there's little bits of paperwork around the room, I know it's in a folder like that . . . but do you think I could find it when they were asking about it?*"

However, just because a document is stored electronically does not guarantee it will be easy to find. Retrieval is dependent on the amount of metadata available for a given document, and the accuracy with which such metadata is entered into the system. For example, if the metadata is limited to a filename, and that is entered incorrectly, this can cause problems when retrieving it, as described in the following quote from the participant from Office 8: "*it's getting more and more electronic and it's a lot easier because it's easier to search and find stuff—if you've called it the right name. But that's where* *document management systems are really good, whereas here if you've called it the wrong name, or you've mistyped something in, it might be lost forever*". Even if the filename was entered correctly, retrieving that file still requires recall of that filename, or where the file was stored: "*sometimes we have to say 'can anyone remember where this one was put?'. . . . you can probably use the search field if you can remember what your document was called*" (Office 6).

These issues can be mitigated to a degree though full-text searching of documents, which is usually possible for common document formats. This method was used in Office 3, but could be time consuming, taking 5–10 minutes. However, developments in desktop search (e.g., Minack et al., 2010) aim to improve searching of documents on desktop systems, for example by maintaining an index of documents for faster searching.

Duplication of documents

It was observed, primarily in Office 1, that a single document would be photocopied several times so that it could be filed into multiple categories. This is one way of working around the fact that documents can generally be classified into multiple categories, as discussed previously in the classification difficulties section (Section 2.3.1). Problems with this approach are that it increases storage space requirements, costs time and money to make copies, and increases the difficulty of maintaining consistency between multiple copies.

Similarly, in the same office, it was described how documents would be copied so that they could be taken away from the file (while leaving the original in the file). However, it would later be forgotten whether a given document was the originally filed version or a copy, and therefore whether or not it needed to be refiled. This could lead to multiple copies of the same document being filed in the same place.

Loss of documents

Another problem with paper documents (and other physical artefacts, such as books), was that they could go missing. This could be within the office itself, or in its communication with other external entities.

For example, regarding the parallel of paper and electronic document management systems in Office 3: "*the two technically should marry up, but there's* *sometimes glitches where forms get lost*". Similarly, on the topic of items checked out of the reference library (Office 5B): "*unfortunately things do tend to go missing, people leave and don't give things back*".

In terms of external entities losing documents, this was a problem in the past for Office 3: "*at one stage . . . the main ITO*[2] *that we deal with were so inefficient that we used to copy every piece of information that came in here; you'd photocopy, you'd send it off to them and it would get lost. And you'd have projects that would be held up for weeks—payment would be held up for weeks—because 'we haven't got your paperwork'*".

The participant in Office 5A noted that drawings did not go missing in their main drawing library, but they may go missing from the libraries of drawings kept on-site at each power station: "*at the station, yes, they go missing. They might have a contractor come in and whiz the drawing out of the folder and use that and not put one back. We really discourage that. We have done audits, but it's a huge job, it can be done we can get a complete list of the database with all the latest issues*".

Storage of irrelevant documents

One of the reasons that participants found it difficult to locate documents within their office was that a large proportion of irrelevant material was kept, which obscured the relevant material. The participant from Office 1 noted: "*irrelevant or not, everything is kept. Which makes it difficult to find things, because . . . you might have 10% relevant stuff and 90% irrelevant stuff. So to find that 10%, you have to waste so much time going through 90% of things*".

One reason for keeping irrelevant material was the difficulty of deciding what was relevant and what was irrelevant, especially when there was potential disagreement about this. As the participant from Office 1 remarked: "*I could decide which is [irrelevant], but [her co-workers] would have different ideas of what is, and what isn't irrelevant*". She gave an example of correspondence information being kept from ten years ago, and pointed out that even Inland Revenue do not require material over seven years old to be kept.

Another reason for keeping irrelevant material was that it requires discipline to take the time and effort to purge irrelevant material. "*But I think it's just the discipline that we need to purge it when it is past its use-by date. You tend*

[2] Industry Training Organisation

to go in spurts and you do it, and then you don't do it for quite some time, and you think 'Jeez, this is starting to get unmanageable'" (Office 3).

Other studies have noted the effect of irrelevant materials on document retrieval, and the difficulty of deciding what is irrelevant material and purging it (e.g., Barreau, 1995; Sellen and Harper, 2003; Whittaker and Hirschberg, 2001).

User specific

Another potential problem that was observed was that both paper and electronic filing systems can be very user-specific and, while they may make sense to the person who creates and uses them, they may not make sense to someone else. This may be an issue if that person is away or leaves the company. As noted by the participant from Office 1: "*there's lots of different places you can file it on there, and what happens is that I know it, I know where they're all filed, but someone new comes into the system, if I'm not here or I leave, then someone else is gonna try and find that*".

This underscores the findings of others, such as Cole (1982) who noted that filing systems are often idiosyncratic and individualised.

Managing dynamic content

An area where electronic documents excel paper documents is with respect to dynamic or frequently changing content. This has been discussed previously by others (e.g., Sellen and Harper, 2003, p. 31), and was confirmed by observations during this study. For example, the following quote from the Office 6 participant identifies problems with maintaining an up-to-date paper version of a contacts database. "*Another paper-based system that's proving to be quite archaic is our contacts database. . . . What we really need to do is try to have an electronic system that would allow all of these computers to access a central database . . . cause the single database would be more up-to-date than the paper copies that everybody's working from. Otherwise you've got to print off a copy every time, or set a copy for all of the desks each time*".

4.3.4 Fragmentation of paper and electronic document management

In all the offices visited, both paper and electronic documents were used to some degree, and therefore the severity of information fragmentation between paper and electronic forms varied between offices. At one extreme, such fragmentation was almost non-existent within Office 5C, which was a mostly paperless office, with only contract records stored in paper form. At the other extreme, Office 8 had a mix of information in both paper and electronic forms that was used in daily work, and this caused some problems due to information fragmentation.

One of the issues with maintaining information in both paper and electronic form is keeping the two organisational systems consistent. For example, in Office 3: "*the two technically should marry up, but there's sometimes glitches where forms get lost*". Likewise, in Office 8: "*if everything's been done right they should match up pretty well, and hopefully what you've got in paper form you should have in the system. But sometimes that doesn't happen*".

Linking between paper documents (or organisational units of paper documents) and metadata (e.g., an entry in an index, either electronic or hard-copy) was provided by some form of unique identifier. This was, in most cases, something descriptive (e.g., the name of the organisational unit) or a code that had some meaning (e.g., the hierarchical drawing numbering scheme used in Office 5A). Office 6 gave an example of the inverse of this, where paper documents contained links back to their electronic originals: "*we're trying to, on the bottom of our documents, put the path so we can find them in the future*". However, this approach was met with some difficulty, since "*once you rearrange your G drive, if that's an old document it might not be valid anymore*".

Such fragmentation, both between multiple electronic software and between paper and electronic forms, led to inefficiency when looking for information: "*cause you probably have to look in about 4 places; so you look in the paper file, you do a search for any of the electronic copies, or you'd search in the email*" (Office 8).

Fragmentation of paper document collections was also observed in Office 8, where it was noted by the participant how, due to mergers, several different physical filing systems had been combined. This led to inconsistency in the ways in which these were managed, with some being sorted by name and others by sequential numbering ("*which is just awful*" since "*you have to know what the number is before you can find it*").

4.4 Discussion

Despite the prevalence of electronic technology and its use in all of the offices that were studied, paper documents still have a variety of roles in these offices. The offices that were studied represent a range of office types, such as small business, school, reference library, and offices within larger companies. This shows that paper still has a role in many modern offices.

The roles that paper documents served in the offices that were studied were varied. At one end of the spectrum were roles such as authoring and communication in which paper documents were ephemeral, serving purposes such as proofing or marking-up before being disposed of. At the other end of the spectrum were documents that had long-term value, such as those that served as records or reference material.

Therefore, such offices need to have organisational systems in place to manage both paper and electronic documents. The organisational systems currently used in the offices that were studied have been discussed in Section 4.3.2.

The problems that users of these organisational systems encountered were also investigated by this study. One of the major problems was that organisation systems were labour intensive and time consuming to use and maintain, especially paper document management systems. Reasons for this included the limitations of paper documents, and the fragmentation of document management across several electronic tools and between paper and electronic forms. Other problems included difficulties finding and retrieving documents (especially those in paper form), inefficient document management strategies due to use of paper (e.g., keeping multiple copies of a single document, either intentionally or accidentally), loss of documents, and the storage of irrelevant documents.

Based on these problems that were observed, a set of requirements for systems that integrate paper and electronic document management have been identified. These requirements, which are discussed below, aim to address these limitations.

4.4.1 Requirements for an integrated system

From the study the following requirements have been identified for a system that integrates management of paper and electronic documents. While they are referred to as requirements, they are more akin to guidelines, since not all offices will have the same set of requirements. They are grouped into six categories, the first five of which correspond to the document management topics introduced in Section 2.2.5 (i.e., organisation, maintenance, control, storage, and retrieval), while the sixth category covers the requirements that fall outside this range.

Organisation

Support filing a document into multiple categories. Integration of paper and electronic document management would enable a single copy of the document to be kept in paper form and classification into arbitrary categories to be performed electronically. This would remove the need for multiple copies of a document to be filed in order to cope with multiple classification.

Integrate fragmented paper document collections. An integrated system would assist with reducing not only the fragmentation between paper and electronic document management, but also fragmentation of paper document management systems, by allowing these fragmented systems to be combined into a single integrated system.

Link paper documents to electronic documents. An integrated system could provide a means of linking between a physical document and its related electronic documents (e.g., the electronic original of a printed paper document).

Maintenance

Automate maintenance of collection indexes. By integrating the paper and electronic systems, the maintenance of indexes (e.g., those of Offices 4A, 5A, and 8) could be automated using electronic metadata. This could include the generation of automated metadata, (e.g., the time when each document was added to collection, or the history of who has accessed each document).

Simplify filing and refiling of documents. An integrated paper and electronic system could assist in reducing filing time by, for example, simplifying the organisational system, removing the need to make copies of a document (as described above), aiding in selecting, locating, and returning the necessary document folder to file the document in.

Better support for purging irrelevant information. The amount of irrelevant material kept can potentially be reduced by integrating electronic rules to assist with purging. For example, by tracking properties of documents and folders such as the age, frequency of access, and type of documents it would be possible to present the user with a list of out-of-date documents that could be purged.

Reduce manual labour. An integrated system could potentially reduce the amount of labour required to maintain paper document management systems. For example, the task of starting a new document folder could be simplified by automating the generation and printing of labels using the electronic metadata that is entered for that folder.

Control

Track documents and document folders. In order to support improved integration between paper and electronic documents it is important to be able to track physical documents and folders. This has been the subject of previous research, as discussed in Section 3.2.2. Tracking documents and folders is an important part of providing other functionality such as assisting the user to locate them, and providing alerts as documents progress through the office work-flow.

Storage

Reduce storage space requirements By improving the organisation of the system, for instance by removing the need for filing a document into multiple categories, and reducing the amount of irrelevant material being stored by facilitating the purging process, an integrated system can as a result reduce storage space requirements.

Retrieval

Integrate locating of paper and electronic documents and document folders. A significant feature of an integrating paper and electronic document management systems is integrating browsing and searching for paper and electronic documents. This is closely linked to automating collection indexes (described above), since it is likely to be these collection indexes that act as metadata for such searching and browsing.

Reduce time locating documents and document folders. By tracking documents and/or folders, as described above, it would be possible to provide assistance to users in order to reduce time taken when locating documents and folders.

Reduce difficulty of finding documents and document folders in an unfamiliar collection. By improving the way documents are found, for example, using metadata as described above, it would be possible to make it easier for someone who is unfamiliar with a filing system to be able to use it.

Miscellaneous

Ability to retrofit existing collections. Many offices have legacy paper documents. Therefore, there is a need that a system that integrates paper and electronic document management supports retrofitting of these existing collections of documents.

Ability to generate and display alerts. Several of the offices had to work to deadlines, or had a work-flow where actions would be triggered by certain events. Such offices would benefit from the ability to display alerts to the user. Rule based alerts are not a new concept, and have been previously suggested in a similar context by O'Neill et al. (2006). This study has confirmed that such features would be useful in other office environments beyond the patent office scenario described by O'Neill et al.

4.5 Summary

This chapter described an observational study that was conducted in order to gain a better understanding of paper and electronic document management in present-day offices. The study had four main components: identifying the roles of paper and electronic documents, the ways they are managed, problems encountered when managing them, and current ways of dealing with the divide between paper and electronic document management in present-day offices.

The study found that paper documents are still commonly used in modern offices, and that people still encounter a number of problems in managing their documents (in part due to the fragmentation between paper and electronic documents). This, therefore, confirms the need for better integration between paper and electronic document management. As such, a set of requirements for systems that aim to integrate paper and electronic document management was identified. The next chapter describes how such a system could be designed to satisfy these requirements.

Chapter 5

Design of an Integrated System

In the previous chapters the problem of information fragmentation was introduced, and the purpose of this research—to provide better integration of paper and electronic document management—was outlined. Overcoming the divide between paper and electronic document management can be viewed as part of a solution to the larger problem of information fragmentation.

This chapter provides a more detailed look into methods of integrating fragmented information. In information management literature, the term unification is also used in this context, and in this thesis, unification is considered to have the same meaning as integration. This chapter begins with a discussion of approaches to integrating information in electronic form (Section 5.1). This is followed by an overview of how an integrated electronic system can be extended in order to integrate management of physical artefacts (Section 5.2). The chapter ends with a summary (Section 5.3).

5.1 Integrating Electronic Document Management

It was established in Chapter 2 that the fragmentation of information is a significant information management problem. Information is often fragmented not only across paper and electronic forms, but also across the numerous software applications used to manage information electronically. It is, therefore, important to investigate ways of integrating information management in order to overcome, or at least mitigate this problem. Karger (2007) describes three categories of approaches to integration of fragmented electronic information.

Visual unification. With this form of unification, information from separate applications is displayed side-by-side on-screen. Making use of the visual unification one could thus read an email address from an address book while typing it into an email client.

Standard data types. By using standard data types applications can interact with each others' information. For example, allowing copy-and-paste from one application to another, or allowing an application to read the files of another.

Metadata. When following this approach, the content of the objects being managed is disregarded, and external information about them (metadata) is used to manage them instead.

All three of these approaches are important research areas, tackling different parts of the information fragmentation problem. However, the metadata approach is the most relevant to this thesis because it is the most applicable to integrating management of documents in both paper and electronic form.

Systems that use metadata to integrate management of information include *Placeless Documents* (Dourish et al., 2000) and *Haystack* (Karger, 2007). These systems both allow the user to manage units of information (documents in the case of Placeless Documents, and information objects in the case of Haystack) regardless of their information content.

The Placeless Documents system is implemented as a layer that integrates the management of documents within numerous existing document repositories (e.g., document management systems, IMAP servers for mail, the World Wide Web, and filesystems). A "Placeless Document" comprises a reference to a document in an underlying document repository, and metadata in the form of static and active properties. These documents can be grouped into dynamic collections, which themselves are a type of document. The Placeless Documents system provides the means for higher level applications to interface with it.

Haystack applies a similar concept, but rather than managing documents as whole files, these are instead "shredded" into individual information objects (e.g., individual contacts, calendar entries, or email messages). References to information objects and the relationships between them are recorded using the Resource Description Framework (RDF). Haystack is able to support a wide range of information object types as it focuses on the metadata associated with these objects, rather than the objects themselves. The Haystack user interface reassembles shredded information into user customisable views (see, e.g., Bakshi and Karger, 2005). Users are able to navigate through the information, following links between information objects, and perform queries over object metadata.

These systems provide examples of how the metadata approach can been taken in order to overcome fragmentation of electronic information. As systems taking this approach function independent of the types of information being managed, they therefore have the potential to support management of information stored in paper form, as discussed in the next section.

5.2 Integrating Physical Artefact Management

The discussion in the previous section outlined three approaches that can be taken in order to integrate information that is in electronic form. The third approach, making use of metadata to manage information, could conceivably be extended to support management of information on paper, in addition to electronic information. This is made possible by the fact that the information objects being managed are treated as opaque. That is to say that the system does not need to be able to read or understand them, it simply needs to be able to hold a reference to them.

When discussing management of information in paper form it is important to consider not only paper documents, but also the infrastructure used to manage them (e.g., folders, archival boxes, etc.). Therefore, this section discusses management in terms of *physical artefacts*, which includes both the documents and the physical infrastructure used to manage them.

Figure 5.1 shows how a common metadata layer would allow integration of not only fragmented electronic information, but also physical artefacts as well. In this example, several physical artefacts and several electronic information objects are each linked to their own metadata.

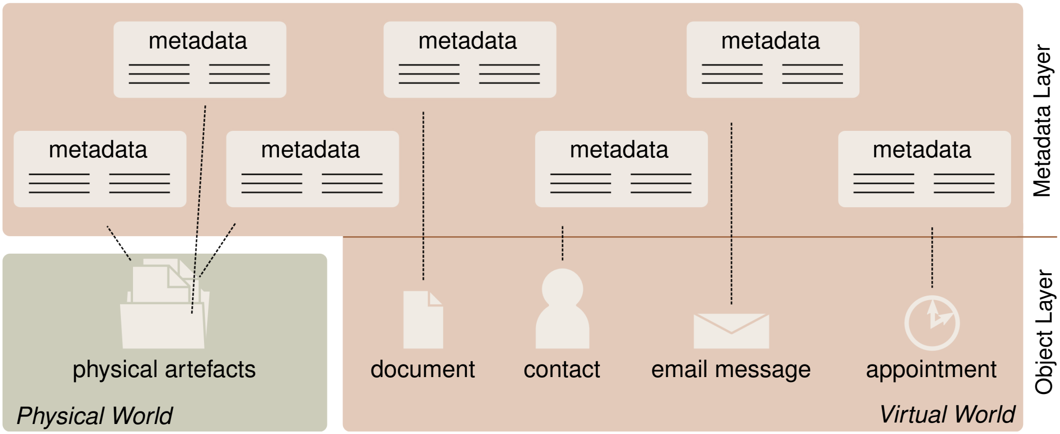

Figure 5.1: Integration of fragmented information objects using a common metadata layer. Each information object is associated with a metadata element, which can then be managed electronically.

5.2.1 Linking between physical artefacts and metadata

While the approach described above would enable integration of paper document management with electronic document management, it raises an additional problem that is not experienced with purely electronic systems: the crossing of the divide between the physical and virtual worlds. This was not an issue for the systems described above as they only manage electronic information, and thus only have to link between electronic metadata and electronic information objects. Incorporating management of physical artefacts with such systems, however, requires the ability to link between electronic metadata and physical world artefacts.

There are two main aspects to providing this link between physical artefacts and electronic metadata. One aspect is making the electronic system aware of the physical artefacts and their physical properties (e.g., their location). The other aspect is the presentation of electronic information back into the physical world.

A range of approaches can be taken in order to make the electronic system aware of physical artefacts, as discussed in Chapter 3. This can be as simple as giving each artefact a unique ID and using this as a key to metadata for that artefact. A more useful method is to automatically generate metadata about the object that reflects its physical state. For example, this might include information about its location that is generated by some form of tracking system. As such, information about artefacts in the physical world can be represented in the electronic world as metadata, and can be used as necessary by the electronic document management system.

The other aspect of providing seamless interaction is providing a way of presenting electronic information back to the user. The most common way this is done with present systems is through GUIs on computers. However, this does not provide very seamless interaction between the physical and electronic document management systems. Other alternatives include using overlaying electronic content on the physical world (e.g., using a projector, as in Raskar et al., 2004), or integrating interactive elements with the physical elements of the system (e.g., incorporating LEDs with library shelves, as in Choi et al., 2006).

5.2.2 Incorporating the physical document management system as a tangible/embodied user interface

The approach to integrating the physical and electronic elements of the document management system that is taken in this thesis is for the physical elements of the document management system to be incorporated as a tangible/embodied user interface. With this approach, the physical document management artefacts (e.g., folders) and infrastructure to support these (e.g., filing cabinets) are augmented with electronic circuitry that allows them to be linked to the electronic document management system. The circuitry enables the electronic system to be aware of such artefacts and infrastructure, to get information about their physical state (e.g., their location), and to enable them to act as tangible interfaces to the electronic document management system.

A initial set of requirements for systems taking this approach has been identified in order to guide prototype development. These requirements are listed below.

Unique identification of artefacts. In order to differentiate between individual artefacts, a means of uniquely identifying each artefact is required.

Artefact tracking To be able to assist the user in locating an artefact in the physical world, the system would need a way of determining an artefact's physical world location.

Artefact-level output. Each artefact would need some form of output, at its simplest this could be an LED to provide visual display of its location. This supports interaction from the virtual to physical world (i.e., output).

Artefact-level input. For the system to support user interaction it also requires an input capability. Therefore, each artefact would need some means of determining when a user is interacting with it.

Power supply for artefacts. In order for the artefact circuitry to function it requires power. Therefore, some means of supplying power to this is required.

Communication with artefacts. In order for artefacts to be integrated with the electronic document management systems, a means of communication between the artefacts and the electronic document management system is required.

5.3 Summary

This chapter provided an overview of three methods of integrating electronic document systems, and discussed how one of these methods, using a common metadata layer, could be extended to incorporate management of paper documents. This was followed by a discussion of an additional problem that must be addressed when incorporating paper document management with such an approach: the crossing of the divide between the physical and virtual worlds. Two main aspects of linking physical document management artefacts with an electronic document management system were discussed: making the system aware of the physical artefacts, and presenting electronic information back into the physical world. A discussion of the approach taken by the systems that have been developed as part of this thesis followed, and finally several initial requirements for such systems were identified. These requirements have guided the development of the prototype systems that are discussed in the next three chapters.

Part III

Development

Chapter 6

Prototype of an Integrated System

The previous chapter discussed how existing systems for integrating paper and electronic document management either lack support for interacting with the digital elements of the system beyond a traditional computer-based user interface, or have requirements that may be impractical (e.g., carrying a mobile projector). Therefore it was proposed that a system be developed that combines paper and electronic document management, and integrates its user interface with the physical elements of the document management system. The design and development of such a system is described in this chapter.

Also discussed in the previous chapter were the requirements to be met when designing such a system. These requirements can be summarised as follows:

- the ability to uniquely identify artefacts;
- the ability to track artefacts;
- artefact-level output;
- artefact-level input;
- power supply for artefacts; and
- the ability to communicate with artefacts.

The rest of this chapter describes the design and implementation of a proof-of-concept prototype system that was developed to meet these requirements.

At this point it should be noted that to avoid confusion between the terms *physical hardware* and *electronic hardware* the following terminology has been used in the remainder of this thesis:

electronic hardware or *circuitry* is used to refer to electronic circuitry, and

physical hardware is used to refer to the physical objects that are to be augmented with circuitry.

6.1 System Design

The design of the system centres on the augmentation of physical document management hardware with electronic circuitry that enables it to fulfil the requirements listed in the previous section. The selection of the hardware to augment is closely linked to the design of the power supply and communication circuitry, discussed below.

6.1.1 Selection of artefact type

As described above, the design of the prototype system is based around the augmentation of physical document management hardware with electronic circuitry. The purpose of the augmentations is primarily to make the document archive act as an embodied user interface that can be linked to electronic document management systems. Such an interface would provide both input and output through the physical artefacts as opposed to simply tracking documents, and as such it is important for this interface to be visible and accessible to the user. In Chapter 4 it was observed that documents are, in most cases, organised in some form of hierarchical structure comprising a range of different document management artefacts (e.g., folders, binders, and archival boxes). Therefore, if the documents themselves were to be augmented with a user interface, this would essentially be obscured by the artefacts that the documents are stored within. As such, it would be of greater benefit for the user interface to be part of the document management artefacts, rather than the documents themselves.

There are several other benefits gained by augmenting the document management artefacts rather than documents. For example, in the case of folders, the additional size and weight of the circuitry can be more easily absorbed by a folder because the relative size and weight of the circuitry is much lower in comparison to a folder full of documents than in comparison to the documents themselves. Similarly, as there are usually several documents per folder, there is a significant decrease in the number of augmentations required when folders rather than individual documents are augmented.[1]

For these reasons, the prototype system discussed in this chapter is based around the augmentation of folders, and the filing cabinet that they are stored within.

6.1.2 Power supply

In order for the electronic circuitry of the augmented folders to function it requires some form of power supply. The selection of the folder power supply was an important factor that influenced the overall design of the system.

There are several options for providing power supply to the folders:

- by a battery;
- through a wireless (e.g., inductive) connection; or
- through a wired, physical connection.

Each of these approaches has advantages and limitations. Batteries have the advantage of providing power to the folder even when it is out of contact with the rest of the system. However, they have the disadvantage of a limited capacity. This is potentially a significant issue, since archives may be kept for many years, and thus outlive the folder batteries. Therefore, these batteries would need regular charging or replacement, which would be impractical to carry out for individual folders in anything but a trivially small archive. In any case, the charging infrastructure would have to be provided through a wired or wireless connection. Additionally, batteries would add extra cost, weight, and size to individual folders that would likely prove impractical for large archives.

An alternative to including a battery on every folder is to supply power to the folders wirelessly. The most suitable form of wireless power supply for this type of application would be an inductive power supply (see Baarman and Schwannecke, 2009). The basic principal of an inductive power supply is that electricity moving through a coil (also known as an inductor) in the power supply unit generates a magnetic or electro-magnetic field that induces a current in a coil on the device being powered. This follows the same principle

[1]The augmentation of individual documents is discussed in Section 11.3.3.

as a transformer, but with greater separation between the coils. Passive RFID (AIM Inc., 2001) tags make use of inductive power.

Wireless power supplies have the advantage of not requiring direct physical contact in order to function. However, they have several limitations that would make them largely impractical for this application. Passive RFID tags are limited in the power that is available to them, as the tags are designed for identification rather than interaction. Wireless power supplies for higher power loads (e.g., wireless battery chargers) are designed for close proximity, and there is a noticeable drop-off in power as the device being powered moves or rotates away from the power supply. As such, there are no current solutions that are able to meet the requirements of the system being developed.

An alternative to wireless power supplies is to supply power through direct physical contact. Such a wired connection removes the need for batteries, and in most cases is able to supply more power than a wireless power supply. The main limitation of a wired connection is that it requires physical contact in order to work, and once this contact is broken power is no longer supplied. This limitation is not viewed as critical in the design of this system as it aims to support interaction with folders while they are in a storage location such as filing cabinet, rather than when they are being carried around by users. Therefore, for the prototype system described in this chapter a wired connection was chosen for power, and communication as described below.

6.1.3 Communication

Along with power for the folders, the ability to communicate with them is a requirement of the system. As with power supply, communication can take place across either a wired or wireless connection. There are many factors that affect the choice of connection and the communication protocol to be used, including communication speed, reliability, and security. In general, the choice of connection type is closely linked to the choice of power supply as, for example, it would be counter-intuitive to use a wired communication protocol with a wireless power supply. In most cases where power is supplied physically it is appropriate to use wired communication. However, there may be cases where physical constraints make a wireless communication protocol a more viable solution.

For the prototype discussed in this chapter, wired communication was chosen.

Using wired rather than wireless communication reduces the complexity of the design and is simpler to implement. Since this choice was made during design of the system, it was possible to design the physical connection such that wired communication was possible. The communication protocol is discussed further in Section 6.2.2.

6.1.4 Folder circuitry

In order for the system to be able to uniquely identify folders each folder needs to be augmented with an Integrated Circuit (IC) that includes some form of memory, in which is stored a unique ID. Ideally this memory would be one-time-programmable so that once a folder is given an ID this will not change. The ID need not be set by the user, and can be factory programmed when the IC is manufactured.

Another requirement of the system is for the folder to have some form of output mechanism. In this prototype system the output is provided visually, in the form of an LED attached to the top of the folder such that it remains visible when the folder is placed in the filing cabinet. This enables the folder to visually draw attention to itself, for example to display its location when the user wishes to identify and retrieve it. This visual display reduces the accuracy with which the location of the folder needs to be determined by the system. Instead the system could show the location of a given folder with low-grain accuracy (e.g., which drawer of which filing cabinet the folder is in), while the LED on the folder would display its location to the user.

Although the output device used in this implementation is visual, other types of output devices are also possible. These include a buzzer to provide an audible cue in case the folder is visually obscured, or a vibration motor to provide tactile feedback. However, in order to function, such devices may require more power than an LED and the choice of devices may depend on the power supply used.

In order to support user interaction in the input direction, each folder includes a button that allows the user to trigger events in the software system (e.g., displaying folder information on-screen).

6.1.5 Filing cabinet circuitry

The filing cabinet acts as the interface between the folders and the rest of the system. Physical contact between the cabinet and the folders supplies power to folders and allows for communication with them. The cabinet circuitry is able to connect to a computer so that it can be integrated with the electronic document management system.

In order to enable further interaction with the system, the cabinet also includes some input components, in the form of a dial and a button. These allow the user to interact with the document management system enables to, for example, scroll through information on a computer display by turning the dial.

Although the cabinet circuitry did not include any output devices in this implementation of the prototype, such devices could be included in future designs. For example, it may be desirable to draw attention to a given cabinet visually using an LED. Alternatively, the output could be in the form of an LCD screen that displays electronically stored information associated with the physical contents of the cabinet. However, this was assumed to be beyond the requirements of this initial prototype system.

6.2 Implementation

This section describes the implementation of the prototype system. It begins with a description of the physical hardware that forms the basis of the prototype (Section 6.2.1), followed by a description of the protocol that was selected for communication between the filing cabinet and folders (Section 6.2.2). Finally, the circuitry that was implemented for this system is described in Section 6.2.3.

6.2.1 Base hardware

The physical hardware that forms the basis of the prototype system described here is a vertical (or suspension) filing cabinet, and a corresponding set of folders. Since this system was to be a proof-of-concept rather than a full-scale implementation, a desktop file storage unit was used (shown in Figure 6.1) in order to increase its portability. This is referred to in the rest of this chapter as the *filing cabinet*.

Figure 6.1: The filing cabinet and folders that form the basis of the prototype system.

Folders are suspended in the filing cabinet by hooks on each side that sit on the the metal rails which run the length of the cabinet. Each folder has four hooks—two at the front, and two at the back (see Figure 6.2). Each pair of hooks is part of a plastic cross-member that runs the width of the folder.

The cabinet rails are metal and are insulated from each other. As a result, they can be used for communication between the cabinet and folders, providing two separate parts of electrical contact. The folder cross-members are plastic, and therefore they do not create a short circuit between the rails. However, since the folder hooks are non-conductive, they need to have some conductive material applied to them to contact the rails.

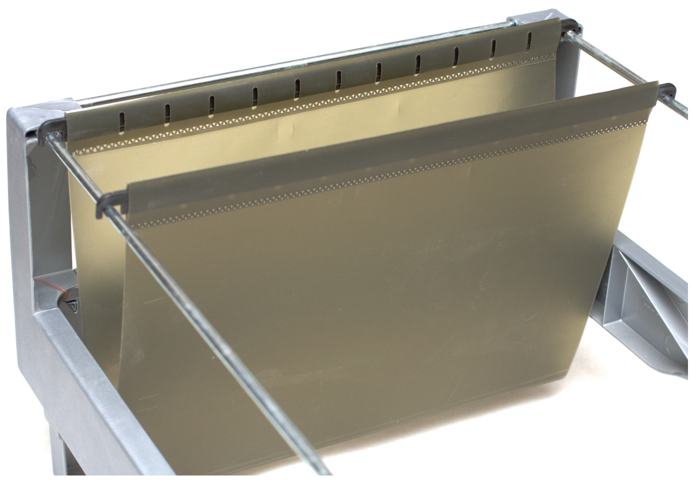

Figure 6.2: One of the folders that form the basis of the prototype system.

6.2.2 Communication protocol

A significant part of the implementation of this prototype system was providing a means of communication between the filing cabinet and the folders. As noted above, the two rails of the filing cabinet form the communication medium. This requires the use of a communication protocol that supports power and communication over a single wire (since the other wire needs to be the ground connection in order to complete the circuit). Because there may be multiple folders on the rails at any given time they will, therefore, form a bus network topology. Therefore, the communication protocol must be able to support a bus topology. Also, in order for the system to keep track of which folders are in the cabinets and detect the addition and removal of folders, the communication protocol also needs to support dynamic detection of addition and removal of devices on the bus, and needs to be able to read the unique IDs of these devices.

The most appropriate protocol that is able to work within these constraints is a protocol known as the *1-wire protocol* (see Linke, 2008). The naming of the 1-wire protocol can be misleading. The "one" in 1-wire actually refers to the fact that only one wire is required for both power and communication (whereas this typically would require at least two wires). However, two wires are in fact needed by this protocol, since a second connection is required in order to complete the circuit.

The 1-wire protocol operates with a master/slave topology. A single master and an arbitrary number of slaves may be connected to the network. Dallas-Maxim produce numerous 1-wire slave devices that perform a variety of functions, such as identification, digital and analog input and output, and temperature sensing. Each slave device has a unique 64bit ID hard-coded into it when it is manufactured. This allows it to be uniquely identified on the 1-wire network so that the master can direct commands to individual slave devices in order to read their status, as well as controlling them.

6.2.3 Circuitry

The circuitry used for this prototype is based on the design of the circuitry used in the VoodooIO project (Villar and Gellersen, 2007; Villar, 2007). VoodooIO is a *malleable control structure*, a tangible interface that can be dynamically altered to suit the needs of the user. It consists of two main parts, a flexible multi-layer substrate, and a number of controls such as buttons, dials, and sliders. The controls have pins underneath that puncture the substrate, serving the dual purpose of holding the control in place and creating an electrical connection between the control and the substrate (see Figure 6.3). The substrate has two conductive layers internally, while each pin of the controls is co-axial. The substrate is connected to the computer by a USB dongle. Communication between the USB dongle and the controls is conducted using the 1-wire protocol.

The technology of VoodooIO was ideally suited to form the basis of this prototype system, and was thus adapted to perform this role. Folders are equivalent to VoodooIO controls, with their hooks performing the function of the VoodooIO pins. Consequently, the rails and hooks method of creating a connection is equivalent to the substrate and pins of VoodooIO.

Folder circuitry

Each folder has a single DS2406 IC (Dallas Semiconductor/Maxim, 2009a), which has a 1Kbit user-programmable, write-once non-volatile memory and two General Purpose Input/Outputs (GPIOs). The LED and button of the folder connect to the GPIOs of the DS2406. The left and right hooks of the back cross-member of each folder are coated in a conductive material (Figure 6.4) that is wired to the appropriate pins of the DS2406 IC (data and ground). The wiring of the circuitry of the folder is shown in Figures 6.5 and 6.6. When the folder is constructed, a code is written to the non-volatile memory of the DS2406 that identifies it as being a folder.

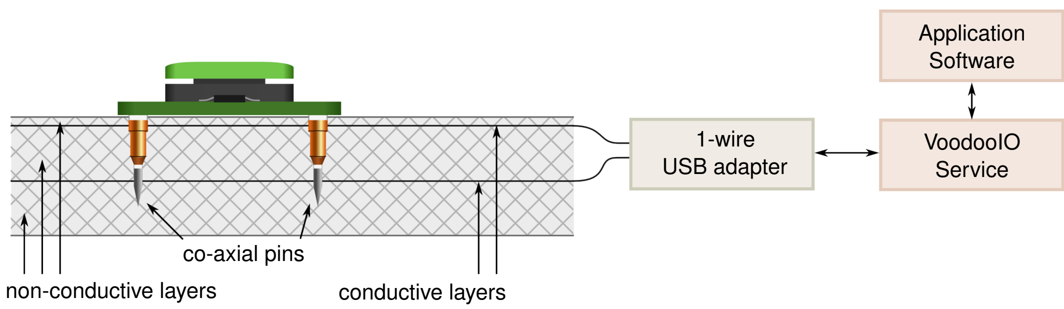

Figure 6.3: VoodooIO is a tangible interface which is made up of controls (e.g., buttons, dials, etc.) that can be pinned into a substrate. Co-axial pins on the controls contact conductive layers in the substrate, connecting the control to a 1-wire network.

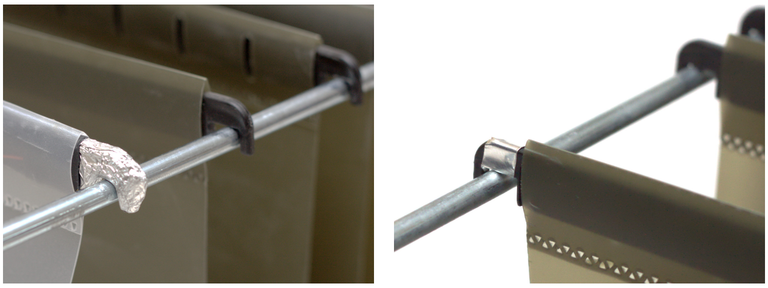

(a) Aluminium foil **(b)** Adhesive aluminium tape

Figure 6.4: As the hooks on the folders are made from plastic, these need to be coated in conductive material and connected to the 1-wire IC. Two conductive materials that were tested are shown here.

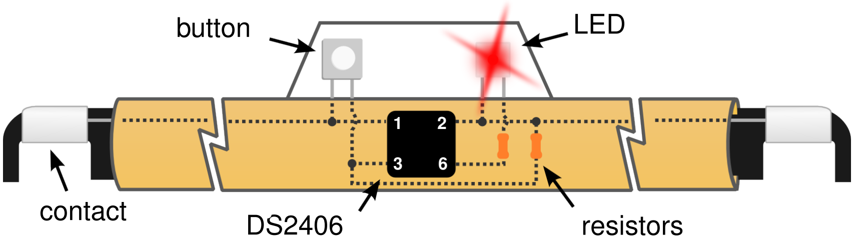

Figure 6.5: Folder circuitry. Dotted lines indicate wires running inside the folder. To increase the clarity of diagram the pins of the DS2406 are not shown in their physical positions.

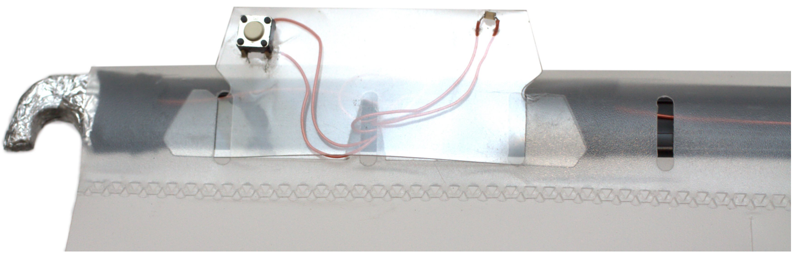

Figure 6.6: Close-up of the folder augmentations, showing the button and LED. The wires from these run to a DS2406 1-wire IC behind the plastic tab.

Cabinet circuitry

The cabinet circuitry also includes a DS2406 for identification. When the cabinet is constructed, the memory of this DS2406 is programmed with a code that identifies it as being part of the cabinet. Since one of the input devices of the cabinet is a dial (which is an analog component) the cabinet circuitry requires an additional Analog-to-Digital Converter (ADC) IC. As a result, the cabinet circuitry also includes a DS2450 (Dallas Semiconductor/Maxim, 2006b) 1-wire ADC. These components (the DS2406 and DS2450) connect to the same network bus as the rails of the cabinet, as shown in Figure 6.7.

The two wires of the 1-wire network terminate at an RJ-12 socket on the side of the cabinet. The socket is wired so that it can be directly connected to a standard 1-wire adaptor. The 1-wire adaptor used for this prototype was the Dallas-Maxim DS9490R USB adaptor (Dallas Semiconductor/Maxim, 2009b). Connecting the cabinet to the computer using a 1-wire adaptor allows the software to query the 1-wire network to determine the IDs of the 1-wire devices, and the cabinet and folders that are present. It is able to read the status of their input devices (buttons and, in the case of the cabinet, dial), and to control their output devices (folder LEDs).

The cabinet augmentations are shown in Figures 6.7 and 6.8.

6.3 Software

The previous sections have discussed the design and implementation of the physical and electronic hardware of the prototype system. In order for the

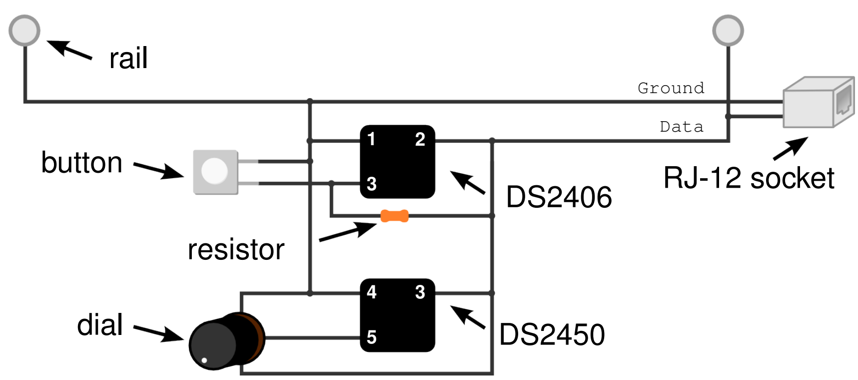

Figure 6.7: Cabinet circuitry. To increase the clarity of diagram the pins of the DS2406 and DS2450 are not shown in their physical positions.

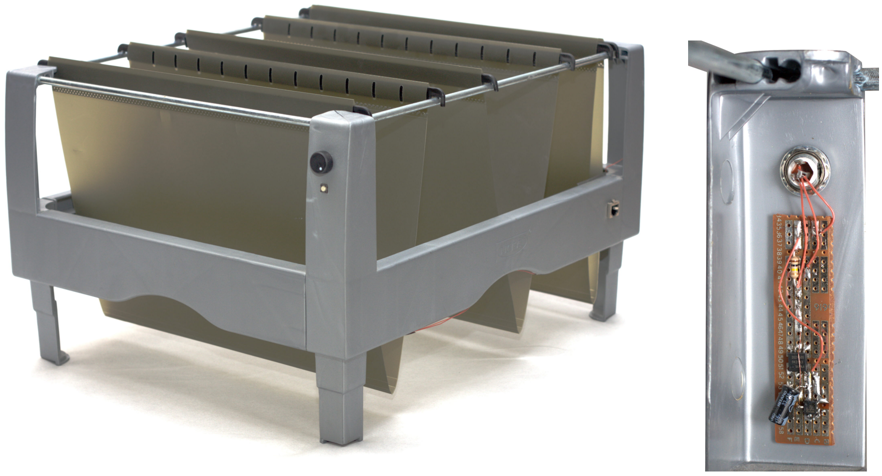

(a) Front (b) Back

Figure 6.8: The augmented cabinet showing the dial and button on the front, and RJ-12 jack on the side for connection to a 1-wire USB adaptor.

system to function it also needs a software component. The software is separated into two layers, the *middleware* (Section 6.3.1) and the *client application* (Section 6.3.2). The middleware layer is tightly coupled to the functioning of the electronic hardware and is independent of the client application implementation. On the other hand, implementation at the client application layer is dependent on the specific application of the system. The relationship between these software layers and the electronic hardware is shown in Figure 6.9.

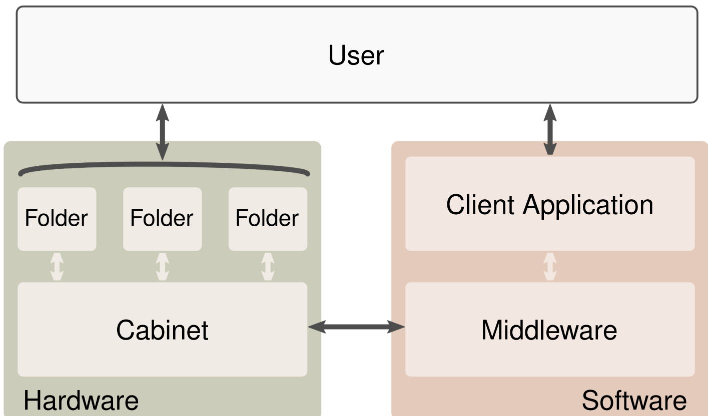

Figure 6.9: Block diagram showing the interaction between the software and hardware elements of the system, and the user.

6.3.1 Middleware

The middleware is the low-level software layer that is required for the electronic hardware to function. On its low-level side, the middleware interfaces with the hardware, polling the 1-wire network to detect the addition and removal of folders, and receive events from the folder and cabinet input elements. On its high-level side, the middleware communicates with the client applications.

Initially a slightly modified version of the VoodooIO service (described in Villar, 2007, chap. 6) performed the function of the middleware for the prototype described in this chapter. The VoodooIO service is the middleware element of the VoodooIO system that interfaces between the 1-wire network to which the VoodooIO controls connect and the software applications that receive input from these controls. At the low-level it is responsible for detecting the addition and removal of controls on the VoodooIO substrate, and the manipulation of these controls (e.g., buttons presses, or turning of dials). At the high-level it communicates this to application software using a simple plain-text communication protocol over TCP/IP. Slight modifications were made to the VoodooIO service in order to enable it to recognise folders and cabinets as types of VoodooIO controls.

However, since the VoodooIO service was not targeted towards the needs specific to this prototype, a replacement was eventually developed. The replacement middleware was able to be simplified, since much of the complexity of the VoodooIO service was not required. For example, the VoodooIO service is designed to prioritise manipulation events, whereas for the system described in this chapter such prioritisation was not required.

6.3.2 Client application

The client application layer is where the physical and electronic components of the system are integrated. While the middleware effectively acts as a bridge between the physical components of the system and the client application, the client application is where electronic content is actually associated with the physical folders.

The middleware provides an easy to use Application Programming Interface (API), and therefore the development of client applications is not tightly coupled with the development of the system itself. This is important as the requirements of a client application may vary significantly between different use cases, and therefore the design of a client application may need to be tailored to the specific case where it is to be used.

The next section describes an example of how the prototype system may be integrated with an electronic document management system.

6.4 Integration with Aroo

The prototype described in this chapter was adopted by the Media Interaction Lab at the Upper Austria University of Applied Sciences for use in the Aroo project[2]. Aroo is a sub-project of the Office of Tomorrow research project[3], whose other sub-projects include Coeno (Haller et al., 2005) and the Shared Design Space (Haller et al., 2006). Coeno is an environment that allows users to work collaboratively using a shared augmented reality tabletop interface. The Shared Design Space is a similar collaborative environment that includes augmented reality tabletops and whiteboard interfaces that users can interact with using Anoto digital pens (see Silberman, 2001).

Aroo (Seifried, 2007)—which is also known as the Smart Filing System (Seifried et al., 2008)—is the implementation of a system that combines paper and electronic document management. Using the prototype system described in this chapter as its basis, Aroo adds client application software and support for digital pen-based annotation and interaction on an augmented reality tabletop interface.

6.4.1 Client application

A client application was developed by the Office of Tomorrow group, based on Microsoft OneNote™(Figure 6.10). Each of the physical folders corresponds to a virtual folder in OneNote. This creates a link between related physical content (e.g., printed documents stored in the physical folder) and electronic content (e.g., electronic documents stored in the virtual folder).

Two-way interactions between the physical interface and the on-screen display is possible. In the output direction, selecting a virtual folder in OneNote causes the LED of the corresponding physical folder to be lit. Conversely, pushing the

[2] http://mi-lab.org/projects/aroo/
[3] http://mi-lab.org/projects/office-of-tomorrow

button on a physical folder causes its corresponding virtual folder to be selected in OneNote. Additionally, turning the dial on the cabinet scrolls through the virtual folders on screen while also illuminating the LED of the selected folder.

Using this client application it is possible to perform a search over all the electronic content, for example to locate a desired document, and then display the location of the folder that matches this search using its LED in order to retrieve a physical copy of the document.

6.4.2 Incorporating other technologies

The prototype system described here may be combined with other technologies in order to add extra functionality and to allow the folders to be integrated with other systems. This is demonstrated by Aroo, which adds extra augmentations to the folders so that they can be annotated and used on an augmented reality tabletop interface.

Annotation support was added to the folders using Anoto technology. An annotation area made from Anoto paper (see Silberman, 2001) was attached to the front of each folder and any annotations on this made with an Anoto pen could be captured and transmitted to the computer (Figure 6.11). The

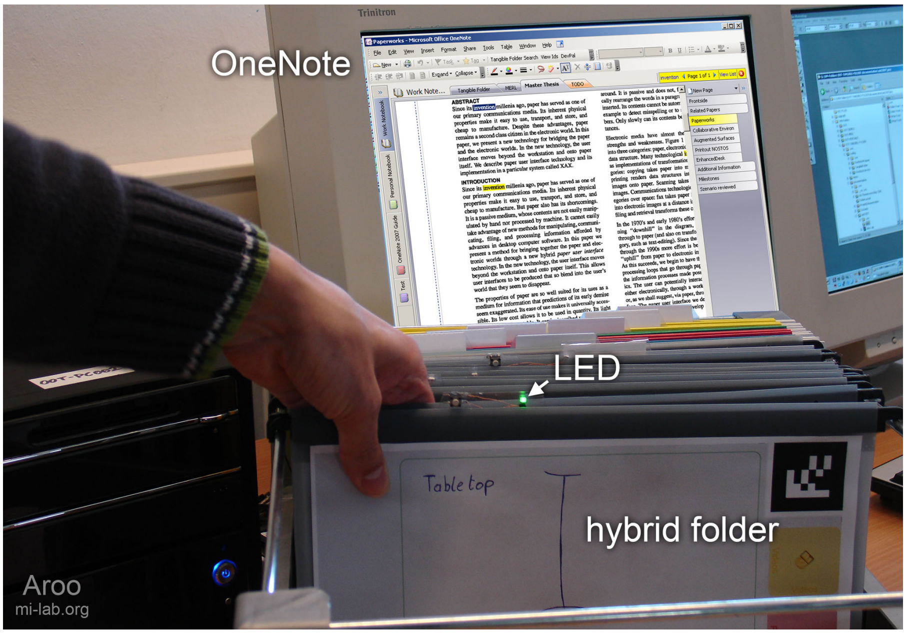

Figure 6.10: Aroo combines the prototype system described in this chapter with a Microsoft OneNote based client application (Image source: The Media Interaction Lab, 2007).

captured annotations are imported into the client application and associated with the virtual folder corresponding to the physical folder that is annotated.

Previous research has also investigated ways of using Anoto and other similar technologies to integrate paper documents with electronic systems. Examples of such research include linking from text on a paper document to electronic content (Norrie et al., 2006), and automatically applying annotations on printed paper documents to their electronic originals (Weibel et al., 2008). Such functionality is complementary to that provided by the system described in this chapter, and could be incorporated with this system if required.

Aroo also demonstrates the ability for folders to be integrated with augmented reality systems. ARTag markers (Fiala, 2005) were added to the folders so that they could be registered by an interactive tabletop interface (e.g., Haller et al., 2006). This allows electronic content to be projected onto them, as shown in Figure 6.12.

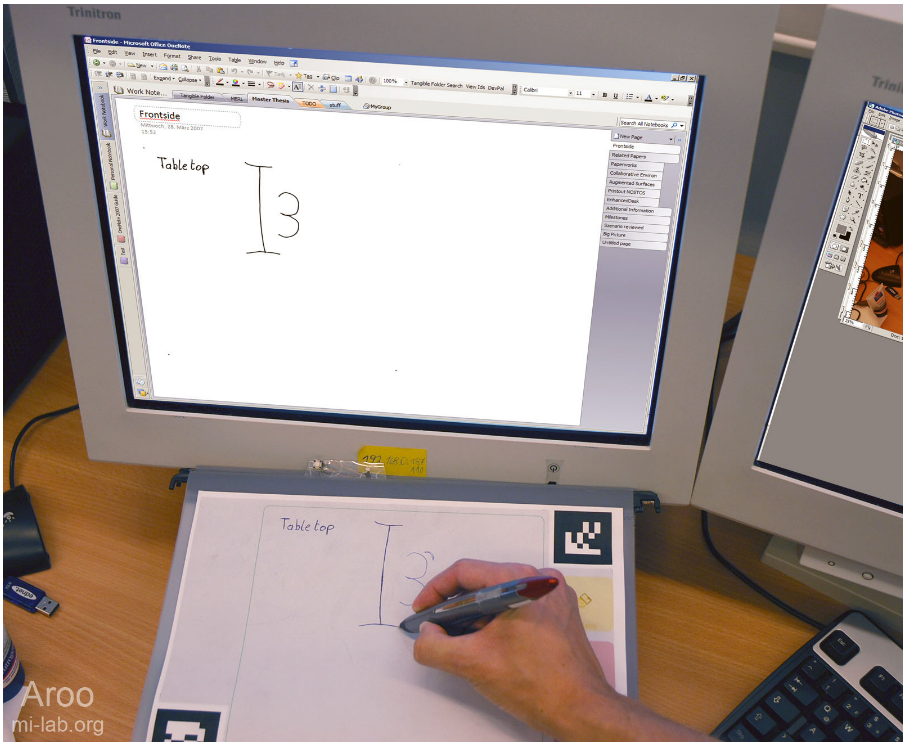

Figure 6.11: Other features, such as support for synchronising physical annotations with the application software using Anoto technology (shown here), can be easily incorporated with the prototype system described in this chapter. (Image source: The Media Interaction Lab, 2007).

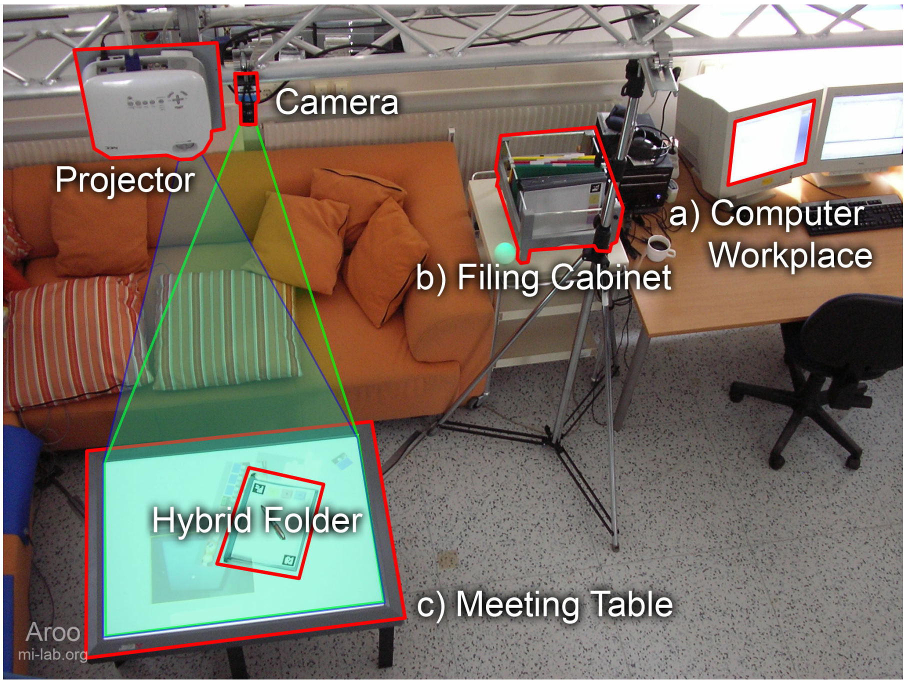

Figure 6.12: Aroo combines the prototype system described in this chapter with application software, and adds support for using the folders on an augmented tabletop interface. (Image source: The Media Interaction Lab, 2007).

6.4.3 Pilot study

An informal pilot study was conducted by the Office of Tomorrow group to evaluate the usability of the Aroo system (for details see Seifried et al., 2008). The study was conducted using a prototype of the Aroo system that comprised one cabinet and five folders. Participants were university colleagues of the Office of Tomorrow group who were not affiliated with the project.

The participants' response to the system was positive. They found it easy to locate folders using either search or synchronised navigation, and found the LED to be a useful aid as it immediately showed the location of the folders that they were looking for.

6.5 Discussion

The development of the prototype described in this chapter demonstrated the feasibility of developing an integrated paper and electronic document management system. It also allowed the evaluation of the underlying design concept, and its potential benefits and possible flaws.

The development of the client application by a third party, on the other hand, demonstrates that intimate knowledge of the functioning of the hardware is not required when a high-level API is available. As such, the developments in the next chapters have continued to focus on the layers up to, and including, the high-level application interface, rather than the development of any specific client applications.

The Aroo system also demonstrated the possibility of extending the prototype discussed here with other technologies to add extra functionality and allow it to be integrated with other systems. An example of this was adding ARTags to the folders so that they could be tracked on an augmented tabletop. Such augmentations are not limited to folders, and could equally be applied to the documents stored within them. For example, this system could be combined with an RFID document tracking system (such as described by Arregui et al., 2003). In this case, the system described in this chapter would provide a physical user interface embodied in the folders, while RFID tracking would allow the system to determine which documents are in each folder

6.5.1 Limitations

From the development of this prototype a number of insights were gained that would help the design and development of future prototype systems. These insights mostly came from observing the areas where the system had limitations. Such limitations would need to be overcome in order to develop a full-scale solution, and are discussed below.

Application constraints and scalability limits

A significant limitation of this prototype system is its small scale, and to some extent application specific design, which means that it would not be suitable for deployment in many offices. As was evidenced in Chapter 4, the physical infrastructure used to manage documents varies both within offices and between different offices. Vertical folders, as used by the prototype, are only one of a number of different type of document containers. Other commonly used document containers include lateral folders, manila folders, ring binders, and archival boxes.

Additionally, the system was not designed with scalability as a priority. Since it was designed as a proof-of-concept, it was only made to support a single cabinet with several folders in it. However, a system deployed in the real-world offices would need to scale to support an arbitrary number of folders that could number into the hundreds or thousands.

Based on these observations, it was decided that a future system would benefit from a more abstract design that could be applied to various possible physical implementations, and would scale to an arbitrary large size.

Hardware and circuitry limitations

The prototype has several other limitations caused by the hardware and circuitry used in its implementation. These limitations stem from the design of the connection method between the folders and the cabinet, and limitations of the 1-wire protocol.

Several of the prototype's limitations are due to the use of 1-wire technology for communication with the folders. Because 1-wire slave devices are designed to draw power through the data line, they need to keep their power consumption to a minimum, and therefore 1-wire slave devices only provide simple functions such as memory, GPIO, analog-to-digital conversion, and do very little internal processing. There are a number of cases, however, where it is beneficial to have more processing capacity available at the slave device. For example, to implement a continuously blinking LED using a DS2406 1-wire device would require the 1-wire master to send a command to the slave every time the LED needs to change its state, rather than simply sending a 'blink' instruction to the slave device.

A similar issue is that the communication protocol used by 1-wire is limited in the speed it can operate at. The visible implication of this to the user is that there can be a noticeable delay between the occurrence of an event (e.g., removal of a folder from the filing cabinet) and the detection of this event by the application software.

The limited power available to the folders reduces the range of user interface elements that can be incorporated into them as well. While a single LED provides the ability to draw attention to a folder and show its location, it is limited in the functionality it can support. For example, it may be desirable to display searches from different users simultaneously, or to have a different colour LED for displaying alerts.

Furthermore, the implementation of the hooks and rails method of electrical connection between the folders and the filing cabinet in the prototype system was found to be unreliable. As such folders are not always detected as present in the cabinet, or spurious removal and addition events are generated. Similarly, if a folder is in the cabinet, but the folder contacts are not fully in contact with the rail then the folder is not detected as being present in the cabinet by the system, and therefore its button and LED are unusable. This problem is largely due to the prototype nature of the system, and could likely be overcome by, for instance, experimenting with a more suitable profile for the hooks and rails, and a better combination of materials. However, such a design would still be specific to a vertical filing cabinet, and would still be limited to a two wire interface (as there are only two rails).

6.5.2 Designing the next generation

From the discussion of the limitations above, it is possible to derive a list of requirements and specifications for the design and development of the next generation prototype system.

The implementation of Aroo has shown that it is possible for a third-party to easily adopt the system and integrate it with their own client software applications. For this reason, the design of the next generations of the systems presented in this thesis continue to focus on providing the underlying physical infrastructure, along with a generic application interface to interact with it.

In summary, in order for the next generation of the system to be useful in a real-world office environment it needs to be more flexible in its support of document management artefacts, scalable to support a large number of artefacts, and more responsive in terms of communication speed. Additionally, support for more advanced user interface elements on the folders would be beneficial. To be able to provide this, the method of providing physical connectivity needs to be revised, and alternatives to the 1-wire protocol considered.

6.6 Summary

This chapter described a proof-of-concept prototype system for integrating paper and electronic document management that takes a different approach to previous systems. The system described in this chapter manages documents at the folder level, rather than as individual documents, and is thus able to add features to those folders that have not been provided together in other existing systems: the ability to detect the presence of a folder in a filing cabinet, the ability to detect user interaction with it via an input device, and the ability to display its location to the user using an LED. However, this system has several limitations, which prompted the design of a second generation prototype, which will be described in the next chapter.

Chapter 7

Architecture and Implementation of a Revised Prototype System

In the previous chapter a prototype system for integrating management of paper and electronic documents was described. The chapter ended with a discussion of the system, which noted that the prototype had successfully demonstrated the feasibility of such a design, and described the significant findings gained from the development and testing of the prototype.

One of the most significant findings from the development of the previous prototype was the need for a more flexible design in order to be able to support a wide variety of document organisation schemes suitable for different types of offices. Taking this into account, a new architecture was designed that follows the same principles as the previous prototype, but abstracts away application specific details in order to create a framework in which similar types of integrated paper and electronic document management systems can be developed. Based on this architecture, a prototype system was developed. This new prototype system is called SOPHYA, which is short for 'Smart Organisation of PHYsical Artefacts' and alludes to Sophia, the Greek goddess of wisdom.

The approach taken by the prototype described in this chapter is similar to the design specifications given by Stanfield et al. (2007) in their patent. However, the existence of this patent was discovered after the design and development of the system described here, which was carried out independently.

This chapter begins by outlining the motivation for developing this prototype system (Section 7.1). This is followed by a description of the architecture (Section 7.2), and discussion of the implementation of SOPHYA (Section 7.3). The chapter concludes with discussions (Section 7.4) and a summary (Section 7.5).

7.1 Motivation

The prototype described in the previous chapter was developed to investigate the feasibility of the potential technology, rather than aiming for the flexibility to support the varying needs of real-world offices. While it fulfilled its purpose, it also showed a number of areas that could be improved. Therefore, the main objectives of developing a second prototype were to improve the following aspects of the system:

- flexibility, so that it could support a wider variety of applications;
- scalability, so that it could support arbitrarily large document collections;
- reliability, so that it would be more robust and less prone to failure;
- usability, so that it would respond to user interaction more rapidly (e.g., reducing input latency); and
- functionality, by supporting more advanced levels of user interaction.

The design of the previous prototype did not make a conscious effort to separate the application specific elements of the design (e.g., document management software and document container type) from the elements that would be common across all applications (e.g., the electronic hardware), instead its focus was on a single application. However, as was evidenced by the study presented in Chapter 4, the physical artefacts that are used to manage documents vary both within offices and between different offices. Commonly used artefacts include not only vertical folders as in the previous prototype, but also lateral folders, manila folders, ring binders, and archival boxes. Additionally, the requirements of the users of an integrated document management system vary between offices in a number of ways, including levels of access control, tracking, and logging that are required.

Because of this variation between offices it cannot be assumed that a ‘one size fits all’ approach to system design will be successful. Therefore, the approach that has been taken in the development of the system described in this chapter has been to design a basic high-level architecture that separates the elements that may vary between different offices (e.g., the choice of document management artefacts, and the electronic document management system) from those elements that are not application specific (i.e., the electronic hardware and low-level software used to provide the link between the physical artefacts and the electronic document management system). The design of this architecture is described in the next section.

7.2 Architecture

In order to meet the objectives described above, an architecture was designed that allows the separation of the application specific elements of the physical and electronic document management system from those that are common across different applications. To achieve this, the architecture is separated into four layers, shown in Figure 7.1.

Physical interface. This layer includes the physical and electronic hardware required to integrate the management of paper documents and their associated infrastructure with that of electronic documents.

Middleware. Provides a layer of abstraction between the Document Management System (DMS) software and the physical interfaces.

DMS software. The layer where management of electronic documents occurs, and where electronic content is mapped to paper documents.

Client software. The software that the user interacts with.

Each of the component of this architecture is described in more detail below.

7.2.1 Physical interface

The physical interface has two main components: the physical document management hardware such as folders and filing cabinets, and the electronic hardware that the physical hardware is augmented with. In order to provide support for as many different types of physical document management systems as possible, the architecture is designed such that it is not specific to a single type of document container (e.g., the vertical folders of the previous proto-

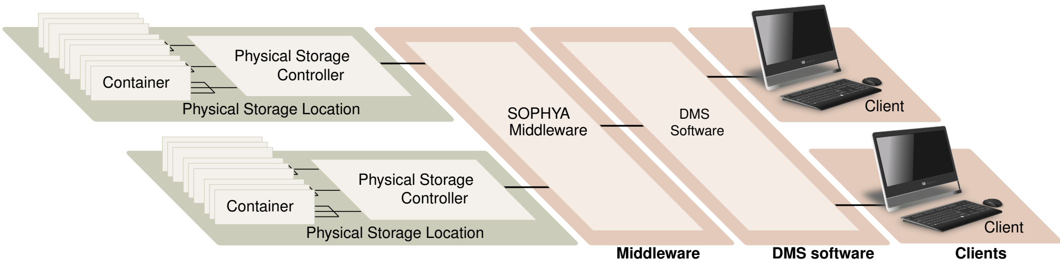

Figure 7.1: Depiction of the architecture of the system as discussed in this Section 7.2. The architecture has four major components: the physical interface and middleware (which are common), and the DMS software and clients (which are application specific).

type). Rather, the architecture is based around the concept of a generic document container, which could, for example, include folders, binders, archival boxes and books. These containers are augmented in a similar manner to the previous prototype. Furthermore, rather than focusing on a single type of storage location for these containers, a range of physical storage locations (e.g., filing cabinets, shelves, and desktops) would be supported . These would be augmented with electronic circuitry, and provide an interface between the containers and the document management software. The physical interface component of the architecture is shown in more detail in Figure 7.2.

Similar to the first prototype, the electronic circuitry that the document containers are augmented with provides them with a unique ID, and some user interface devices. However, the circuitry used in this case is completely different to that of the previous prototype, and has been designed specifically to support this architecture. The new design of the circuitry allows it to be more flexible in its support of functionality (e.g., its memory and user interface devices). As shown in Figure 7.2, a container may include some form of User Interface (UI) devices, some memory, as well as the necessary *container logic* (the circuitry and firmware that is responsible for interfacing between these elements and the physical storage location).

Containers placed in a physical storage location connect to a container network (this is separate from, and not to be confused with, the Local Area Network, LAN). A single storage location may include an arbitrary number of container networks. Every container network has a single container network controller, which acts as the arbitrator for that network. It is responsible for polling

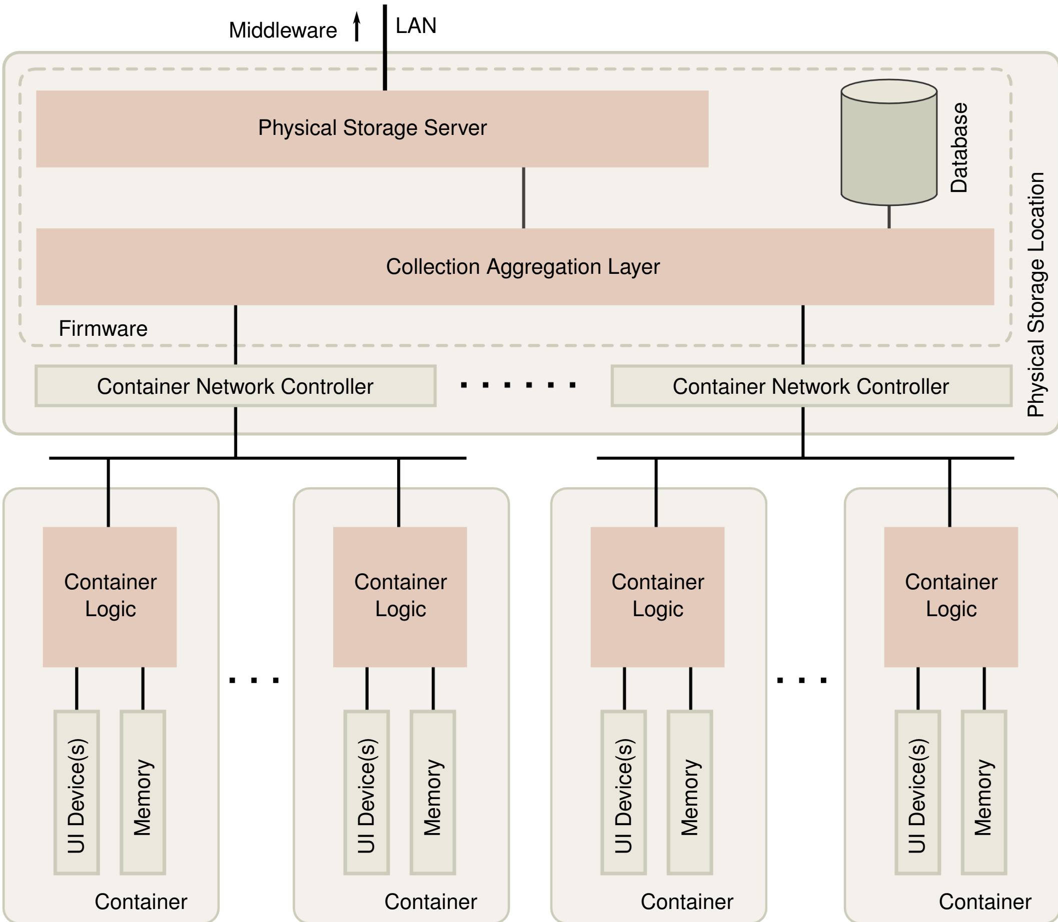

Figure 7.2: Detailed view of the physical interface sub-architecture, showing how containers such as folders connect via 'container networks' to the physical storage location, which is in turn connected to the middleware via the LAN.

the container network to detect the addition and removal of containers, as well as communicating with the containers in order to access their UI devices and memory. Each container network has a unique ID that enables it to be differentiated from other container networks.

The physical design and number of container networks belonging to a given physical storage location is dependent on the physical hardware of the containers and the storage location. Although it is possible to determine which containers are on a given container network it is not possible to determine their position or order, and therefore having a larger number of smaller networks, rather than fewer large networks, would result in more fine-grained location determination. Additionally, a lesson learnt from the development of the previous prototype was that it is better to keep the number of containers on the network bus as small as possible to provide better reliability. For these reasons, a physical storage location would typically have several smaller container networks rather than a single large one.

The next layer of the sub-architecture, the *collection aggregation layer* is responsible for aggregating the information coming from the individual container network controllers. The collection aggregation layer maintains a database of containers, storing metadata about each, which is updated whenever the addition or removal of a container is detected. The information about each container stored in this database includes:

- whether or not it is currently present in the storage location;
- the ID of its current container network (if that container is present) or the last container network that it was on (if it is now absent); and
- the time it was last added or removed.

The top-most layer of the sub-architecture, the *physical storage server*, provides an interface between the collection aggregation layer and the LAN. Higher-level software connects to the physical storage server via the LAN and is able to query the collection aggregation layer in order to access information about the containers. Through the physical storage server, the software is able to:

- get a list of all container networks belonging to the physical storage location;
- get a list of all containers present in the physical storage location (or any container network of the storage location);
- read container metadata;
- control the user interface devices on the containers; and
- receive events when containers are added or removed.

At the physical interface layer there does not yet exist any mapping between the physical containers and electronic information (aside from some basic metadata described above), rather each container is simply referred to by its unique ID. The reason for this is that the physical interface layer is designed to be as generic as possible, such that it could be deployed in any office regardless of the types of containers being managed or the software being used to manage them. The mapping of electronic information to the IDs of the physical elements is carried out in a higher layer.

7.2.2 Middleware

The middleware layer sits between the physical interface and the application specific document management software (as shown in Figure 7.1). Its function is to provide the document management software with a unified view of the arbitrary configuration of physical interfaces that make up the physical part of the document management system.

Although it is possible for the DMS software to connect directly to physical interfaces, the benefit of having a separate middleware layer is that it provides greater flexibility for expansion. For example, if a new type of physical interface was implemented that provided different features, the middleware would make it easier to integrate this with the document management system, since it could present a consistent interface to the DMS software.

Similar to the other parts of the architecture that aim to be generic rather than application specific, the middleware does not attempt to associate electronic information with the physical elements. It is simply responsible for mediating the communication between the physical interfaces and the document management system.

7.2.3 Document management software

This layer is application specific and is where electronic information is mapped to containers and physical storage locations. It is also responsible for making this information available to the next layer up, the client layer. Its implementation depends largely on the office it is to be deployed in. For example, in offices with an existing DMS it may be integrated with this in order to provide integration between management of paper and electronic documents.

7.2.4 Clients

In order to access the information stored by the document management system, some form of interactive client software is required. For smaller applications, where access is only required from a single computer, the client may be implemented as part of the document management system. However, in most cases, especially those where documents are being shared between multiple people, multiple clients will be required.

Client software may take on a variety of forms. It may simply be part of a desktop file-management system, allowing users to access both the electronic documents and the physical containers' electronic representations from their workstations.

Alternatively, clients could be embedded into physical storage locations themselves. For example, a filing cabinet may be augmented with a touchscreen interface with client software. This could allow the users to browse or search through the contents of the folders within the cabinet until they find the ones they are looking for, at which point they could, for example, activate LEDs on the folders to show their locations within the cabinet. Furthermore, touch sensors on the folders could allow the user to select a folder by touching it, and have information about that folder be displayed on the output display of the cabinet. These potential extensions of the system are discussed in Section 11.3.2.

Similarly, with the proliferation of mobile computing platforms such as smart phones and tablets, it would be possible to have portable clients that can browse the contents of paper document repositories within an office through a wireless network.

7.3 Implementation

This section describes the implementation of SOPHYA, which is based on the architecture described above. SOPHYA is an extensible system developed to demonstrate the capabilities of the architecture. It incorporates several major changes from the prototype system described in the previous chapter, including a redesign of the physical hardware, the development of a new communication protocol to replace the 1-wire protocol, and the design and implementation of the corresponding circuitry and firmware. These are changes are discussed in more detail below.

7.3.1 Container network protocol

The protocol used for communication on the container networks is an important factor in providing the functionality required by the system. In order to meet the requirements of this prototype, the communication protocol needs to include support for:

- bi-directional communication over a single wire (bus topology);
- master/slave topology with a single master and multiple slaves;
- dynamic discovery of the addition and removal of slave devices on the network; and
- transmission and reception of data to and from the slave devices.

Although a number of existing protocols were first considered, none of them were able to support the needs of SOPHYA. When considering existing communication protocols, the most obvious candidate was the 1-wire protocol, that was used in the previous prototype (see Section 6.2.2). However, the 1-wire protocol has several limitations, chiefly its communication speed is limited (which results in latency when detecting user interaction), and the 1-wire slave devices that are available are limited in their functionality and processing capability.

No other well known protocols were able to meet the requirements listed above. For example, most protocols did not provide bi-direction communication without requiring multiple data lines, or lacked dynamic discovery of slave devices. Therefore, a new protocol was designed specifically to support the architecture described above.

The protocol that was developed for this prototype is described in detail in Appendix B. The most significant deviation, in terms of hardware, from the 1-wire protocol used in the previous prototype is that it requires three wires, as opposed to two. Therefore, this had to be taken into account when designing the containers and physical storage locations. The three wires are used for ground, power, and data.

7.3.2 Containers

SOPHYA bears some similarity to the previous prototype in its use of folders as the containers. However, it differs from the previous prototype in its use of lateral, rather than vertical, files. This change was made in order to overcome the limitations of the hook-and-rail method of contact encountered in the previous prototype. The significant limitations of the hook-and-rail method, as described in the previous chapter, were that it did not provide a reliable medium for communication and was limited to two points of electrical contact between the filing cabinet and the folders (i.e., the two rails).

SOPHYA uses conductive strips that run around the outside of the folder to provide electrical contact between the folders and the storage location they are placed in. The current design of SOPHYA makes use of three conductive strips in order to support the communication protocol described above. These are made from adhesive aluminium tape, as it is easy to attach to existing folders. However, while this has worked satisfactorily for the prototype that has been implemented, it would not be suitable for extended use as the opening and closing of the folder soon fatigues the metal causing invisible breaks along the fold. Figure 7.3 shows a comparison of a folder of the previous prototype design with those used by SOPHYA.

The conductive strips run around the outside of the folders, both on the front and the back, and because of this, when a folder is placed in a storage location its conductors will make contact with not only the conductors of the storage location, but also those of the folders in front of and behind it. This provides greater surface area and reduces the likelihood of a folder losing its connection to the container network.

Another benefit of the conductors wrapping around the folders is that it makes it possible to stack folders while maintaining their connectivity to the container network, as shown in Figure 7.4, so, for example, if a physical desk surface is augmented with SOPHYA storage location circuitry, folders can be stored as a pile in plain view, retaining their accessibility and ability to remind, while remaining able to be located electronically and having their user interfaces still able to function (enabling them to, for example, provide visual indication of

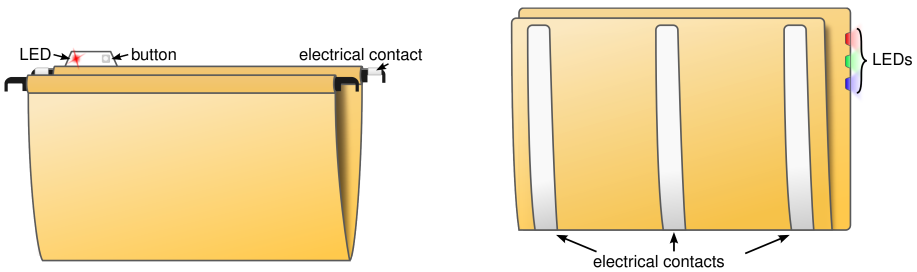

(a) Previous prototype folder (vertical). **(b)** SOPHYA folder (lateral).

Figure 7.3: Comparison of previous prototype folder and SOPHYA folder.

their location).

The current implementation of the folders supports a simple user interface in the form of three different coloured LEDs (red, green, and blue). These can be toggled individually, and turned on in any combination. This implementation does not include an input device such as a button. However, if this was required in future, the folder circuitry and container network protocol could easily support the addition of a button, touch sensor, or other input as well as output devices.

Two variants of SOPHYA containers have been constructed. While they are the same electrically, they are based around different physical folders. The first design, shown in Figure 7.5a, is based on a cardboard Codafile folder, while the second, shown in Figure 7.5b, is based on a plastic clamp folder.

Although the folders of the SOPHYA may look similar to the folders of the previous prototype, the design of their electronic circuitry is in fact totally different. In order to support the new communication protocol the electronic circuitry has been completely redesigned. This involved the design of custom circuit boards and firmware for the folders and storage location, which is discussed in Appendix C, Section C.1.1.

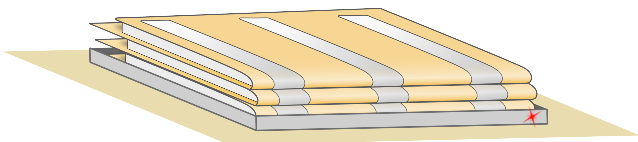

Figure 7.4: Folders are able to be stacked while remaining in contact with the storage location since the conductors wrap around them.

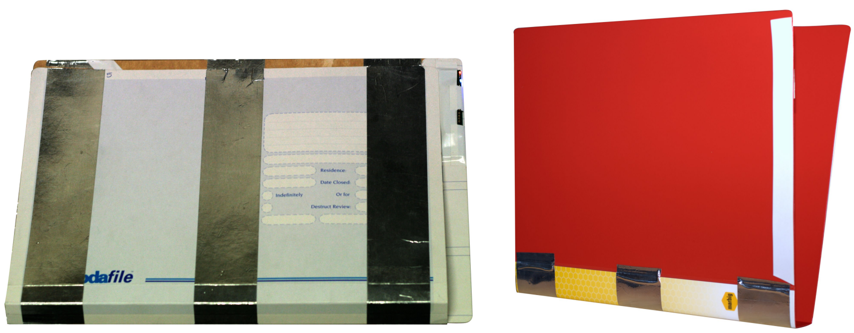

(a) Codafile **(b)** Clamp folder

Figure 7.5: Two variants of SOPHYA containers that have been constructed.

7.3.3 Storage locations

Four different types of SOPHYA physical storage locations have been constructed as part of this prototype. These are filing cabinet, shelf, document tray, and desktop, as shown in Figure 7.6.

The filing cabinet and shelf store folders with their spines contacting the conductive strips within the storage location, while the document tray and desktop are designed to have folders stacked on them and make contact with the front or back of the folder. In order to ensure that folders stand up when placed on the shelf, adjustable plastic dividers hold them in place.

Conductors

The design of the conductors that make contact between the containers and the storage network is a significant factor in ensuring that the communication is reliable. Having a third conductor requires that all three conductors are aligned correctly because any significant variation will result in not all conductors

(a) Filing cabinet

(b) Shelf

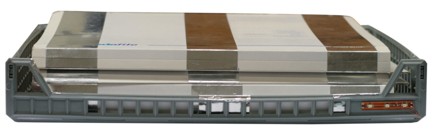

(c) Document Tray

(d) Desktop

Figure 7.6: Four variants of SOPHYA storage locations that have been constructed.

making contact (as shown in Figure 7.7).

Although ensuring the correct alignment of the conductors is mostly a factor of container design, the choice of material used in the storage location also has a bearing on the reliability of the connection. Initially aluminium tape was used for the storage location conductors, and this worked well for the document tray and desktop units where a larger area of the folders and storage locations would make contact. However, the aluminium tape was not reliable in the case of the filing cabinet and shelf, where the contact is made by the spine of the folder only. The problem with the aluminium tape was that it was too thin, and any slight misalignment of the folder conductors would result in unreliable contact. After some experimentation, braided wire (e.g., grounding strap, shown in Figure 7.8) was used for the conductors.

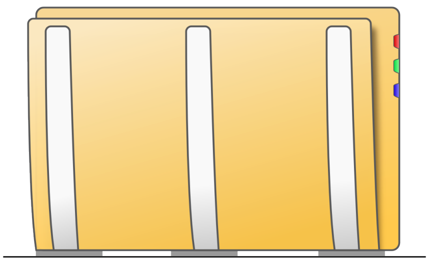

(a) Correctly aligned conductors resulting in a good connection.

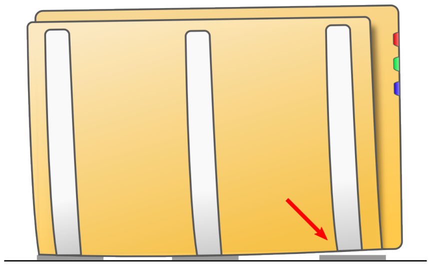

(b) Incorrectly aligned conductors resulting in an unreliable connection.

Figure 7.7: Conductors need to be correctly aligned as any significant variation will prevent all three conductors from making contact.

(a) Conductors in the filing cabinet.

(b) Close-up of conductors used on the shelf.

Figure 7.8: Conductors used in the filing cabinet and shelf implementations of SOPHYA.

Physical storage controller

The physical storage controller is the electronic part of the physical storage location. It comprises the container network controllers and firmware parts of the sub-architecture shown in Figure 7.2. For this implementation the physical storage controller was implemented as a stand-alone unit (Figure 7.9) that could support up to three container networks at a time.

The control logic of the container network controller is implemented in a Field Programmable Gate Array (FPGA, see Kuon et al., 2008). The FPGA and associated circuitry that makes up the container network controller is laid out on a circuit board that was custom designed for this purpose. More details of the design of the container network controller can be found in Appendix C, Section C.1.2.

The firmware that provides the collection aggregation layer, database, and physical storage server is executed by an embedded system. The system used in this prototype was a Gumstix Verdex (see Gumstix, Inc., 2007) embedded computer running the Linux operating system. The container network controllers connect to the verdex via USB. The verdex includes support for connecting to Ethernet, WiFi, and/or Bluetooth networks, and this is used to provide LAN connectivity for communication with the higher-level software.

The physical storage controller firmware is implemented as a Linux daemon. It continuously receives the output generated by the container network con-

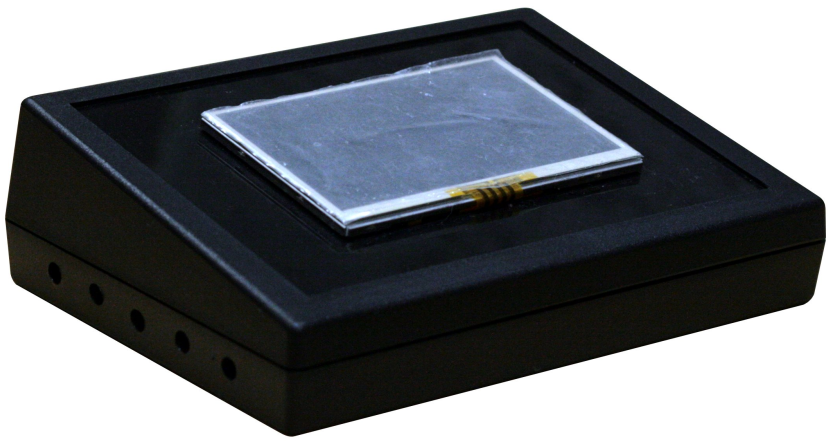

Figure 7.9: The physical storage controller unit. This unit is able to control up to three container networks, and includes an embedded GUI for testing and demonstration.

trollers as they poll their container networks. This is filtered and processed in order to detect when containers have been added to or removed from a container network. When this is detected the internal database of the firmware is updated, and an event message is transmitted to the higher-level software.

Communication with the higher-level software is via the LAN. The firmware operates a TCP/IP server which the software can connect to. A simple text-based protocol allows the software to read container information from the database, receive event notifications when containers are added or removed, and control the user interface devices of the containers. The current implementation of the prototype does not provide any security in terms of access control or encryption of communication. However, this could easily be added, for example using Secure Shell (SSH) for authentication and tunnelling.

7.3.4 Middleware

The purpose of the middleware is to allow the application specific DMS software to connect to multiple physical storage locations without requiring it to maintain connections to the individual storage locations. The middleware uses the same text based communication protocol as the firmware, and so is effectively transparent. Because only a single physical storage controller has been implemented for this version of the prototype, a middleware was not initially developed.

7.3.5 Higher-level software

As discussed in Section 7.2, the DMS and client software elements of the high-level architecture are application specific. The focus of this chapter, however, has been on the development of a prototype that implements elements of the architecture that are common across different applications. Therefore, the prototype described in this chapter does not include a DMS or client software. Instead, it provides an API that enables it to be integrated with such software.

Since, it was necessary to develop some software in order to be able to test and demonstrate the system, a simple software application was developed to perform these tasks. This software application had two main functions: displaying the status of the physical interface (e.g., presence or absence of containers, current or previous location of containers, etc.), and controlling the user interface devices (i.e., LEDs) on the containers. A screenshot of this application is shown in Figure 7.10.

The physical storage controller also included some interactive software for testing purposes, as shown in Figure 7.11. The physical storage controller included an LCD touchscreen so that users could interact with software running on the Verdex embedded computer. The application that was implemented was able to connect to the physical storage server and would display, in table form, the information stored in the database of the physical storage controller. It also enabled the user to control the LEDs on the containers.

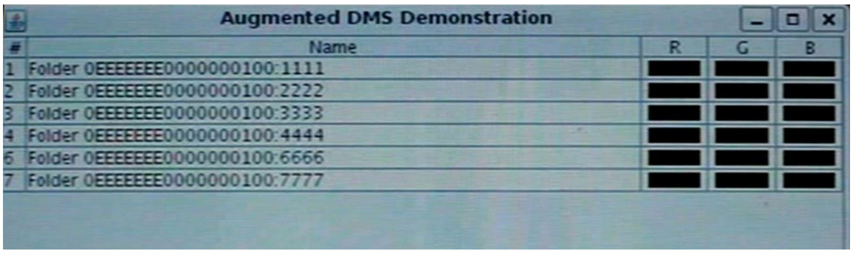

Figure 7.10: Screenshot of the software application that was developed for testing and demonstration of SOPHYA on PCs

Figure 7.11: Screenshot of the software application that was used for testing and demonstration of SOPHYA running on the physical storage controller.

7.4 Discussion

This chapter described an architecture that has been designed to separate the application specific components of an integrated physical and electronic document management system from those that are common across different applications. Based on this architecture, a prototype system was designed and built to implement these common components, as well as some application specific components for demonstration. This section discusses some of the observations and knowledge gained during the design and implementation of this prototype.

7.4.1 Changes to user interface components

The buttons that formed part of the user interface on the folders of the previous prototype are no longer present on the containers of this prototype. The decision not to include the button as part of the container user interface was made for several reasons. Aesthetically, the size (and weight) of the button, and the fact that it is a mechanical part, meant that it would not fit well with the container design, since the number of electronic components would ideally be kept to a minimum. Additionally, removing the button from the design simplified the construction of the container circuitry, which was important since it was assembled manually for this prototype.

It is possible to achieve the function of the button through other means. For example, the front edge of a folder could be augmented with capacitive touch sensing. Alternatively, user interaction equivalent to a button press could be provided by the user lifting a container (to break contact) then placing it back again after a short delay. However, since such functionality was not required for testing the fundamental design of the system, it was not implemented.

7.4.2 Testing

As described in Sections 7.3.2 and 7.3.3, several different types of containers and storage locations have been developed using the technology described in this chapter. During their development these prototypes have undergone a continuous process of testing and revision in order to improve reliability and ensure that they function correctly. The testing has been conducted at several levels, from low-level testing of communication between containers and storage locations using an oscilloscope, to high-level testing using a specifically designed software application (as described in Section 7.3.5). The primary aspects of the system that was tested were the reliability with which containers were detected by the system, both when they are initially placed into a storage location and while they remain in the storage location. Another aspect of the system that was monitored during testing was the time it would take to detect addition and removal of containers.

7.4.3 Reliability and responsiveness

As previously discussed, the construction of the physical hardware and conductors of the containers and storage location has a significant impact on the reliability of communication between the containers and storage location. Reliable communication is important to ensure that containers are detected in the storage location in which they are placed, and to enable their user interfaces to function. The clamp folders (Figure 7.5b) and shelf (Figure 7.6b) are the most recently implemented containers and storage location, and are the most reliable constructed to-date. Their reliability was suitable for conducting an evaluation of the system, which is described in Chapter 10. It is envisaged that if such a system were to be deployed commercially, the reliability of communication could be further improved through industrial design (which is beyond the scope of this thesis).

Closely linked with the reliability of the system is its responsiveness to container addition and removal events. Addition events can be detected and communicated to high-level software with a latency that is almost unnoticeable to users. This is because the container network controller is continuously polling the container network, and completes several iterations of this polling loop every second. When a container is added to the storage location it will be detected by the next polling cycle. Since each container ID includes a checksum, whenever a valid container ID is detected that was not previously present, then this is clearly a container that has been added to the storage location.

However, the detection of removal events is more difficult, and more closely linked to the reliability of the contact between the containers and the storage location. A naïve approach to detecting removal events would be to declare a container as removed when it is not detected in a given polling cycle. However, it is common for a container to 'disappear' for several polling cycles before reappearing, and therefore if such an approach is taken container addition and removal events are generated every time this happens. Therefore, the system needs to ensure that the container is not detected for a number of cycles before declaring it to be removed. The selection of the number of cycles is a trade-off between responsiveness to the actual removal of a container on the one hand, and generation of spurious removal events on the other. The current prototype implementation has selected a number of cycles that provides few spurious removal events and detects removal within several seconds. This has been sufficient for the testing and evaluation of the system to-date and could be improved on by increasing the reliability of the system as described above, and enhancing the communication protocol as described in the next section.

7.4.4 Limitations

One limitation of the system is related to the communication between folders and storage locations. With the current version of the prototype the data rate of the protocol is limited. This is partly due to the physical properties of the conductors, and the fact that a conservative bitrate has been used in order to ensure reliability. While the current bitrate of the protocol is enough to support detection of folders and controlling a simple user interface device on them, it would not be able to support accessing a large amount of memory on the folders. One method of overcoming this would be to use a hybrid approach, with wired communication providing power and location information, and wireless technology providing higher data rate communication.

Another limitation of the SOPHYA prototype is that, while it is able to determine which containers are currently present on a container network, it is unable to determine their order, or precise location in relation to one another. This information may be useful in situations where the ordering of the physical containers is important, such as books on a library shelf. Including this ability would allow the system to determine when containers have been placed incorrectly (e.g., a book is placed out of order), and could be combined with display on the storage location to assist users when placing containers. A method of providing this, and a prototype system that implements it, is discussed in the next chapter.

7.5 Summary

In order to overcome the limited scalability and flexibility of the previous prototype an architecture was designed which allowed the separation of application specific parts of the system from those that were common across all applications. Thus the common parts could be developed for use with application specific parts as necessary. A prototype system based on this architecture was implemented. It included the common parts of the design (the physical storage controller and container circuitry), as well as application specific physical hardware (folders, filing cabinet, etc.).

To improve reliability and allow for more advanced user interfaces a new method of communication between the storage location and containers was designed. A new communication protocol that uses three conductors was designed and implemented. Additionally, a new method of making contact between containers and storage location was implemented.

Chapter 8

Prototype System for Managing Ordered Document Collections

The two previous chapters described the design and implementation of two successive systems for integrating paper and electronic document management. The design of these systems was based on the augmentation of document containers (e.g., folders), and the storage locations in which they are kept, with electronic circuitry in order to give them an electronic manifestation and enable them to act as an embodied interface to electronic document management software.

One of the limitations of these systems, however, was that they were unable to determine the ordering of containers within a storage location. They were able to determine the location of containers with coarse-grained accuracy (e.g., the drawer of a filing cabinet which a container was in), and this enabled them to assist users when retrieving containers using this information combined with the output display components of the containers.

The design and development of a third prototype system is described in this chapter. This system builds on the work done in the development of the previous systems, but adds the ability to determine the position (and therefore order) of containers within storage locations. Since this system follows a similar high-level architecture to the previous prototype it is also referred to as SOPHYA. In order to differentiate between the two SOPHYA systems, however, the system described in this chapter will be referred to as the *ordered* SOPHYA system (since it is designed to manage ordered collections of containers), while the system described in the previous chapter will be referred to as the *unordered* SOPHYA system from the point on.

This chapter begins with a discussion of the motivation for developing a system that is able to determine ordering of containers (Section 8.1), followed by a descriptions of the design (Section 8.2) and implementation (Section 8.3) of the ordered SOPHYA system. Finally, the chapter ends with discussion (Section 8.4) and summary (Section 8.5).

8.1 Motivation

The prototype described in the previous chapter did not provide a means of detecting the position or ordering of containers within storage locations (see Section 7.4.4). However, this limitation was mitigated by the inclusion of an output display component on the containers, making it possible for users to quickly locate specific containers without relying on their ordering within the storage location. Therefore, containers could be placed into the storage location without regard for their order (a form of piling, as discussed in Section 2.3.1), and the system would aid the user in retrieving them. There are, however, a number of situations where maintaining order of items in physical collections is necessary or beneficial beyond simply facilitating the retrieval of a given item.

An example of such a case, which was introduced at the end of the previous chapter, is the management of books in a library. The strict ordering of library books serves not only to facilitate their location and retrieval by library patrons, but also to keep related books together. This enables efficient retrieval of multiple books of the same topic, and allows browsing of related books on the shelves—potentially leading to serendipitous discovery of useful material. Therefore, a library would very likely gain more benefit from a system that is able to detect and help maintain the ordering of books on the shelves, as opposed to a system that discards ordering information and follows a piling strategy.

As the library example shows, there are benefits to be gained from maintaining the ordering of items within a collection, in this case the ability to browse and retrieve related items. There are other benefits that are gained from maintaining a strict organisation—as opposed to the flexible organisation afforded by piling—including the ability to support spatial memory by ensuring that items are always stored in the same place. This could be especially useful in environments where multiple people share a collection, and therefore the person who returns an item to the collection may not necessarily be the next to retrieve it.

While the ability to determine the ordering of containers within storage locations enables the system to support strict organisation of containers, it does not require this. It is possible for such a system to support a flexible piling strategy as in the previous prototypes, while making use of the fine-grained position information in the document management software. This information could, for example, enable the development of visualisations that allow users to browse through the collection on a computer, displaying containers in their current physical ordering, or it could enable software to be developed that detects relationships between containers by considering which ones are frequently placed together.

Similarly, the ability to determine the order of containers allows for new possibilities for the design of the tangible user interface components of the system. For example, where previously the output display components of the user interface were embedded into the containers (Figure 8.1a), the addition of position determination to the system enables these to be embedded into the storage location (Figure 8.1b). This makes it possible for the output display to function even when containers are absent, for example allowing the system to indicate where an absent container should be placed on the shelf.

8.2 Design

This section covers the theoretical aspects of the design of the next prototype system. It begins with a discussion of how the ability to determine the position of containers allows greater flexibility for the design of the user interface elements of the system. This is followed by a description of the method used for determining the positions of containers within storage locations, and then a discussion of the design of the electronic hardware required for this method.

8.2.1 Method for determining container positions

The systems described in the previous chapters were unable to determine the positions of containers with fine-grained accuracy. However, with some mod-

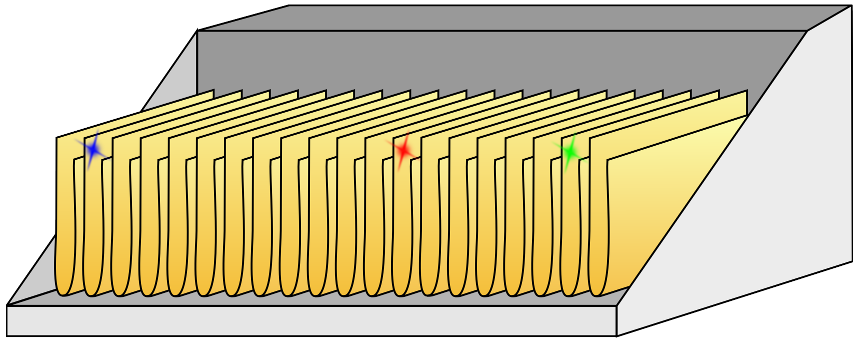

(a) Ordering unknown, user interface output component on containers.

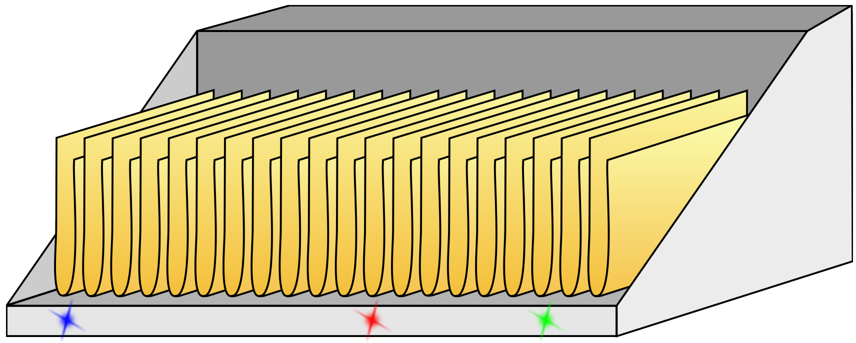

(b) Ordering known, user interface output component on storage location.

Figure 8.1: When the ordering of the containers is known, the user interface output elements can be integrated into the storage locations, rather than the containers. This enables new possibilities for information display to users, such as showing the location where a container should be placed when returning it to the storage location.

ification to their design, it is possible to add support for this feature. The SOPHYA prototype discussed in the previous chapter is able to provide some position information, but this is very coarse (e.g., which drawer of a filing cabinet a container is in). The design of this prototype is such that storage locations typically have several container networks, and each container network can support the connection of a number of containers (Figure 8.2a). However, while the system is able to determine which containers are on a given network, it is unable to determine their order.

The ability to determine the position of containers with fine-grained accuracy can be added by altering this design so that storage locations have many container networks, each of which has no more than one container present on it at any given time (Figure 8.2b). Each container network has a unique location associated with it. Therefore, because it is possible to determine which container network each container is on, and the physical location of each container network is known, it is possible to determine the location of each container.

8.2.2 Physical interface architecture

In order to implement the method described above, some changes need to be made to the physical interface sub-architecture (the previous version of which is described in Section 7.2.1). The altered version of the physical interface architecture (Figure 8.3) is designed for a system where the output elements of the user interface are part of the storage location, rather than the containers (as discussed in Section 8.1).

The architecture is split into two independent components. One component is responsible for determining the positions of containers in the storage location, and passing this information to higher-layer software, while the other is responsible for controlling the output display devices of the physical storage location. The reason for this split is that the two components can then be physically separate when mounted on a physical storage location. For instance, the position modules could be placed near the back of the storage location (discussed in Section 8.2.3), while the display modules can be mounted in an easily visible location (e.g., the front of the storage location).

(a) Unordered system topology: containers in the physical storage location share a common bus network.

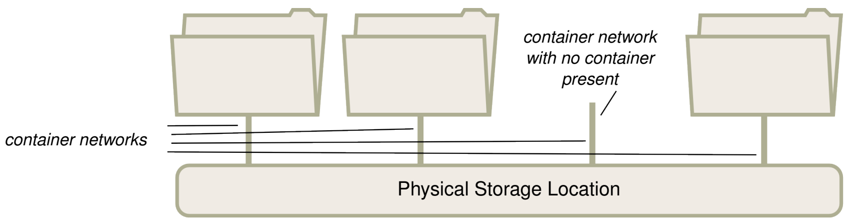

(b) Ordered system topology: each container in the physical storage location connects to a separate container network.

Figure 8.2: Using multiple networks with a single container per network, rather than a single network with multiple containers, allows the system to determine the position of a container based on which network it is currently present on.

Each of the two components of this architecture has two main elements: a controller and a number of modules. The modules are responsible for implementing the container network controllers (in the case of the position component) or the output display (in the case of the display component). The number of modules used can be varied depending on the size of the storage location. The modules connect via a bus network (e.g., I^2C, NXP Semiconductor, 2007) to the controller, which interfaces between them and the processing software layer.

Position detection

The modules of the position component of this sub-architecture each have multiple conductive pads (the exact number is an implementation decision). Each of these pads has corresponding container network controller hardware. The position module logic is responsible for continuously polling each of the

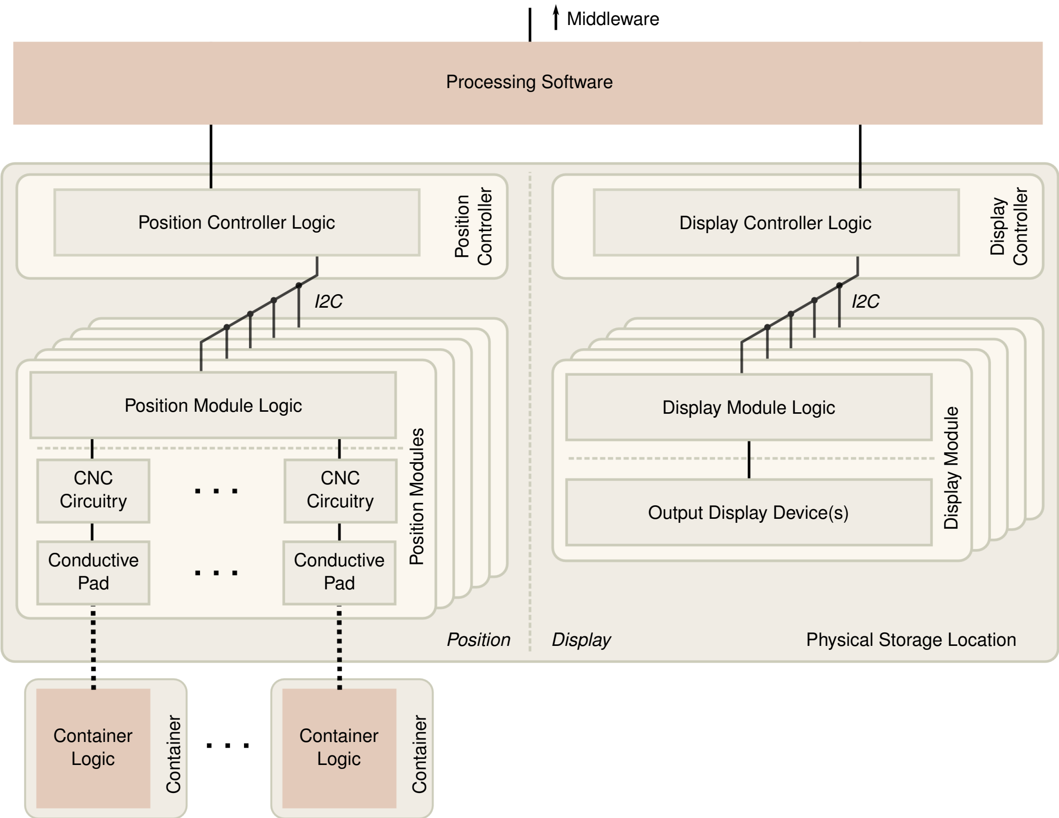

Figure 8.3: Architecture of the physical interface component of this prototype. This corresponds to the physical interface component of the architecture described in Section 7.2.

container networks in turn to detect present containers. The controller continuously polls the modules, reading from them the status of each of their container networks, and this information is passed up to the processing software.

Since the containers of this prototype do not need to support any user interface devices, their electronic hardware can be simplified in comparison to that of the previous prototypes. Therefore, each container only needs to have some read-only memory that stores a unique ID, and some logic to enable this to be read by the container network controller. This means that the container network protocol can also be simplified in comparison to the previous prototype. Furthermore, due to low power requirements of the container, a communication protocol that uses only two points of connection can be used.

Output display

The display component of the architecture follows a similar design to the position component. The major difference is that the display modules have output display devices (e.g., LEDs) in place of the container networks and location pads of the position modules. The processing software issues instructions to the display controller to set the state of its output devices. The display controller determines which display module the instruction needs to be sent to and passes it on. When the display module logic receives an instruction from the display controller it updates the state of the affected output display.

Scalability

The hierarchical structure that results from the design being split into controllers and modules makes it possible for the design to scale to arbitrary sizes, and therefore enables it to support arbitrarily large storage locations. The processing of the data from the location controllers can easily be parallelised. Each location pad can be considered independently (except in cases where a container bridges multiple location pads). Thus the processing of a storage location can be split up and performed in parallel as long as some care is taken to factor in the bridging of multiple pads.

8.2.3 Layout of conductors

The physical layout of the conductive pads that make contact between the containers and storage locations is an important consideration when designing this system. The number, size, and placement of the conductive pads defines the accuracy of the system. A top-down view of the layout of the conductive pads of the physical storage locations (also referred to as location pads) is shown in Figure 8.4a, and similarly the conductive pads on the underside of the containers are shown in Figure 8.4b.

As shown in Figure 8.5, the location pads are placed at the back of the storage location. This allows for containers that vary in length, as they will normally be aligned at the back of the storage location. The location pads are offset

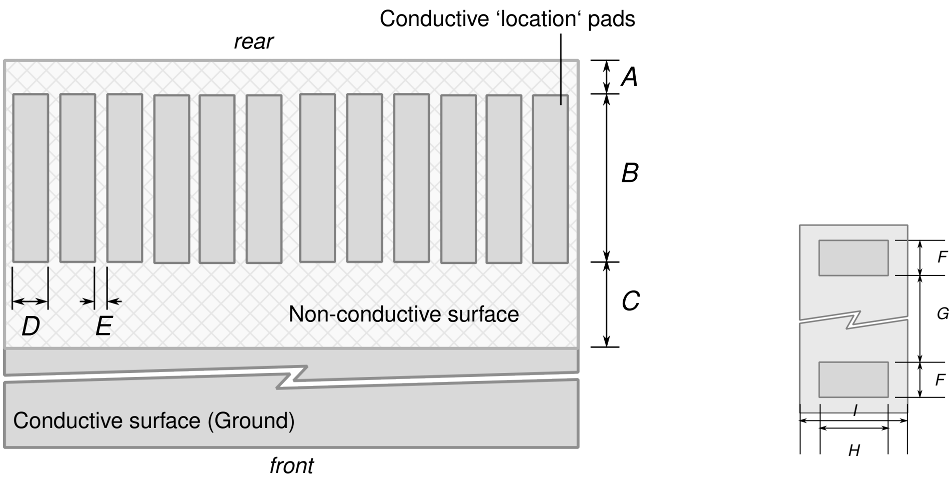

(a) Physical storage location **(b)** Container

Figure 8.4: Layout of conductive pads on physical storage locations and containers.

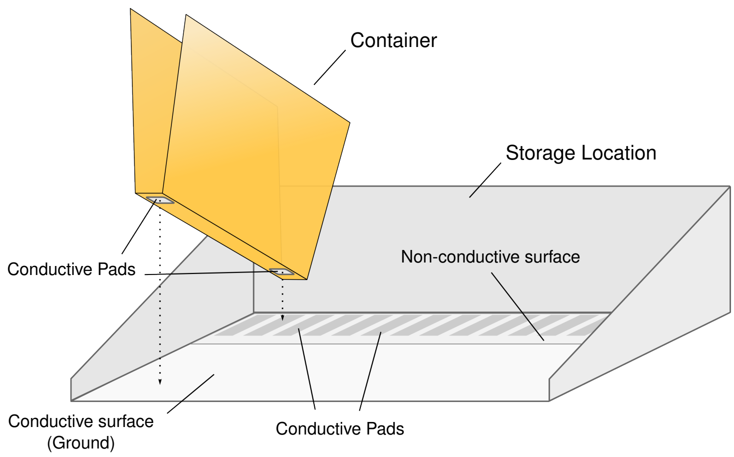

Figure 8.5: Placement of the conductive pads within the storage location.

by a small amount (A) from the back of the storage location as a precaution against short circuits between the container and the rear wall of the storage location.

The flexibility of the longitudinal position of the containers within the storage location is determined by the length of the location pads (B). The longer these pads are the more the containers can be moved forward or back within the cabinet and still be detected. A small width of non-conductive surface (C) prevents a short circuit if a container is placed with its conductive pad overhanging the location pad. The length of this non-conductive surface needs to be greater than the length of the conductive pads on the containers (F) to guarantee prevention of shorts.

To ensure that the containers' conductive pads are always in contact with a location pad, the width of the conductive pads on the containers (H) needs to be greater than the gap between the location pads (E). It is important to note that, as a result of this, it is possible for a container to contact more than one location pad. This needs to be taken into account when implementing the electronic hardware and the polling algorithm of the container network controllers (see Sections 8.3.2 and 8.3.3).

In order to be able to differentiate between all containers in a storage location, it is necessary to prevent more than one container from contacting a given location pad at a time. If there are multiple containers contacting a location pad then it is not possible to determine their individual positions. To prevent this from occurring, it is necessary for the width of the location pads (D) to be less than the minimum possible distance between the container conductors (which can be calculated by subtracting the width of the container conductors (H) from the width of the containers themselves (I)).

These layout considerations were taken into account when implementing the system described in the next section.

8.3 Implementation

This section describes the implementation of a system based on the design described in the previous section. Similar to the previous systems, folders were chosen as the physical hardware for the containers. Also, a shelf was chosen as the physical hardware of the storage location (Figure 8.6) since this was easier to work with for testing purposes than a filing cabinet. This section describes the implementation of the containers, and the position detection and display components of the storage locations.

8.3.1 Container implementation

The 1-wire protocol was chosen for communication between the position modules and the containers. Although, several limitations of 1-wire were previously discussed in Section 6.5.1, due to the architectural changes described in Section 8.2.2 these limitations of 1-wire do not affect the implementation of the system described here. In the previous prototypes each container network needed to support multiple containers, and each container had to support some form of a user interface. However, in this prototype, there is a maximum of one container per container network and each container only needs to support very simple functionality (i.e., reading of its ID). Therefore, issues such as the limited power and features of 1-wire devices do not have a negative impact on this system.

Each container is augmented with a DS2401 *silicon serial number* IC (Dallas Semiconductor/Maxim, 2006a). The data pin of the DS2401 is wired to a conductive pad at the back of the folder, while the ground pin is wired to a conductive pad at the front. The design of the container circuitry is simplified in comparison to that needed for the previous prototypes. In fact the DS2401 and the conductive pads are the only electronic hardware that is added to the containers.

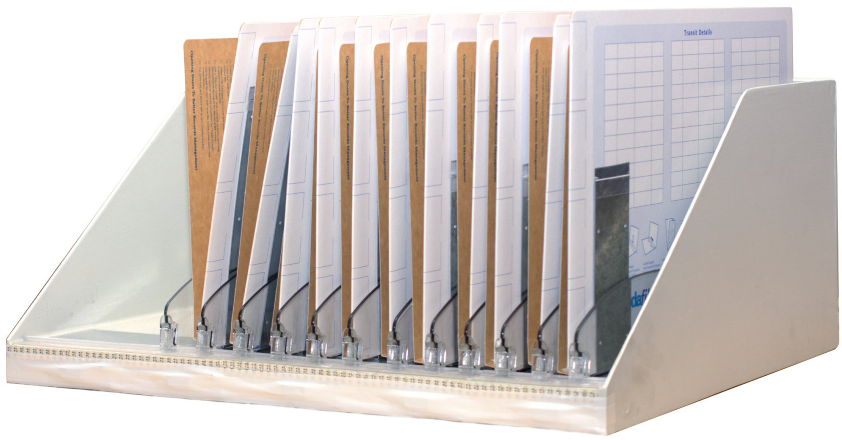

Figure 8.6: The completed prototype system. The physical hardware chosen for this system was a shelf type physical storage location, and Codafile folders.

The physical containers that are augmented in this prototype system are Codafile folders (Figure 8.7a). Several variations of conductor design were tested, and the conductor type chosen for the current version of the system is shown in Figure 8.7b. The conductors are made from a solid metal plate with a right angle bend in it, such that the conductor runs underneath and up the back of the folder. This enables the containers to sit upright without assistance, which reduces the likelihood of a container falling over and losing contact with the container network. The measurements of the conductors are given in Table 8.1. The conductors protrude from the bottom of the folders to ensure that they make contact with the storage location, since the location pads of the storage location are flush with its surface.

Table 8.1: Measurements of the implemented storage location and containers, corresponding to Figure 8.4

	Description	Length (mm)
A	Rear offset	3
B	Length of location pads	20
C	Front non-conductive surface	8
D	Width of location pads	5
E	Location pad spacing	1
F	Length of container pads	50
G	Container pad spacing	215
H	Width of container pads	18
I	Width of container	20

8.3.2 Position detection hardware

As described in Section 8.2.2, the position detection component of the architecture consists of position detection modules and a position controller. The

(a) Container hardware. **(b)** Close-up of container contacts.

Figure 8.7: Codafile folders were used as the physical hardware for the containers in this prototype.

implementation of the position modules (Figure 8.8) comprises a PIC microcontroller (see Microhip Technology, Inc., 2002), which provides the position module logic, and associated circuitry. Because each container network only has to support a single DS2401 IC, the network controller circuitry can be simplified—consisting only of a pull-up resistor (in the order of 1kΩ) and the location pad. Each container network is wired to a GPIO pin on the microcontroller, which acts as the 1-wire master for each container network. The schematic and Printed Circuit Board (PCB) layout for these modules is given in Appendix C, Section C.2.2.

Due to the design of the conductive pads of the containers and storage locations, it is likely that a container will contact several location pads, and create a short circuit between them. This effectively joins them together, forming a single container network with multiple container network controllers. These container network controllers are still able to function as normal, provided that only one of them is active at a time. It is the responsibility of the position module firmware to ensure that this is the case. The current implementation achieves this by polling a single location pad a time, as discussed in the next section.

An important component of each position module is its PCB. The location pads are part of the PCB layout, and are gold plated during construction to maximise their longevity. The PCB is very thin (0.2mm) so that it can simply be affixed to the base of the storage location, as shown in Figure 8.9.

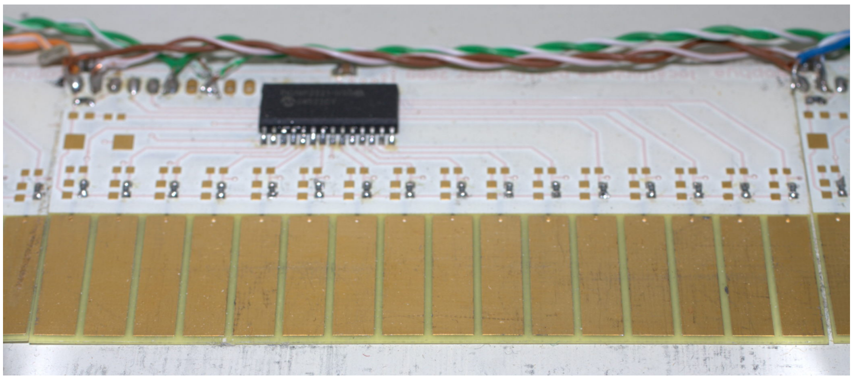

Figure 8.8: Close-up of a position module installed on a shelf physical storage location. The gold contacts at the front are the location pads, the IC is the microcontroller, and the wires at the back connect the element to power and the I^2C bus.

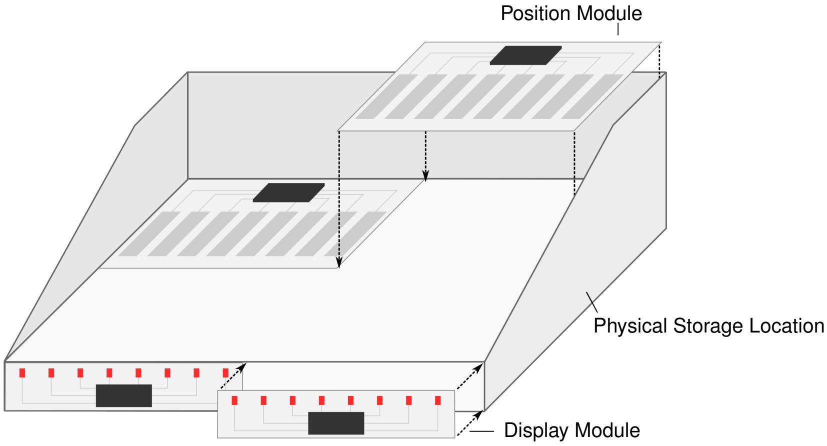

Figure 8.9: Installing position and display modules in a physical storage location. Each of the display modules has an LED that corresponds to a location pad.

8.3.3 Position detection firmware

The microcontroller of each position module executes firmware that is responsible for continuously polling its container networks and passing this information up to the position controller. The firmware of each position module microcontroller runs in a continuous loop, executing a 1-wire *read rom* command on each container network in turn to determine whether that network has a container present and, if so, reads the ID of the container. The IDs of detected containers are stored in an area of memory referred to as the *container ID list*. This area of memory is an array, each element of which corresponds to a container network. The order of the array elements corresponds to the order of the location pads.

Because each container network is read in turn, it is not possible for neighbouring location pads of the same position module to be active simultaneously. While it is possible for the edge pads of neighbouring position modules to be active simultaneously, this can be overcome by synchronising the polling of the elements, for example by having the position controller issue a reset command so that the position modules start polling at the same time.

The position controller interfaces between the position modules and higher-layer software. Position modules connect to the position controller via an I^2C bus, with the position controller acting as bus master. Currently the position modules each have an ID number that identifies their physical position hard-coded into their firmware. However, in the future, the design of the modules could be updated so that their ordering could be detecting dynamically. The firmware of the position controller loops continuously, reading the container ID list of each position module in turn and, as it does so, outputting this to the higher-layer software. The output is in the form of a key-value pair, where the key is the ID of the container network (derived from the position module number concatenated with the number of the location pad within that element) and the value is the ID of the container detected at that position (or an ID of -1 if no container was detected[1]).

8.3.4 Output display control

The architecture of the output display component of the system is similar to that of the position determination component. In fact the circuitry of the position controller is identical to that of the display controller. Similarly, the same circuitry is used for the common elements of the position and display modules (above the horizontal dashed lines shown in Figure 8.3). Similar to the position circuitry, the display circuitry is split up into *display modules* (Figure 8.10). Each display module consists of some display devices, and the logic required to control these. The output display devices used in the current implementation are single LEDs.

Each display module corresponds to a position module, and the number of LEDs on each display module is the same as the number of location pads on each position module. Therefore, when they are mounted, the position and display modules are aligned, as shown in Figure 8.9. This means that there is an LED that can be controlled for every detectable container position.

Each LED is individually controllable by the higher-layer software, and has four possible states: continuously off, continuously on, slow blink, and fast blink. The display modules connect to the display controller using an I^2C bus. The position controller receives commands to change the state of LEDs from the higher-layer software. Each state change command includes the ID of the LED whose state is to be changed, and the value to set its state register to. From the ID that is given, the display controller determines which display module it needs to address, and issues a command to it via the I^2C bus. The firmware of the display modules listens for these commands, and when they are received, it updates the state of its LEDs.

[1]Since -1 is not a valid 1-wire ID, it is not possible for a container to have an ID of -1.

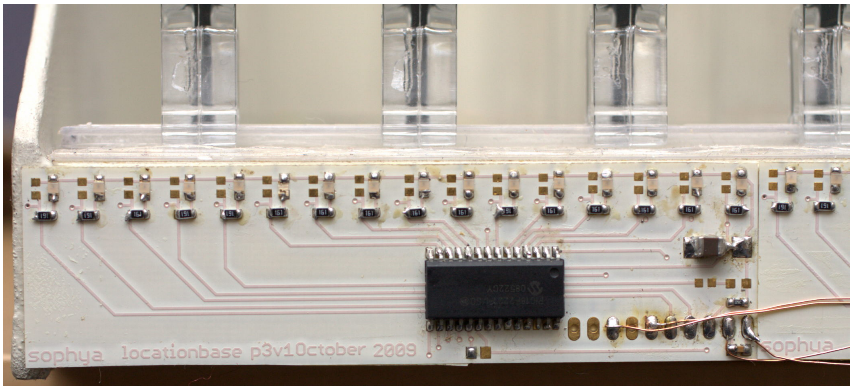

Figure 8.10: Close-up of a display module, installed at the front of a shelf physical storage location. The display modules utilise the same PCB as the location elements, but with the location pads removed.

8.3.5 Processing software

The position and display controllers connect via USB-serial bridges to a computer system that runs low-level processing software. This low-level software is responsible for processing the raw data coming from the position controller in order to maintain an internal database of container state and transmits event notifications to the middleware when there is a change in state. Conversely, the middleware can pass commands to the processing software in order to control the state of each LED of the display modules.

8.3.6 Filing assistance

Unlike the prototype systems described in the previous two chapters, which support a piling organisational strategy where the order of the containers is not important, the system discussed in this chapter supports a filing strategy in which the containers are stored in a specific order. To be able to support this strategy, the system needs to be able to guide users when they place containers in the storage locations, in order to ensure that the containers are put in the correct place.

One of the advantages of the system described in this chapter over those of the previous chapters is that it enables the user interface output components of the containers to be moved to the storage location. Using these user interface components the system is therefore able to show users where containers should be placed as the user returns them to the storage location. The SOPHYA implementation described in this chapter uses LEDs as the output display, and one or more of these LEDs can be illuminated to show the user where to place a container.

A method for the user to trigger the display for the appropriate folder is also required. The current implementation of the system provides this using RFID tags. In addition to the SOPHYA augmentations described above, folders are tagged with RFID tags. A short-range RFID tag reader is mounted near the storage location, and when the user swipes a folder past the RFID tag reader, this instructs the DMS software that the user wishes to return the folder to its appropriate place. The DMS then instructs the appropriate storage location to illuminate an LED to show the user where to place the folder. This sequence is diagrammatically shown in Figure 8.11.

8.3.7 Software

Although, as with the implementation of the previous prototype, the focus of the development of this system was on the hardware and low-level software components of the design, however, a simple software application, similar to that described in Section 7.3.5, was implemented for testing purposes. This application provided a simple GUI that listed containers and their current state, and allowed for control of their corresponding LED on the physical storage location.

For the evaluation of the system, a more advanced software application was developed that incorporated document management functionality. This is discussed as part of the evaluation described in a later chapter (see Section 10.6).

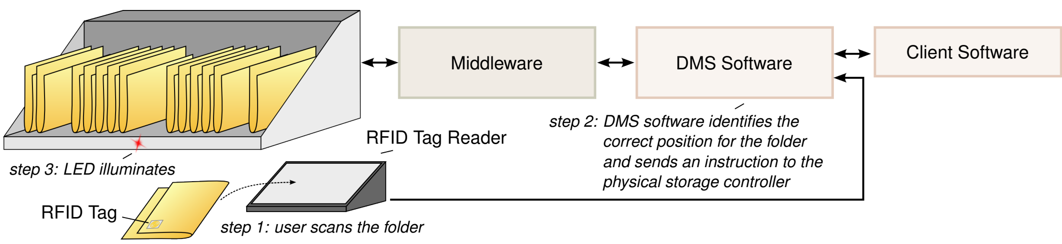

Figure 8.11: Folders are tagged with RFID tags. When the tag is scanned by an RFID reader the system illuminates an LED to show where the folder should be placed on the shelf.

8.4 Discussion

The design of the prototype described in this chapter passed through several iterations, with continuous hardware and software testing at each stage, in order to provide the most robust and reliable system possible within the prototyping constraints. These tests involved evaluating different materials and hardware designs to find the most reliable.

The reliability of the detection of containers in the storage locations is an important factor in ensuring that the system is usable. For the system to work reliably, it is important that there is good contact made between the conductive pads of the containers and those of the storage location. The containers must be seated properly in the storage location in order for them to be detected, and if they fall over or sit too far forward it is not possible for the system to detect them. This problem can be partly mitigated by having the system provide feedback to the user when they place a container into the storage location, for example by blinking an LED when the container is detected. Thus, if the user ensures the LED blinks when they place the container, then they can be sure that the container has been detected by the system.

As with the previous prototype, the responsiveness of the system is closely linked to its reliability (see Section 7.4.3). With the current implementation of the system the latency between adding a container to a storage location and the detection of this by the system is insignificant. Similarly, the latency when controlling the output display is also insignificant. The only perceptible latency is when removing containers, which is several seconds. This is similar to the latency issue discussed in Section 7.4.3.

A related problem that was encountered during development was that the folders would not sit upright when on their own. Instead, they needed to have folders around them to keep them up. Consequently, if they fell over their conductors would lose contact with the location pads and they would be detected as being 'removed' from the storage location. Two steps were taken to solve this problem. First the conductors were designed, as discussed in Section 8.3.1, to reduce the likelihood of the folder falling over. Second, adjustable plastic dividers were added to the storage location to hold up the folders, as shown in Figure 8.6. Another option would b e to incorporate an RFID tag reader into the storage location. While this would not be able to provide accurate position detection, it could indicate when containers are in the storage location but are not detected on a container network (e.g, due to misplacement).

The current design of the conductors works sufficiently for the prototype system, and it is envisaged that the reliability problems could be overcome with better industrial design of both the container and storage location conductors. Such industrial design, however, is beyond the scope of this thesis.

8.5 Summary

This chapter described the design and implementation of a prototype system that followed a similar concept to the previous prototype, but also added support for determining the position, and therefore order, of containers within storage locations. The ability to determine ordering of containers within a collection can be useful in a number of cases (as described in Section 8.1), for example in libraries where the collections are maintained in a strict order. In order to add this support, the design of the physical interface component of the architecture described in the previous chapter was modified (Section 8.2), and a prototype system was implemented based on the new design (Section 8.3). The reliability of the system was discussed in Section 8.4. The next part of this thesis discusses the evaluation of this prototype and the prototype discussed in the previous chapter.

Part IV

Evaluation and Conclusions

Chapter 9

Evaluating the Functional Requirements of SOPHYA

The previous three chapters discussed the development of three prototype systems. The first of these systems was a proof-of-concept prototype, described in Chapter 6. This system, however, had several limitations which necessitated the development of a second prototype system, which has been described in Chapter 7. While this system improved on the design of the first prototype, it was not able to determine the ordering of document containers in their physical storage locations. In order to provide this functionality, a third prototype was developed, as described in Chapter 8. This system followed a similar high-level architecture to the second prototype, but added support for determining the ordering of containers.

This chapter begins with a comparison of the second and third prototype systems (Section 9.1), as these are the two main systems developed during this research. The comparison outlines the differences between the two systems, and shows how they complement each other. This comparison is followed by three examples of potential applications of these systems in Section 9.2. A discussion of how the high-level architecture that these systems fit into enables flexibility in selecting which (or both) of the systems to use for a given application follows (Section 9.3). Section 9.4 provides a review of how well SOPHYA implements the functional requirements of integrated paper and electronic document management systems that were identified in Section 4.4.1. Finally, the chapter concludes with a summary (Section 9.5).

9.1 Comparison of SOPHYA Systems

This section provides a comparison between the prototype systems introduced in Chapters 7 and 8. Because these two systems are both variants of the same high-level architecture, they are both regarded as variations of the SOPHYA system. The first SOPHYA prototype is referred to as the unordered SOPHYA system, since it is suited for managing collections that are not maintained in a specific order, while the second SOPHYA prototype is referred to as the ordered SOPHYA system, since it supports management of ordered collections.

Although the ordered SOPHYA system was developed subsequent to the unordered SOPHYA system, it does not supersede the unordered system, since each of these systems has its own advantages, and they are able to serve complementary purposes. An overview of the differences between the two prototypes is given in Table 9.1.

Table 9.1: Comparison of unordered and ordered SOPHYA prototypes.

Technical differences	**Unordered**	**Ordered**
Detection of container order supported	No	Yes
Contact between containers and storage location	Spine and sides	Spine only
Container networks per storage location	Few	Many
Containers per container network	Multiple	One
Stacking of containers supported	Yes	No
User interface devices on container	Yes	No∗
User interface devices on storage location	Partial†	Yes
Application differences		
Organisational strategy	Piling	Filing
Level of storage	Hot, Warm	Cold

∗ Not currently supported, but technically possible, see Section 11.3.2 for discussion.

† Storage location-based user interface components of the unordered system cannot show the position of containers, and thus can only be used to support container-based user interface components.

To be able to support detection of order, the storage locations of the ordered prototype need to have many container networks (at least as many as there is space for containers), and it must not be possible for multiple containers to contact the same container network. This is in contrast to the unordered system, in which the number of container networks per storage location is small (as few as one), and there can be many containers on a container network. A result of this is that the unordered system enables greater flexibility of container placement, for example supporting stacking of containers (as discussed in Section 7.3.2) which the ordered system cannot.

Another difference between the systems is the placement of user interfaces devices—the unordered system is designed for container-mounted user interface devices, while the ordered system is designed for storage location-mounted user interface devices. While it is technically possible for the ordered system to support container-mounted user interface devices, this has not been implemented in the current prototype (for a more detailed discussion of this see Section 11.3.2). Conversely, since the unordered system is unable to determine the position of containers within the storage location, it is only possible for it to display the position of a given container to the user with container-mounted output devices. Thus, although storage location-mounted user interface devices are possible, these are only able to act in a supporting role in determining the locations of containers.

The technical differences between the systems result in practical differences that make them suited to different roles. The unordered system is better suited to supporting a piling strategy (see Section 2.3.1), where the order of the elements being organised is not important. The ordered system, on the other hand, best supports a filing strategy (see Section 2.3.1), where the elements are organised in some order. Additionally, the unordered system is more suited for *hot* or *warm* containers (as discussed in Section 2.3.1), that are likely to be frequently moving around, through numerous storage locations such as shelves, document trays, and desktops. The unordered system is better able to support this type of use since it allows for more flexible placement due to support for stacking by allowing contact between container sides as well as their spines. Alternatively, the ordered system is better able to support *cold* containers which are stored in archives, and where some ordering is often useful. Therefore, in most cases where both piling and filing strategies are followed, and/or a combination of hot, warm, and cold containers are used, the two systems need to be used in combination.

The next section provides three examples of how the two SOPHYA systems can be applied in different environments.

9.2 Example Applications

This section outlines three examples of environments that could benefit from SOPHYA. The first example is based on the observation of an office that was visited as part of the observational study described in Chapter 4. The second is a typical library environment, and the third example demonstrates how the two different SOPHYA systems can be used together in an office to support the management of different levels of information.

9.2.1 Production department

The first example of an environment that could benefit from integration of paper and electronic document management using a system such as SOPHYA is based on observation of a real-world office. The following scenario is loosely based on the work-flow of Office 2, a newspaper production department where advertisements are designed, as described in Chapter 4.

> *The work-flow begins with sales representatives in the field getting jobs from clients. A 'job sheet' is then filled in for each job, and is brought back to the office along with any related physical material. All material relating to a job is placed in a 'job-bag'. Information from the job sheet is then entered into the job management software and the job is assigned a unique ID. The job-bag goes to whoever is working on the job, and may get passed around if more than one person needs to work on it, though only one person can have it at a time. Job-bags for jobs that are incomplete, but not currently being worked on are stored on shelves at the centre of the office. These shelves provide a quick visual indicator of how much work remains to be done.*

Such a work-flow could be supported using the technology of the unordered SOPHYA prototype, as shown in Figure 9.1a. The physical in-trays and desktops of the designers working in this office are augmented with physical storage location circuitry, and the job-bags are augmented with container circuitry. SOPHYA enables the physical containers (i.e., the job-bags) to be integrated with the existing job management system used in the office. Since, each job (and thus job-bag) is already assigned a unique ID, the link between the physical job-bag and the the job's entry in the job management system can be created by mapping between the job ID and the container ID of the augmented job-bag. This could, for instance, be done when entering the electronic information about a job into the job management system. While doing this the sales representative could place the job-bag on a container reader (which could, for example, be a document tray designated for this purpose) to create a mapping between its job ID and container ID.

The development of a more advanced job management system, and associated client software, would enable the office to take advantage of the extra functionality made available by SOPHYA. For example, it would be possible to track job-bags as they are processed and moved between different people, making it possible to dynamically view the location and history of job-bags. The job-bags could be augmented with user interface devices, such as LEDs, so that, for example, as deadlines approach, the LEDs on the job-bags could signal which jobs have the highest priority.

9.2.2 Library

The second example of an environment that could benefit from using SOPHYA is a library. Libraries are a good example of an environment where the management of physical artefacts (e.g., books and journals) can benefit from integration with electronic systems, such as electronic cataloguing and lending systems. In a typical library, each physical artefact has an entry in the catalogue, and the physical and electronic manifestations are linked by a unique ID that they share. This unique ID is commonly encoded as a barcode, which is attached to the artefact and is used to identify it (e.g., when checking books in and out). However, barcodes are generally not used for tracking the location of artefacts in the library. Instead, when artefacts are retrieved from the shelves, they must be located manually using serial numbers that are visible when the artefact is shelved. The lack of integration with an electronic tracking system means that if an artefact is, for example, placed on the wrong shelf it has to be found by manual searching. This is especially an issue for large libraries. Furthermore, previous research has indicated that library patrons often have difficulty physically locating books (even those that are shelved correctly) within the library (see McKay and Conyers, 2010).

These problems could be reduced by incorporating a system such as SOPHYA for tracking the position of artefacts on the library shelves, and providing vi-

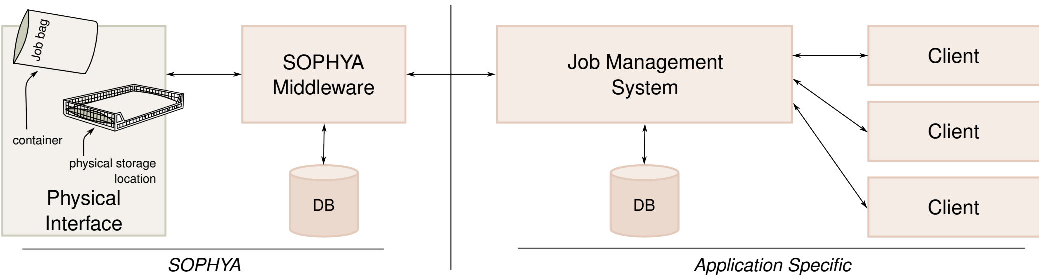

(a) Production department application (described in Section 9.2.1)

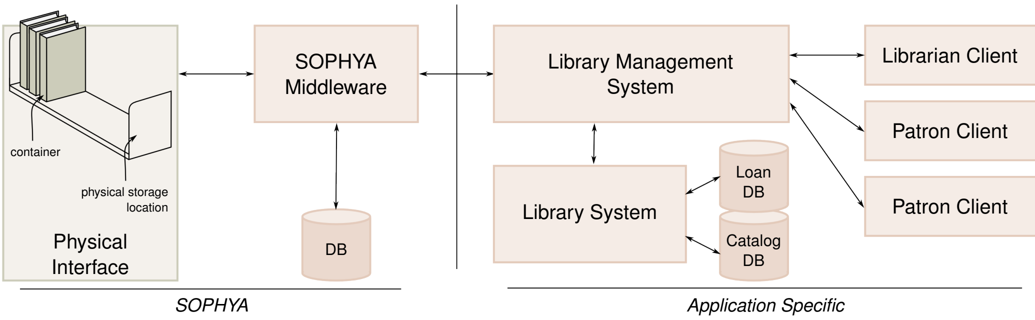

(b) Library application (described in Section 9.2.2)

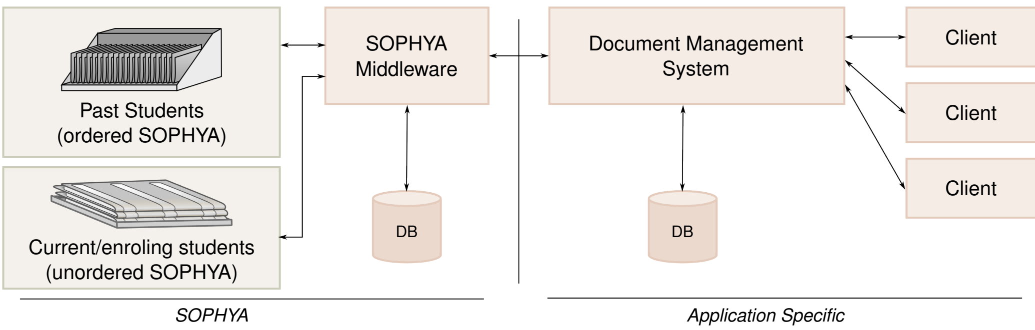

(c) University faculty office application (described in Section 9.2.3)

Figure 9.1: Examples of different applications of SOPHYA.

sual output that is embodied in the library shelves. Because libraries typify an environment where collections are ordered, they are suited to using the ordered version of SOPHYA. Figure 9.1b gives an example architecture for integrating physical artefacts and the electronic catalogue and lending systems. The physical interface element of the architecture consists of the *containers* (e.g., books) and *storage locations* (e.g., shelves) which are augmented with SOPHYA circuitry. The *SOPHYA middleware* links these to the *library management system*.

The *library system* element of the architecture represents the existing cataloguing and lending systems used by the library. The *loan database* keeps track of items that are on loan, while the *catalogue database* provides catalogue information about the items (e.g., title, author). The library management system is responsible for combining location information from SOPHYA, with catalogue and lending information from the library system, and making this available to the client software.

Clients access integrated library collection information through the library management software. There may be different clients for different purposes, allowing different levels of access. For example, the librarian client would allow librarians to access loan information for all library patrons and to add entries to the catalogue database, while the patron client would only be able to view information about books and loan information for the logged-in user.

The integration of the physical and electronic elements creates possibilities for clients that would not otherwise be available, such as on-shelf visualisations using the output display devices on the shelves, and remote visual browsing of the shelves electronically based on the real-world position of the artefacts.

9.2.3 University faculty

The previous two examples have shown how the two SOPHYA variants could be used individually to support specific document management needs in different situations. However, there are many cases where the systems can both be used within the same office, serving complementary purposes, such as managing different levels of information. One example of an office where this could potentially be useful is the office of a university faculty. In such an office, SOPHYA may be used for the management of student records.

It was noted in Section 2.3.1 that people commonly work with multiple levels of information (e.g., archives of cold documents, warm documents kept close to hand, and hot documents that are in active use). In the case of a faculty office, the records of past students may be considered to be cold, those of present students could be considered to be warm, and records of students who are currently going through the enrolment process could be considered hot. In supporting the management of these records, it may be desirable to store the cold past student records in a filing cabinet with strict ordering, and the warm present student records in order on a shelf. The hot records for students in the process of enrolling, however, would likely be better suited to the less rigid structure afforded by a piling strategy.

These needs could be met by using both SOPHYA variants in combination, as shown in Figure 9.1c. The ordered SOPHYA system could provide the necessary support for maintaining the filing organisation of the cold and warm records, while the unordered SOPHYA system would better support the hot records (e.g., allowing them to be piled in document trays or on an augmented desktop). The ability to support such a combination of SOPHYA systems is enabled by the high-level architecture, as described in the next section.

9.3 High-level Architecture

These above examples showed how the different two SOPHYA implementations could be used for different applications. It is also possible that both systems could be used together within the same application, but serving different purposes. It is therefore important that both systems can be supported within the same architecture.

The architecture on which the SOPHYA prototypes are based was first introduced in Section 7.2. This architecture is divided into four components (shown in Figure 9.2): physical interface, middleware, DMS software, and client software. As demonstrated by the examples in the previous section, the DMS and client software are both application specific, and would vary depending on where the SOPHYA systems are deployed. The SOPHYA middleware component provides a level of abstraction between the DMS software and the physical interfaces. It enables arbitrary combinations of SOPHYA physical interfaces to be used depending on the needs of a particular application, and provides the DMS with a single hardware-independent view.

9.3.1 Middleware interface

The interface that the middleware exposes to the DMS software layer enables such software to do the following:

- Get a list of storage locations.
- Get a list of containers (both those currently present in a storage location and those that have been previously present but have since been removed).
- For each storage location:
  - get a list of which containers are present (in the order that they are placed in the storage location if supported by the SOPHYA hardware);
  - control the user interface output devices on the storage locations (if included); and
  - read from the user interface input devices on the storage location (if included).
- For each container:
  - determine which storage location it is currently in (if included), or was previously in (if no longer present);
  - determine the time of its most recent state change (i.e., addition or removal);

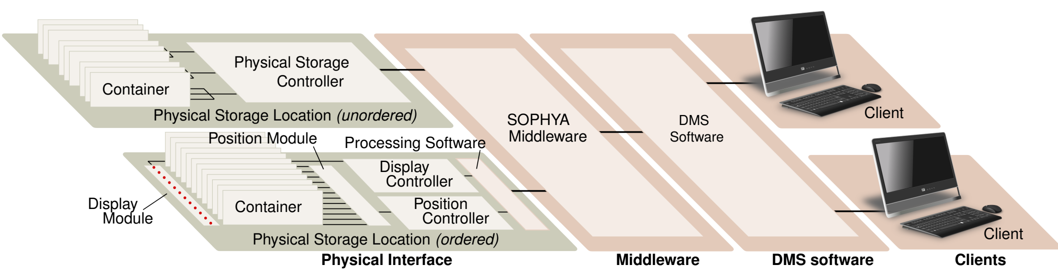

Figure 9.2: High-level view of the architecture of the systems described in Chapters 7 and 8, showing how the two systems are able to be used together.

  - control its user interface output devices; and
  - read from its user interface input devices (if included).
- Receive events when a container is added to, or removed from, a physical storage location.
- Receive events when a container is scanned by a connected RFID tag reader.

The current middleware supports most of these features. It is currently designed with support specifically for the container and storage location user interfaces that have been implemented. As further user interface devices are added to containers and storage locations, the middleware would need to be updated to be aware of these. The design and implementation concerns surrounding more advanced user interfaces are discussed in Section 11.3.2.

9.3.2 DMS software implementation

Using the middleware interface described above, DMS software can be developed that links electronic metadata to the physical containers. The software can make use of the user interface devices on the containers and storage locations to perform functions such as showing users where to retrieve containers from, where to place containers, and to display alerts.

The event notifications and corresponding information available through the middleware interface can be used to add event driven features to the DMS software. Examples of such features include logging of container history, identifying frequently accessed containers, as well as more application specific uses such as detecting when a container progresses to the next stage in the office work-flow. Similarly, for instance, the event generated when a container is scanned by the tag reader can be used to trigger display on the appropriate storage location to show where it should be placed.

The next section revisits the requirements introduced in Section 4.4.1, and discusses to what degree the features that SOPHYA provides to the DMS software can be used to implement a system that would meet these requirements.

9.4 Fulfilment of Requirements

This section provides a discussion of how and to what extent the SOPHYA systems described in the previous chapters fulfil the functional requirements specified in Section 4.4.1. It is worth noting that SOPHYA itself is not a stand-alone system, but rather a component that would be combined with other components (e.g., DMS software) in order to form a complete system. Therefore, the sections below discuss the specified requirements in terms of the functionality that SOPHYA supports, and other components that have to exist in order to fulfil these requirements.

9.4.1 Organisation

Support filing a document into multiple categories

In order to remove the need to file multiple copies of a single document, it should be possible to file a single paper document into multiple categories using electronic metadata. Although SOPHYA does not work at a document level, it provides part of the solution to this problem. SOPHYA allows linking a container to metadata, and the container metadata can include links to metadata for documents within the specific containers. This metadata may include the multiple categories to which the document may belong. Therefore, it is possible to associate a single document with an arbitrary number of categories. This is similar to how some manual systems work (e.g., Wilson, 2001) but SOPHYA provides improved linking between the physical and electronic elements of the document management system through its embodied interface.

Integrate fragmented paper document collections

In the study described in Chapter 4, it was observed that fragmentation can occur not only between paper and electronic document management, but also when different paper document management systems are combined (for example when companies merge). SOPHYA can assist in integrating these fragmented paper document collections in the same way that it bridges the fragmentation between paper and electronic document management. Therefore, this requirement is met implicitly as part of the design of SOPHYA. Because the fragmented paper document management systems would be managed by a common metadata layer, the fact that their previous organisational systems are fragmented is rendered irrelevant.

Link paper documents to electronic documents

Another requirement was that the system should have support for linking paper documents to electronic documents. This is another document tracking problem, and therefore SOPHYA would not assist with this directly. To meet this requirement some form of individual document tagging (e.g., with RFID tags) would be required. This option will be discussed in Section 11.3.3.

9.4.2 Maintenance

Automate maintenance of collection indexes

The observational study identified that three of the visited offices used manually maintained indexes to assist in locating folders within a collection. In order to improve this, one of the requirements was to automate the maintenance of such indexes using electronic metadata, both manually entered and automatically generated.

As described in Section 7.2.1, SOPHYA operates at the document container (e.g., folder) level. The indexes observed in the study were also maintained at container level. SOPHYA is therefore ideally suited to enabling such indexes to be maintained electronically and is able to provide tighter integration between the collection index and the physical containers that it indexes. The extent to which this is supported depends on the software implementation.

Simplify filing and refiling of documents

It was noted during the observational study that filing (and refiling) documents was a time consuming task. Therefore, a requirement was to simplify this process. Aside from removing the need to have multiple copies of a document, as described above, SOPHYA can aid in simplifying the task of filing and refiling. SOPHYA enables DMS software to be implemented that can aid the user in selecting the right container to file a document in. In the case of filing, the selection of the container could be done manually, or could be automated with suggestions based on document metadata. In the case of refiling, it is likely that the document would simply be returned to the same container that it was previously stored in. Once the container that the document needs to be placed in is identified by the DMS software it can use the output display of the SOPHYA component to assist the user in finding the correct container, and returning the container back to its current location once the document is filed in it. For example, when the user needs to refile a document they initiate this process (e.g., if the document is RFID tagged, its tag could be scanned by a tag reader) and the appropriate container would activate its output display, quickly showing the user where to place the document.

Better support for purging irrelevant information

In order to reduce the amount of irrelevant information stored in collections, one of the requirements was to better support purging of unnecessary and irrelevant documents. This can be achieved, for example, by tracking the age and frequency of access to specific documents and containers, along with other properties such as their type, and combine this information with organisational rules and policies to determine when documents, or containers of documents, should be purged. SOPHYA can support this in two ways: firstly it can track access to containers (and documents if combined with document level tracking, as shown in Section 11.3.3), and secondly it can be used to display alerts. The alerts feature could be used, for example, to indicate to users when a container holds irrelevant documents that need to be purged. These alerts could be generated either automatically (e.g., using set dates), or manually when the user decides to purge specific documents (e.g., those that match a certain set of criteria).

Reduce manual labour

An example of a time consuming manual task is starting a new folder. This can be time consuming due to manual tasks such as labelling the folder, and data entry. By automating parts of the data entry, and reducing the amount of labelling required it could be possible to reduce this manual labour and thus the time taken to create a new document folder.

9.4.3 Control

Track documents and document containers

In order for the system to be able to assist the user in locating containers and documents, as well as preventing their loss, it is necessary to have some method of tracking these containers and documents. SOPHYA is able to detect the location of containers (with different levels of accuracy, dependent on the SOPHYA implementation used). The method of connectivity employed by

SOPHYA has the limitation that it is only able to track containers when they are placed in a SOPHYA augmented storage location. However, there are ways of reducing the effects of this limitation, by for instance having an extended range of SOPHYA augmented storage locations, such as desktops and document trays (as shown in Section 7.3.3), so that containers are always stored in a SOPHYA augmented storage location. If greater control is required, it is also possible to add some form of wireless tracking, such as the use RFID tags, which will be described in Section 11.3.4.

To apply the same strategy to finding individual documents would depend on the implementation of the document management system. It is possible, for example, for information about which container each document is stored in to be manually maintained in the database of the document management system if they are not commonly removed from containers, or have set containers to which they belong. If a more dynamic approach is required, SOPHYA could be combined with some form of RFID document tracking, which is described in more detail in Section 11.3.4.

9.4.4 Storage

Reduce storage space requirements

Reducing storage space requirements, especially with regard to paper documents was another requirement. This can be achieved as a result of meeting several of the other requirements described here. For example, by improving purging and reducing unnecessary duplication, the storage space requirements of the paper document management system can be reduced.

9.4.5 Retrieval

Integrate locating of paper and electronic documents and document containers

Another requirement that was identified was to provide better integration when trying to locate paper and electronic documents. The approach that has been taken, as described in Chapter 5, has been to develop a system that enables the use of a common metadata layer that supports the management of documents and information in both paper and electronic form. This common metadata layer enables searching and browsing of metadata for both paper and electronic documents. Once the appropriate documents are identified, the SOPHYA system that links between the metadata and physical document containers can assist with their retrieval (as discussed below).

Reduce time locating documents and document containers

On multiple occasions participants in the observational study remarked on the time consumed trying to locate documents and containers. Therefore another of the requirements identified was that systems integrating paper and electronic document management should reduce the time taken to locate documents and containers.

Through its ability to use the output display components of the containers and storage locations, combined with its ability to link these with electronic metadata (as described above) SOPHYA is able to simplify the process of retrieving containers. For example, the process may have previously involved opening the indexes spreadsheet, finding the appropriate entry in the spreadsheet, finding the key that identifies the container (e.g., its unique ID number or name), copying or remembering this key, going to the archive, and finally locating and retrieving the container with the given key. This could be replaced with a simpler process of opening the document management software, browsing or searching using appropriate metadata to locate the entry for the required container, selecting the entry and instructing the system that it is to be retrieved, and finally going to the archive and retrieving the container whose output display is active.

The fact that SOPHYA operates at a container level, and the implications that this has for the management of documents is discussed above. These implications apply similarly here.

Reduce difficulty of finding documents and document containers in an unfamiliar collection

Another requirement that was specified previously was to make it easier for users who are unfamiliar with a collection to find documents or containers within it. SOPHYA supports this requirement in that it makes it easier to find containers (and to some extent documents) using metadata. By allowing users to browse and search document collections using document metadata it allows them to bypass the user specific organisational hierarchies. However, while SOPHYA supports integration of paper document management with such a system, a DMS that supports these features needs to be used.

9.4.6 Miscellaneous

Ability to retrofit existing collections

Implementing the system such that it is efficient for retrofitting large existing document collections is largely an industrial design consideration. While the task of retrofitting large document collections would always be a labour intensive operation, the labour would be significantly less when retrofitting with SOPHYA technology than tagging individual documents. However, the needs of the organisation would have to be balanced by the cost of this type of operation. For example, if container level management was sufficient and electronic information was already available for the collection then retrofitting would potentially be a matter of transferring the existing documents into SOPHYA augmented containers, and linking these containers to their electronic metadata. Whereas if finer granularity tracking of documents was required, the SOPHYA retrofit could be combined with tagging of individual documents. Designing SOPHYA for retrofitting is described in more detail in Section 11.3.5.

Ability to generate and display alerts

A feature that SOPHYA is able to support is the ability to display alerts using the output display on the containers and storage locations. This could be used, for example, to visually indicate the priority of different jobs in situations where jobs correspond to a container. The electronic DMS can activate these output displays to indicate, for example which containers need attention. With the currently implemented output displays, different priorities could be displayed using various LED colours, or higher priorities could be indicated, for instance, by flashing the LED output display. Future SOPHYA systems may support more advanced displays, and therefore more detailed display of alerts could be supported. For example, if LCD screens were embedded in the storage locations, they could provide a description of specific alerts.

The triggering of alerts would be implemented in the document management software. For instance, they could be triggered based on inputs from the system (e.g., the placement of a container in a given storage location), from other systems (e.g., a document being scanned by an RFID tag reader), or from software events (e.g., receipt of an email).

9.5 Summary

This chapter provided a comparison between the two SOPHYA prototypes that were described in the previous chapters. This comparison showed that the two systems are able to serve complementary purposes, with the unordered SOPHYA system being suited to supporting a piling organisational strategy (as is commonly used for hot or warm documents), while the ordered system is able to support a filing organisational strategy, where some form of ordering is important (such as for archives of cold documents). Three example applications showed how the two different SOPHYA systems could be applied in different environments.

The high-level architecture of SOPHYA that was previously introduced in Chapter 7 was revisited, and was discussed in the context of both SOPHYA systems. The design of the architecture abstracts the details of the physical interface away from the application software, enabling the most suitable combination of SOPHYA hardware to be chosen for a particular application. This architecture allows, for example, both the unordered and ordered SOPHYA systems to be used together simultaneously.

SOPHYA is able to be used as part of a system that fulfils the requirements identified in Section 4.4.1. As previously discussed, SOPHYA needs to be combined with other application specific software components. Therefore, the fulfilment of these requirements is also dependent on these software components. In cases where finer granularity of tracking than container level is required (i.e., at the individual document level) or when tracking of containers outside of storage locations is required, SOPHYA technology would need to be combined with some form of wireless technology (e.g., RFID), which is discussed in Section 11.3.3 and Section 11.3.4.

Although this chapter has demonstrated how SOPHYA, when combined with other software components, is able to fulfil the specified functional requirements, a user evaluation of its usability was also needed. The next chapter describes a user study of SOPHYA that has been carried out to demonstrate the extent to which it supports integration of document management in an office-type setting.

Chapter 10

Evaluation Study

This chapter describes the design, execution, and results of a user study that was conducted in order to evaluate the two SOPHYA systems that have previously been described in Chapters 7 and 8. Both of these systems were evaluated together, since they are able to serve complementary purposes, as described in Chapter 9.

This chapter begins with a description of the study methodology (Section 10.1), followed by sections describing the three parts of the study: the tutorial (Section 10.2), tasks (Section 10.3), and interview (Section 10.4). These are followed by an overview of the study participants (Section 10.5). Section 10.6 describes the application software that was developed for use in this study. The results are then presented in two parts, quantitative (Section 10.7) and qualitative results (Section 10.8). The chapter concludes with a summary (Section 10.9).

10.1 Study Methodology

The first step in designing this study was to chose the type of study to conduct. The type of study needed to balance between, on the one hand, providing users with a realistic environment in which to evaluate the system, and on the other working within several constraints. These constraints include the reliability of the system and the time frame of the research. While the current implementations of the prototype systems provide a level of reliability that is sufficient for a short-term laboratory-based evaluation, they would require further development before being ready to be deployed in a real office. Similarly, while the systems are scalable by design, the current implementations have not had enough hardware constructed to enable their use, for instance, in large archives. Therefore, although perhaps the most beneficial method of evaluating the system would be to deploy it in a real-world office environment and conduct a longitudinal study of its use, this would be impractical given the constraints described above.

Taking these considerations into account, a laboratory-based study was chosen as being sufficient for evaluating the potential application of SOPHYA in an integrated paper and electronic document management environment. However, to add further value to the findings of the study, participants were chosen carefully for their experience of working in offices that manage both paper and electronic documents. The study aimed to get feedback from the participants based on their valuable previous experience.

A number of people from a range of offices participated in the study. The study was conducted individually with each participant, with the same process being repeated for all participants. Each session of the study comprised three parts:

Tutorial. The first part of the session was a tutorial, the purpose of which was to familiarise the participant with the hardware and software that they would be using during the study.

Tasks. Once the participant had completed the tutorial, they were asked to perform five separate document management tasks using SOPHYA. At the completion of each task the participants filled in a short questionnaire.

Interview. After all the tasks had been completed, an unstructured interview was conducted with the participant.

Participants were given a handbook which provided all the instructions required to complete the tutorial and tasks (the handbook is included in Appendix E).

The study was conducted with the approval of the ethics committee of the Faculty of Computing and Mathematical Sciences, University of Waikato. A copy of the approval letter can be found in Appendix A.

10.1.1 Experimental setup

The study was conducted in a usability laboratory at the Department of Computer Science, University of Waikato. All the components of the system used in the study were setup on a desk in the laboratory, as shown in Figure 10.1. The components that were used included:

1. a shelf and 7 folders augmented with the unordered SOPHYA technology;
2. a shelf and 12 folders augmented with the ordered SOPHYA technology (3 additional, unaugmented folders were placed on the shelf as fillers);
3. four unaugmented, conventional document trays (used for organising documents as part of the study tasks);
4. an RFID tag reader; and
5. a laptop computer running the application specific software that was developed specifically for this evaluation (described in Section 10.6).

10.1.2 Task scenario

In order to provide as realistic a simulation of a real-world office as was possible, a scenario of a business office was developed and used as the context for the tutorial and tasks that the study participants performed. The scenario was designed to simulate the type of office in which technology such as SOPHYA would potentially be useful.

The scenario simulated the office of a fictitious landscape design firm, in which

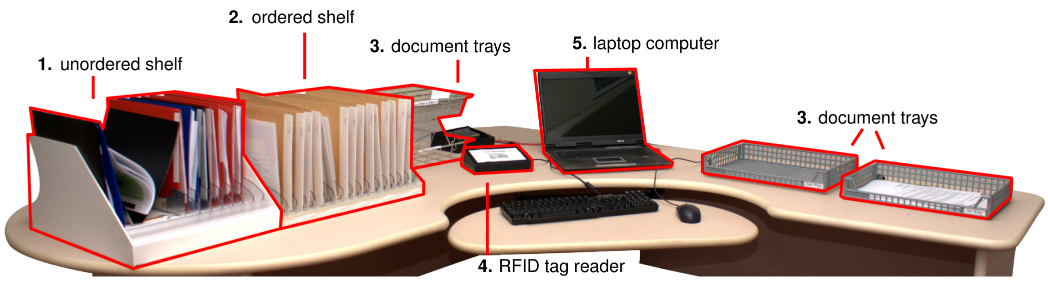

Figure 10.1: Experimental setup for the evaluation study.

the participants enacted the role of a personal assistant to the manager of the firm. The organisation of documents maintained by this fictional firm was job-centric, and as such, documents were arranged into folders where each folder corresponded to a fictitious job. Each folder was augmented with SOPHYA technology, which enabled the physical folder to be integrated with its electronic representation in document management system (described below).

The office that was simulated in this scenario used a combination of the unordered and ordered versions of SOPHYA in a manner similar to that described in Section 9.2.3. The unordered SOPHYA system was used for job folders that were 'in progress' (i.e., jobs that were, according to the scenario, being actively worked on). These were the hot or warm job folders. On the other hand the ordered SOPHYA system was used for the management of folders for 'completed' jobs that had been archived. The archived folders were stored on a shelf in alphabetical order.

10.2 Tutorial

Each study session began with a tutorial that the participant had to complete. This served to familiarise the participants with the systems that they would be using during the study, and to get them accustomed to the environment. The full text of the tutorial is in the evaluation handbook, included in Appendix E.

The tutorial stepped the participants through performing a task that the evaluation handbook described as follows.

> *You have received in your 'in tray' hard copies of the 'Pool and Driveway Maintenance Guide', which need to be filed in the folders of jobs that were started before 8 January 2010 and include a pool and driveway.*

To perform this task the participants had to perform several sub tasks, including:

- searching for jobs that match the given criteria;
- locating and retrieving the job folders for those jobs using the folder LEDs;
- placing the given documents into those folders; and

- returning the folders to the appropriate places on the shelves.

Once the participants had completed the tutorial they began the second part of the study, which required them to carry-out a set of tasks, as described in the next section.

10.3 Tasks

The participants were given five tasks to complete. These tasks required the participants to use SOPHYA to perform a range of filing activities which are normally carried out in a typical office environment.

The instructions for each task were included in the evaluation handbook (Appendix E). These began with an overview of the task, followed by an enumerated series of subtasks that the participant had to perform in order to complete the task. The tasks are described in more detail below.

10.3.1 Task 1 — Retrieving folders using the unordered system

The aim of the first task was to demonstrate the link between electronic search and physical retrieval. The task required the participants to retrieve folders for jobs that matched a given date criterion. The jobs that matched this criterion were the 'in-progress' jobs, and thus the folders used the unordered version of SOPHYA. A summary of the instructions for this task is given below.

> *The manager of Blandscapes wishes to review some of the jobs that have been in progress since 9 April 2010, and has asked you to retrieve the appropriate folders and place them in a tray to be collected later.*
>
> 1. *Use the search facility of the software to locate the jobs which are still in progress, and which were started after 9 April 2010.*
> 2. *Retrieve the appropriate files and place them in the "Task 1" document tray.*

10.3.2 Task 2 — Retrieving folders using the ordered system

The second task was almost the same as the first task, except that the participants had to retrieve folders for ‘completed’ jobs. Therefore, the participants used the unordered SOPHYA system for Task 1, and the ordered SOPHYA system for Task 2. The instructions for this task are given below.

> *You have been asked by one of the designers to retrieve the folders for all jobs completed before 20 June, 2009.*
>
> 1. *Use the search facility of the software to locate the jobs which were completed before 20/06/2009.*
> 2. *Retrieve the appropriate files and place them in the “Task 2” document tray.*

10.3.3 Task 3 — Collating specific documents from different folders

The third task was more complex than the first two tasks, and demonstrated how SOPHYA could be used as part of a larger task. Although the task was clearly broken down into steps for the participants, it required some cognitive effort on the part of the participants since the search for the folders had to be done in two steps. The task required the participants to retrieve three job folders based on two different sets of criteria. From each of these folders the participants had to retrieve a specific document, and gather these documents together into a new folder. The participants also had to return the folders that they had retrieved back to their original places on the shelf.

Since the folders that the participants retrieved for this task were all for completed jobs, the ordered SOPHYA system was used. Therefore, when folders were returned to the shelf, they needed to be placed in the correct position. SOPHYA was used to provide guidance for placement of the folders. Each of the folders had an RFID tag attached, which when waved past an RFID tag reader caused the system to turn on the LEDs at the position where the folder belonged on the shelf (as discussed in Section 8.3.6).

The instructions that the participants were given for this task were as follows.

> *One of the company's sales reps is going to meet with prospective clients. You have been asked to prepare a file of some completed drawings to be shown to clients at the meeting. The sales rep has asked for drawings from all jobs that match* either *of the following criteria:*
>
> - completed *jobs with a pool* and *water feature*
> - completed *jobs matching the keyword "Newton"*
>
> *Repeat the following steps for each of the criteria bullet pointed above:*
>
> 1. *Bring up the* Search Jobs *window.*
> 2. *Use the search function to find the job(s) matching the criteria.*
> 3. *Use the software to determine the folders' real world positions and retrieve them from the shelf.*
> 4. *Retrieve the design from each folder and place it into the meeting folder (you will find this in your "In Tray").*
> 5. *Return the folders to the correct location on the shelf (refer back to the tutorial section* Returning folders to their shelves *to remind yourself of the correct procedure for returning folders).*
> 6. *Close the* Search Jobs *window.*

10.3.4 Task 4 — Combining paper and electronic documents

The fourth task was designed to demonstrate that SOPHYA could be used to integrate the management and use of both electronic and paper documents. In performing this task participants used a combination of paper and electronic documents. Participants were required to retrieve an electronic document from the electronic job folder, as well as paper documents from the physical job folders. The electronic document was, in the context of the scenario, a list of changes that a client had asked to be made to a design. The participants had to compare two revisions of the design (in paper form) and decide whether these changes had been made. The changes themselves were simple and obvious, as the purpose of the task was not to evaluate the ability of the users to 'spot the difference', but rather to demonstrate how SOPHYA could potentially be used in a paper-electronic work-flow.

The text of the instructions for Task 4 is given below. For the complete version of these instructions, including the images they contained, see Appendix E.

> *During the lifetime of a job, the drawings done by the designers will be viewed by the client, giving the client the opportunity to request changes be made to the drawing. The designer will then make the requested changes. Before the drawing is shown to the client again, you are asked to review the changes that have been made against the list of changes provided by the client.*
>
> *This type of task involves working with documents both in the real world (the drawings) and on the computer (the change list). The following steps will walk you through the process of performing this task for the client named "Davie".*
>
> 1. *Browse or search for the* Davie *job using the search window.*
> 2. *Double click on the job name in the list of search results. This will bring up the job information window (shown on the right below). This is the "virtual" representation of the job. It shows information associated with the job that is stored in the database, as well as listing the documents for that job.*
> 3. *The "Change List" is an electronic document, stored on the computer. To view it double click the "Change List" document in the job window's document list (as shown above). This brings up the list of changes requested by the client.*
> 4. *The* Retrieve Folder *button (see above) can be used to retrieve the folder. Click it to illuminate the folder's LED and retrieve the folder.*
> 5. *In the folder you will find hard copies of the two most recent drawings (you can tell which one is newer by looking at the revision number in the bottom right hand corner of the drawing).*
> 6. *Compare the two hard copy versions of the document and check whether all the changes listed in the "Change List" have been applied. You may find the "Guide to Drawing Symbols" in the* in tray *useful.*

7. *If the changes were not applied correctly place the job folder in your "out tray", otherwise return it to the filing cabinet (if necessary refer back to the tutorial section* Returning folders to their shelves *to remind yourself of the correct procedure for returning folders).*

10.3.5 Task 5 — Generating alerts

The purpose of the fifth and final task was to demonstrate to the participants the capability SOPHYA has to display peripheral visual alerts. Since this task was intended as a demonstration of the alert feature only, it did not require the participants to actually setup alerts, but rather asked them to make use of alerts that had, in the context of the scenario, been previously setup.

Two types of alerts were demonstrated, a time-triggered alert that would be generated when a time-based condition was met (e.g., alerting the user to an impending deadline), and an event-triggered alert that would be generated when a given event occurred (e.g., receiving a certain job folder in one's in-tray). The two alerts that were demonstrated were described in the instructions that were given to the participants, which are included below.

> *The software has the ability to generate alerts at given times or under set conditions. When an alert is triggered the LED for a given folder will light up green (if it is an* in-progress *job), or blink (if it is a* completed *job). Potential uses of time-based alerts include alerting you when jobs need immediate action (for example getting close to deadline) or indicating jobs that might need attention (such as jobs that have not had any recent activity). Other alerts may be triggered by the software, for example receiving an email from a given client may trigger an alert for their job folder.*
>
> ***Part I***
>
> *For this task we will assume that an alert has previously been setup to highlight* in-progress *jobs that have had no activity within the past seven days. After seven days of a job having no activity the alert will be triggered and the folders green LED will light.*
>
> *In order to save confusion when performing other tasks, it is possible to turn the alerts off (as they have been in the previous tasks). To*

enable alerts click the Activate Alerts *button in the toolbar at the top of the main window (shown below).*

Now retrieve the folder(s) whose green LEDs are lit (i.e., those that have had no activity in the past 7 days) and place them in your "out tray", as you will give them to the manager to review.

Part II

Another alert was previously setup to be triggered when a given document was returned to you; this was to remind you to file the document back into its original folder. The alert is activated when the document's tag is scanned by the tag reader. When this happens it will cause the LED of the appropriate folder to turn on.

You will find the document, (shown below) in your "in tray". Scan its tag (on the back of the document) and return it to the appropriate folder.

10.3.6 Questionnaires

At the completion of each task, the participants were presented with a questionnaire to fill in. The questionnaire form was included as part of the software application that was used in the study. The questionnaire form needed to be filled out before the participants were able to proceed to the next task.

The questionnaire consisted of five questions, which are listed in Table 10.1. These same five questions were repeated for each of the five tasks the participants performed. Each question was made up of two parts: a rating and a justification for the rating. The participants gave a rating for each question on a scale of 1–7, with 1 being very difficult and 7 very easy. For the second part of the question, participants were asked to give the reason as to why they gave a particular rating.

Table 10.1: Questions that were asked at the completion of each task.

No.	Question
Q1.	How easy was it to understand the task you have just performed?
Q2.	How easy was it to perform the task using SOPHYA?
Q3.	How easy would it be to perform this task again using SOPHYA?
Q4.	How helpful was SOPHYA in assisting you to perform this task?
Q5.	How easy would it be to perform this task manually (without SOPHYA)?

10.4 Interview

After the participants had completed all five tasks, an unstructured interview (see Gorman and Clayton, 2005) was conducted with them. The aim of the interview was to gain an insight into the participant's background, and get feedback from them about the SOPHYA systems, based on their past office experience and the experience that they had just had using SOPHYA to perform the study tasks. To guide the interview a set of open-ended questions was used. These questions were used as a guide only, and were not strictly adhered to. The questions are listed in Table 10.2, where they are divided into three groups: those relating to the background of the participants and their previous office experience, those relating to the usefulness of SOPHYA, and those relating to the adoption of SOPHYA in the real-world.

Table 10.2: Questions used as a guide when conducting the unstructured, open-ended interview.

Background
Do/have you worked with paper documents in your work? How do/did you manage them? How large was the collection of documents? Would others use the documents/files as well? How would you collaborate/share the documents/files? Do/did you deal with sensitive files? How was security implemented? How did you control access to files?
Usefulness of SOPHYA
Based on your experience, do you think a system such as SOPHYA would be useful in the offices you have worked in? Why/why not? Can you see any other situations where it might be useful?
Barriers to adoption of SOPHYA
Can you see any barriers to using such a system?

10.5 Participants

The participants for this study were carefully selected to ensure that they had sufficient experience dealing with paper and electronic documents in real offices, and that they were familiar with the processes of filing, search, and retrieval of folders and documents. The purpose of this selection criteria was to make sure that the participants could provide a realistic comparison between the document management tasks they performed during the study while using

SOPHYA, and their real-life experience of performing similar tasks in their normal work environments.

A total of sixteen participants were selected for the study. Of these, ten were female and six male. Most (14) of the participants were from various units across the University of Waikato. Thirteen of the participants worked in offices where they dealt with paper documents and files on a daily basis, one participant had previously worked as a librarian, one was a researcher in the field of knowledge management, and one worked solely with electronic documents.

10.6 Application software

As stated in previous chapters, the SOPHYA prototype systems that have been developed do not include electronic document management software due to the variation in requirements for such software between offices. Because of this, a software application that provided document management capabilities needed to be developed for use in this evaluation. Also, since the focus of the study described in this chapter was evaluating the functionalities provided by SOPHYA, and not the usability of the application software, the design of the software was kept to a minimum level. Only the type of functionality that was needed for the study was implemented in the software, which was also tailored to the specific task scenario described earlier.

Perhaps the most important element of the developed client application was the search tool, which enabled participants to search for information about the job folders, and the documents they contained, using metadata associated with these jobs (e.g., keywords, start and completion dates, and job types). Figure 10.2 shows the search screen of the client application. Once a user had completed a search activity, the application software enabled the user to locate the relevant folders and documents by clicking on the retrieve button and then visually identifying them in the folder shelves (e.g., by turning their LED on to a specific colour).

As well as implementing the roles of DMS and client software, the application software also collected data needed for the evaluation, for example the time it took the participants to complete each task. At the end of each task, participants were asked to fill in a questionnaire (see Section 10.3.6). The client

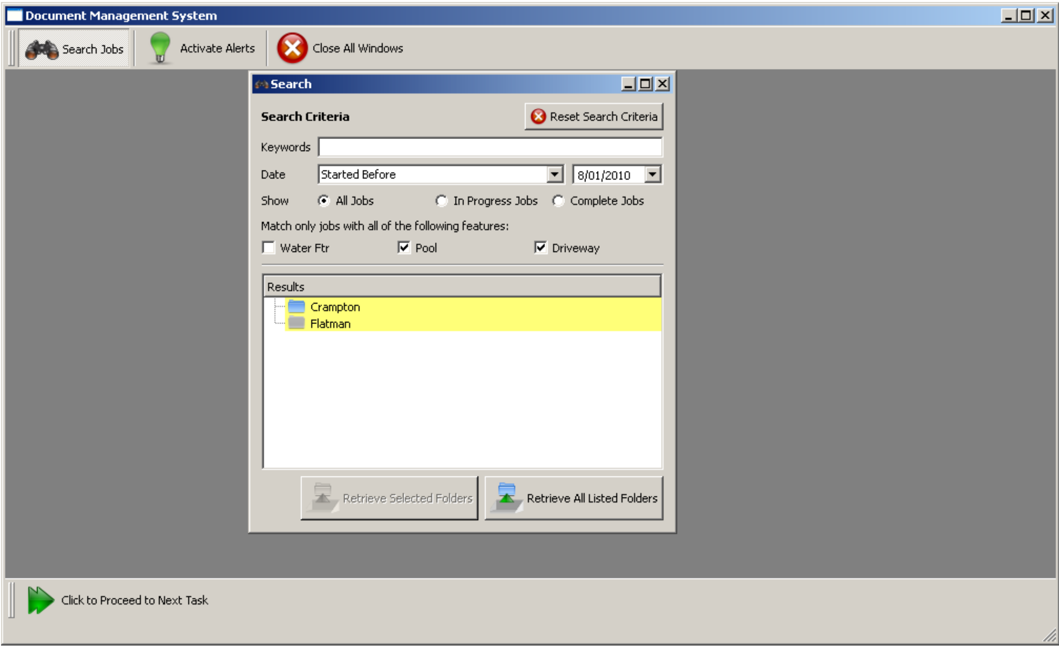

Figure 10.2: The Search form of the evaluation software, showing a search for all jobs that were started before 08/01/2010, and which include a pool and driveway. Two jobs match this search, Crampton and Flatman.

application would display the questionnaire form on-screen (Figure 10.3) for the participant to fill in before they continued on to the next task.

10.7 Quantitative Results

This section presents the quantitative results of the data collected during the study. It begins with the time taken to complete the tasks and the completion accuracy results for each of the tasks, followed by the ratings that participants gave on the post-task questionnaires.

10.7.1 Completion time and accuracy

The quantitative data that was collected during the study included the time taken to complete each task, and the accuracy with which these tasks were completed. The results of the analysis of this data are shown in Table 10.3. The completion times shown are the average across all participants, while the accuracy is the percentage of participants that completed each task correctly.

The time was measured from when the participants began each task (i.e., after the submission of the questionnaire for the previous task) until they began the

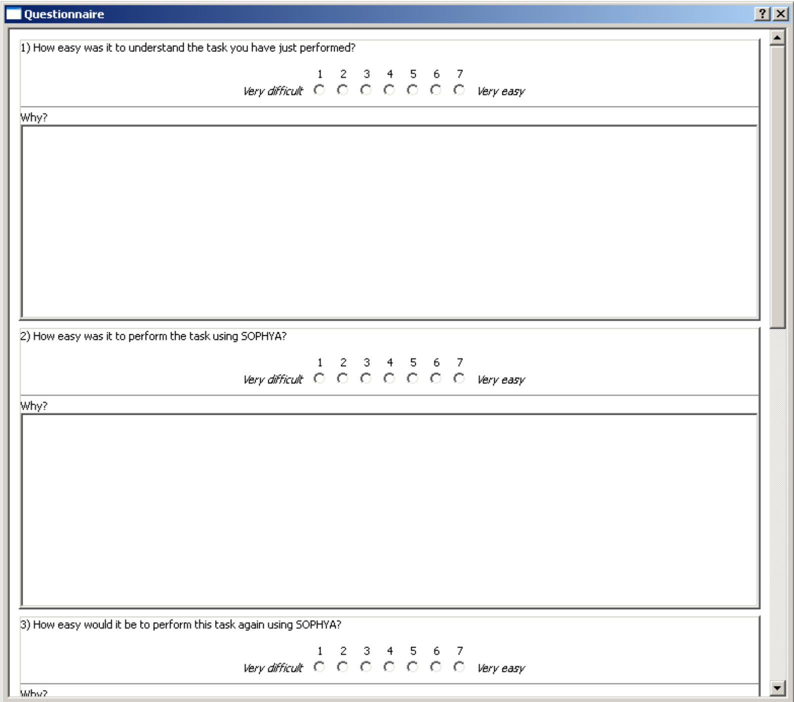

Figure 10.3: Screenshot of the questionnaire form that was displayed to users at the completion of each task.

Table 10.3: Percentage of participants who completed each task accurately (e.g., retrieved the correct folders) and average time taken to complete each task.

Task	Accuracy (%)	Time (mm:ss)	
		Avg.	Std. dev
T1.	94	02:35	01:38
T2.	50	02:55	00:27
T3.	81	05:46	01:18
T4.	94	04:23	01:15
T5.	100	03:04	01:41

questionnaire for that task. The average times taken to perform the tasks were between 2 and 6 minutes. Considering the fact that the participants were not familiar with the organisation of the documents and folders, and therefore had to rely fully on SOPHYA to perform the tasks, these times can be considered to be rather short in duration.

The completion accuracy is the percentage of people who completed each task correctly (e.g., retrieved the correct folders). The accuracy was high for all the tasks except Task 2. Most of the errors that reduced task completion accuracy were caused by what could be considered “clerical” errors, rather than being due to the participants’ use of SOPHYA. For instance, in the case of Task 2, participants were instructed to retrieve folders for jobs completed before 20/06/2009. However, ten of the 16 participants missed this detail when reading the instructions, and entered the correct day and month, but the wrong year (2010 instead of 2009) in the search form. As a result of this the search returned 12 folders rather than only two as was intended. On seeing such an unexpectedly large number of folder LEDs light up, three of the participants re-read the instructions and re-entered the criteria correctly. However, seven participants proceeded to retrieve all 12 folders.

In the case of the third task, participants were required to retrieve documents from several folders that matched some given criteria. For this task, the main difficulty that the participants encountered was working out how to correctly enter the search criteria. Those that did not complete the task correctly either did not realise that two separate searches needed to be conducted or entered the criteria for the second search without first clearing the criteria for the first search from the search form (which was necessary since the search results were updated dynamically as search criteria were entered into the form).

10.7.2 Questionnaire ratings

The average ratings given for each of the questionnaire questions of each the study tasks are shown in Table 10.4. The participants rated SOPHYA very highly in terms of ease-of-use, re-use, and the assistance it provided (Questions 1–4).

Many of the comments made by the participants, both on the questionnaire form and in the interview, supported their positive ratings of SOPHYA’s ease of use. During the interviews, participants commented on how easy or intuitive

Table 10.4: Median and inter-quartile range (in parenthesis) of ratings given for the questions of each task questionnaire (given in Table 10.1).

	Average Ratings				
	T1.	**T2.**	**T3.**	**T4.**	**T5.**
Q1.	6 (*0.640*)	7 (*0.063*)	6.5 (*2.000*)	6 (*1.300*)	7 (*1.000*)
Q2.	6 (*1.300*)	7 (*0.047*)	7 (*1.000*)	7 (*1.300*)	7 (*1.000*)
Q3.	7 (*1.000*)	7 (*0.000*)	7 (*0.480*)	7 (*1.000*)	7 (*0.000*)
Q4.	6 (*2.300*)	7 (*0.480*)	7 (*1.000*)	7 (*1.000*)	6.5 (*1.000*)
Q5.	4 (*2.300*)	3.5 (*3.000*)	3.5 (*2.000*)	4 (*2.000*)	4 (*2.000*)

it was to learn and/or use SOPHYA. While, the simplicity of the tasks and the design of the client application specifically for this evaluation (as opposed to a general purpose document management software application) may have had some bearing on the perceived ease of use, it is likely that many of the features of SOPHYA would remain easy to use even in a more complex setting. Some of the feedback related to ease of use included:

- "*With that [SOPHYA] you just find where everything is and where it should go.*"
- "*Because the light was flashing, it completely removed any kind of thinking process and seemed to speed up the file retrieval process hugely.*"
- "*It seems to me very simple, find it, there goes a light, swipe it, put it back; it's pretty intuitive in that regard.*"
- "*It's very easy to use once you've got the hang of it. And the more you used it the easier it is.*"

The fifth question concerned completion of tasks without using SOPHYA. The participants felt that, based on their experience, the study tasks would be rather difficult to complete without using SOPHYA, especially if carried out in a full-scale office. The three main reasons given by participants for their ratings for Question 5 were:

- SOPHYA saves time over manually looking through all the files;
- SOPHYA provides the ability to search the database for folders matching given criteria; and
- it is easier to find folders with SOPHYA when unfamiliar with the

filing system.

10.8 Qualitative Results

This section presents the feedback that was received from the participants during the study interviews. It is divided into sections based on the structure of the interview questions discussed in Section 10.4. The first section (Section 10.8.1) covers the participants' comments, based on their document management experience, as to whether or not SOPHYA would be useful in their own work environment, and their reasons for this. Section 10.8.2 covers the participants' views on what may be barriers to the adoption of SOPHYA in real-world offices. Finally, Section 10.8.3 describes several improvements that could be made to SOPHYA based on the feedback received.

Direct quotes have been transcribed from audio recording of the interviews and some have had minor editing to improve readability.

10.8.1 Usefulness of SOPHYA

Six of the thirteen participants who worked with paper documents on a daily basis said that they would find SOPHYA useful in their office. The seven participants who did not think SOPHYA would be useful in their office gave reasons for this, which can be grouped into three main categories.

- Their office work-flow focused more on individual documents, rather than folders.
- The number of folders that they worked with was small enough that they knew them well and would not need to use software for assistance.
- The organisational system that they already had in place was serving them well enough.

The first point shows that in order for SOPHYA to be suitable for these offices, it needs to be integrated with some form of individual document tracking to keep track of where documents are (e.g., which folder they are in). This was discussed previously in Section 9.4. To provide such functionality SOPHYA could be combined with, for example, an RFID-based document tracking system such as the one described by Arregui et al. (2003). This is discussed further in Section 11.3.3.

Regarding the second and third points, the organisational systems may work well for their current users, but may not be suitable if other people needed to use them. For example, two of the participants who felt that their current organisational system served them well enough noted that SOPHYA would be useful when accessing their co-workers' filing systems. Similarly, another participant said of his filing system: "*it works for me cause I know it, and I memorised it. If someone else sat in my chair and had to work it all out they'd probably struggle with that. Because the logic that I use, and the labels that I use, make sense to me, but it may not make sense to someone else.*"

Two of the participants who did not feel that SOPHYA would be useful in their office said that although they knew their organisation well, and could easily find folders, SOPHYA could be useful when needing to search based on criteria which they had not previously envisaged. Additionally, two of the participants who said SOPHYA would not be useful in their current office said that it would be useful in a previous office they had worked in.

Usefulness in other offices

The study participants were also asked if there were any other types of offices that they thought could benefit from using SOPHYA technology. Suggestions from the participants included:

- large paper-based offices such as legal offices and accounting firms;
- offices that deal with large, unwieldy documents such as diagrams or maps, which are better suited to use in physical rather than electronic form; and
- offices where there are many casual or temporary staff, since SOPHYA allows a filing system be used by someone who is not familiar with it.

The use of SOPHYA technology to support management of library books has been discussed previously (see Section 9.2.2), and the benefit of SOPHYA to this type of environment was confirmed by one of the study participants who had experience working as a librarian. On the topic of misfiling she pointed out that "*once [a book] is misfiled it's useless, so if this can take out misfiling things it is very appealing*". Regarding SOPHYA she remarked, "*the lights are* *what would save a lot of time and energy*". The use of SOPHYA technology in libraries would simplify the process of shelving and reduce the training required for this task, which would allow skilled librarians to concentrate on other tasks.

Another use for SOPHYA technology that was suggested, though not specifically related to an office environment, was to assist with storage and retrieval of spare parts in the aviation industry. One of the participants described a system that is currently used for storing spare parts, in which parts are allocated randomly to storage bins, rather than being grouped (e.g., by type: screw, bolt, etc.) or ordered (e.g., by size). This means that similar looking parts are not located together, reducing the likelihood of the wrong part being retrieved. The participant stated that SOPHYA would be perfect in this scenario, since this would enable the search for a part to be performed electronically, and the result displayed using a light to indicate the bin that the part needs to be retrieved from. Even though this is not specifically a document management task, SOPHYA is able to support the management of physical artefacts in general.

Potential benefits of SOPHYA

During the interviews, the participants described what they felt were benefits of SOPHYA. This feedback has been grouped into categories and is presented below along with discussion.

Saving time when retrieving and returning documents. A major benefit of SOPHYA that was identified by the participants was that it would save time when retrieving and returning documents. One of the participants noted that "*having a system like this would just cut down that down time of just fumbling through*". Another participant remarked: "*that was what I thought was a big thing, being able to know where to put something back, that must add up to probably weeks of saved time over a year for people*".

One of the participants spoke of his past experience in an office where "*they'd have at any time probably over 1,000 contracts active. Each one had a paper file which had the hard copy of the contract and a bunch of other stuff that had to be there, and usually a bunch of electronic files associated with it. It was an on-going disaster to keep those things matched, getting them back in the right place, losing them, knowing who's got them, etc. I can easily imagine something like this saving that system hours of work. Just being able to 'we need* *to find contract such and such', 'oh there it is', grab it*". He said that he could see SOPHYA being useful in "*organisations in those sort of cases, especially if you've got significant numbers of files and need to keep paper copies, and need to have other people access them and that sort of stuff*".

Talking about the time that could be saved with a system like SOPHYA he said the following: "*for people doing invoicing and stuff that requires interaction with the files it must be like an hour a week or a significant chunk of time when you add up the process of 'where is this file' and then there'd be some cases where it would save heaps ... where there's emails going around the organisation and people hunting high and low. So I would say it's not unreasonable to say it's quite significant, maybe an hour or a couple a week for people that have to interact with those files a lot*".

As discussed in Section 4.4, increasing the efficiency of manual operations such as retrieving and returning documents was a core aim in the development of SOPHYA. The costs of inefficient document retrieval have been described in previous research, such as that of Blair and Gordon (1991) who state that "*businesses pay a real financial price when they cannot locate needed information, retrieve and act upon information that is incorrect or outdated, or spend inordinate effort attempting to locate important information*". There is, therefore, real benefit to be achieved by improving this, and the feedback quoted above shows that SOPHYA has strong potential to reduce the overheads of locating and retrieving paper documents in real offices. This is related to two of the requirements that were introduced in Section 4.4.1: *simplify filing and retrieving of documents* and *reduce time locating documents and document containers*. The feedback from the participants shows that SOPHYA has made progress towards meeting these two requirements.

Integrating electronic searching and real-world display of results. While the use of an output display mechanism such as LEDs enables folders to be quickly located, it was pointed out by some participants that this by itself may not be of great benefit. However, the fact that this ability was combined with electronic searching was described by participants as a benefit. For instance, one of the participants said that "*if it was just 'oh where should I put this file' then it's probably not really going to be that successful, people are going to say 'it's just as quick to find it alphabetically,' but when it's connected to searching for certain things it is really good*".

Similarly, another participant noted that "*for reporting, being able to use the search thing that you've got there [is a benefit] . . . we could look at what contracts are expiring at a certain date and put up all those ones. [It] would be quite useful*".

This feedback highlights the benefit of SOPHYA as a system for integrating physical and electronic document management, since it is the ability to use electronic metadata when performing actions with physical documents (e.g., search and retrieval) that makes SOPHYA useful.

Ensuring folders are returned to the correct place. In environments such as large offices where multiple people share collections with complex filing schemes, the misplacement of files can become an issue, as exemplified by the following quote. "*Often with that amount of files, go to look for a file and someone has used it and has put it just a little bit out of place, because either they've transposed the numbers in their mind as they've put it back, or put it according to the numbers, but then we have special files that are not filed by that numerical number—they have their own separate shelf as well because they are special projects*". Similarly, another participant stated: "*not having to deal with the problem of misfiled stuff is very useful because that's the most time consuming. Other people come and fiddle with my files and they don't necessarily put them back in the right place. Then I'm thinking where the hell's my file, and then I've got to manually scan through each one, instead of just relying on the numeric*".

Without integration between the physical folders and the electronic document management system, once a folder is misfiled it has to be found by manually searching through the collection. However, with a system such as the ordered version of SOPHYA it is possible to assist the users so that they place folders in their correct places. The alert feature can also notify users when a folder is misplaced.

Alternatively, SOPHYA can remove the reliance on ordering as a means of locating folders, since it can guide the user to selected folders using its output display mechanism (in this case LEDs). One participant commented on how this "*changes your mindset about priorities*", regarding feeling in control of the organisation of folders her filing system, "*intuitively you want to sort them, you feel like you need to have them in alphabetical order or date order, or something . . . so that you feel like you have got control over them. But you can't possibly have control*". However, she went on to say that SOPHYA removes the need to have the folders in a specific order: "*you don't need to have them in alphabetical order, you don't need to have them in date order. The LEDs and the tracking tells you where they are. Fabulous*".

The ability of SOPHYA to ensure folders are returned to their correct places demonstrates that SOPHYA is able to *simplify the filing and re-filing of documents*, which is one of the previously identified requirements.

Locating files outside a collection. The ability to locate a file that has been removed from a document collection was seen as another benefit of SOPHYA. For example, "*there's always emails going around 'has anyone seen the file for such-and-such'*". Similarly, another participant said "*with the files I deal with, sometimes I'll have to, say for example, get sign off from another manager. So I will give the files to them. In my tracking spreadsheet I will put down—so I know where it is—'have put on bla-bla-bla's desk'*". Another participant remarked "*the hardest thing is when files get removed from the spot where they live. The first difficulty is if they get taken from the office, then it's all compounded exponentially when they get moved around beyond that. You don't know who's got them, or where they've gone. You don't know how to get them back. That's always the biggest issue, is the disappearance. This would track files brilliantly, it would be fantastic*".

Although SOPHYA itself cannot track the movement of containers outside of storage locations, it can for example log who removed the container. In relation to this, one participant suggested that requiring the users to swipe their ID card when retrieving a folder from the collection would mean that the system could keep track of who has taken which file. This concept has also been discussed in previous research, such as that related to RFID document tracking (e.g., Arregui et al., 2003). In order to provide tracking of containers outside of storage locations, SOPHYA would need to be combined with some other tracking mechanism, such as RFID (as discussed in Section 11.3.4).

Making it easier to use an unfamiliar filing system. Making it easier for people to use an unfamiliar filing system is considered to be another benefit of using SOPHYA. This is something that is demonstrated when people try to use a filing system they are unfamiliar with: "*You really notice how complicated your filing system is or isn't when you have new staff. They just have no idea, can't find anything*". However, SOPHYA could mitigate this problem: "*you could have part time [staff] come in if you had a system where the manager* *could find what he/she needed when someone wasn't there*".

Cole (1982) notes that the more frequent and dynamic a person's interaction with information is, the more familiar they become with it. It is logical, then, that new or temporary staff, as well as staff who do not frequently interact with a filing system (e.g., a manager) will be unfamiliar with it, and therefore likely to find it difficult to locate appropriate information within it. A system such as SOPHYA can enable these sort of infrequent users to perform tasks such as searching, retrieval, and returning of documents (as discussed above) without requiring an intimate knowledge of the filing system. In fact, one of the requirements identified in Section 4.4.1 was that such a system should *reduce difficulty of finding documents and document containers in an unfamiliar collection.*

Ability to display alerts. The ability of SOPHYA to display alerts was considered valuable by a number of participants. One participant mentioned that she would be more inclined to pay attention to flashing lights on the folders than to electronic reminders: "*if there were flashing lights on my rotascan I would so deal to those quickly before anyone else saw them. You're sitting there and your boss comes in and sees lots of flashing lights on these files, they're gonna know you've been slack*".

The ability to generate and display alerts was identified as a requirement of the system in Section 4.4.1. This concept is not new. For example, one participant described the computer-based reminder system used in her office, which is effectively an electronic to-do list updated each Monday for the coming week. Similarly, Arregui et al. (2003) describe how rule-based alerts could be generated in conjunction with RFID document tracking and displayed on a mobile device or on a screen attached to a filing cabinet.

SOPHYA is able to display alerts on the appropriate folders or storage locations, which means that they appear peripherally. Alerts on the computer are considered more "*in-your-face*" and could be easily closed; whereas, it was suggested by one of the participants that if she was required to walk over to the folders in order to acknowledge the alert, then "*I might as well pull the file out now that I have to go all the way over and touch it*".

The feedback regarding alerts shows that this feature of SOPHYA was deemed useful. Reasons for this included the visibility of the alerts (both to the user who they are intended for, and to other people such as their managers) and their peripheral nature. Because the alerts are not displayed in a way that interrupts work, they can act as a persistent reminder until such a time as appropriate action is taken, rather than being dismissed because they are interfering with the user's immediate work (as may be the case with on-screen alerts).

Ability to identify infrequently used folders. Because SOPHYA incorporates logging of the most recent activity of the physical containers (i.e., retrieving from or returning to their storage location), it is possible for the DMS software to detect which folders have not been accessed in a certain period of time. This ability of SOPHYA was demonstrated in the first part of the fifth task, where the participants had to retrieve folders that, in the context of the task, had had no activity for the past seven days. This was highlighted as a useful feature by one of the participants: "*sometimes you're so busy with other students, ones that are knocking on the door all the time, you forget about the ones that don't even bother to turn up to anything. Having some kind of system like that so you can keep on top of everything would be awesome*".

This is an example of the use of metadata that is automatically generated from the users' interactions with the system in combination with the alert feature of SOPHYA. This is again similar to the rule based-alerts described by Arregui et al. (2003). However, in the case of SOPHYA the alerts are combined with the output display on the folders or filing cabinet. Identifying infrequently used folders is just one example of the application of such rule-based alerts.

Changing filing behaviour. It was suggested by one of the participants that because SOPHYA makes the task of filing easier, this could have a positive impact on people's filing behaviour. "*I have a mate . . . and his office is covered in piles of paper waiting to be filed, and he never can find anything. I keep saying 'I sent that to you a month ago, you should have it already, why am I sending it again?' because he never does his filing. But if you change behaviour by making things that easy to do, then that's gotta be a plus—then you're increasing productivity too, and that's always a go for employers*".

Whether or not SOPHYA would be of benefit in such a situation would depend on the reasons behind such filing behaviour. It may, for example, be a means of delaying classification because the person does not have enough information to know how to file a given document, in which case SOPHYA could assist by allowing the document to be physically filed into a 'miscellaneous' folder and then have classifications managed electronically. Similarly, the filing of a document may be delayed in order to keep it visible so that it can act as an ongoing reminder that it needs to be followed up. SOPHYA may be able to assist with this, for example by integrating with an electronic to-do list and using the alerting feature. Another potential reason for delaying filing is that it is time consuming, or perceived as such, and if SOPHYA is able to reduce the time taken to perform this task (or the perception of the time taken) then it may be able to change filing behaviour. It is not possible to make such a conclusion based on the information gathered in this study, but it would be an interesting theory to test in a future study (e.g., a longitudinal study of the effect of SOPHYA on a filing behaviour).

Reduction of duplication in shared and personal archives. Another potential benefit of SOPHYA is that it could reduce duplication of documents in shared environments by making them easier to find. Speaking of his previous office experience, one participant remarked "*it was so much of a pain to find the contracts I ended up printing them myself. Every year it must have been over a 1000 pages of contracts and other stuff that was all in the files but it was just such a performance to find it some of the times that it was just easier to print it out*".

This is similar to the findings of Whittaker and Hirschberg (2001) that one reason as to why people kept personal copies of shared documents was that they did not trust the library or archive system where the shared documents were kept. For example, one of their participants stated: "*I started making copies of the articles I really cared about, and many of the articles in the cabinets in here are those kinds of articles . . . if I had guaranteed going forward access to a library that I knew would supply my needs, I would be happy to donate them to a worthy home*". Therefore, by making it easier for people to find the documents that they need within a shared archive, they may no longer feel the need to keep their own personal copies. This has the added benefit of a reduction in unnecessary printing and use of storage space, which is one of the previously identified requirements of the integrated document management system.

10.8.2 Potential barriers to adoption of SOPHYA

The aim of this study was not only to evaluate whether SOPHYA would be potentially useful in offices, but also to identify any problems that need to be taken into account to enable such a system to be useful in the real world. For this reason, the participants were asked if they could foresee any barriers to the deployment of SOPHYA. The participants provided a range of responses, which are summarised below.

Cost. A potential barrier that was noted by the participants was the cost of deploying and maintaining such a system. There would need to be a demonstrable benefit to using the system (e.g., savings in terms of time and effort, reduced document loss, etc.) to outweigh such costs. As one participant noted: "*being inefficient might prove to be cheaper, even though it might be more frustrating and annoying*". "*You'd have to see if doing it this way uses less energy overall compared to your traditional style of office maintenance which is put it in a pile somewhere, then physically search through it later*".

Costs involved with SOPHYA could include the initial setup—in terms of time and equipment, especially if retrofitting an existing manual system—as well as on-going costs as new folders are added and existing ones are maintained. However, such costs are not specific to SOPHYA and would similarly apply to any method of augmenting the physical document management system with electronic hardware. In fact, because SOPHYA manages containers rather than individual documents, this would require fewer artefacts to be augmented than with a system that tags every document.There is, therefore, a trade-off between the cost of deploying such systems and the degree of functionality that they can provide. A balance between these two factors would need to be found in each office, since the needs of offices vary.

Training is another potential cost of deploying SOPHYA that would have to be taken into account. However, this study has shown that learning to use SOPHYA is not likely to be very difficult or time consuming. Additionally, the training required to use SOPHYA is likely to be less than the training that would be required to introduce staff to a manual system that they are not familiar with.

Data entry. Another potential problem that was mentioned is that a system such as SOPHYA is dependent on accurate data entry. If information about documents is entered incorrectly or inconsistently into the system, it would make it more difficult to find documents. "*Probably management's ability to understand, to drive the discipline of the online system and have it match the offline [would be a barrier]. Often that's where the problem is in this sort of stuff, getting everyone inputting things the same way. And maybe you wouldn't need to have everything in an organisation done that way, but you'd have several key staff in head office to make sure it's done*". Similarly, another participant remarked that: "*you need a good secretary to put all the tags on, enter it all into [the computer]. It requires good organisation on the part of the secretary. If it was standard procedure, it makes everything down the track a hell of a lot easier, but being organised in the first place is probably the toughest thing*".

Reliance on accurate data entry is an issue that is not unique to SOPHYA. There are, however, many ways to reduce this problem through automating elements of the data entry process and providing intelligent assistance. For example, providing auto completion, or suggestions based on context and previously entered data, similar to the "Did you mean" feature provided by Google.

Deeper organisational problems. Yet another potential barrier, closely linked to the issue of data entry, that was pointed out was that if the organisation in which SOPHYA is deployed is not itself organised then just adding new technology is not necessarily going to solve the deeper problem. "*But that's like any technology. If you're not good at running meetings then doing teleconferences isn't going to help you. Your technology emphasises what you're good or bad at rather than fixing something necessarily. But for those ones that do have systems and do have policies around that sort of stuff I'd imagine they could almost use it straight away*".

As evidenced by the quote above, this issue is not directly a problem with SOPHYA, and would similarly apply to other technology. This is an organisational problem, and while deploying a system such as SOPHYA cannot overcome such problems by itself, it may provide an incentive to initiate policy to address such problems.

Staff resistance. The response of the office staff to the system may also be a potential barrier: "*there could be resistance on the PA (or whoever would do this job) side, maybe they don't like any technical advancements*". Similarly, "*with something new, depending on how it was introduced to the staff that had* *to run it [may be a barrier]. It is not a barrier with the technology itself, because it seems to me very simple ... and it's not as if you're saying 'you can't have a coffee break now,' it's taking away a job that is searching through things. So it would depend on how it's introduced and implemented*".

As this participant noted, this is not a limitation of the SOPHYA technology itself, and is a potential issue for any new system introduced into an organisation. This serves as a reminder of the importance of the human aspects of deploying technology in organisations, which are addressed by applying system development techniques such as participatory design (see Gould and Lewis, 1985).

Personalisation. The need to allow individual users to personalise their system to their own way of working was seen as a potential issue: "*everyone might have a different way of dealing with their files. That could be a barrier in terms of trying to find a way to personalise it to every person*".

This is similar to Cole's (1982) finding that "*organisation was idiosyncratic, dependent on the demands on the users and their motivations*". SOPHYA has been designed with flexibility as a primary consideration. Although, the flexibility of a deployed system would be dependent on factors beyond the underlying SOPHYA technology (e.g., the DMS and client software), SOPHYA itself is able to support, for example, both piling and filing organisational strategies.

Risk of breakage. The risk of the system breaking would be yet another perceived barrier. It was pointed out by one of the participants that problems due to the system breaking down could be exacerbated if staff with less filing training became dependent on it: "*then if something went wrong they may not notice a problem or not be able to deal with it*".

This is another potential issue for any system that is developed (e.g., wireless systems may have problems with interference, visual systems with line-of-sight interruptions, etc.). Reliability is an important factor that would have to be taken into account during further revisions of the system and, as has been previously discussed, is an issue that could be addressed by industrial design.

10.8.3 Suggestions for improvements

The participants also gave feedback on issues related to the design of the system that, if not accounted for, may become barriers to its use.

Visibility of LEDs. Some of the comments made regarding the LEDs indicated room for future improvements over the current design. One issue that was pointed out in this regard was the visibility of the LEDs. There were cases where the LED on a folder was obscured by other folders due to the place where the participant was sitting. Once the participant moved so that they were directly in front of the storage location all LEDs were visible.

It is also worth noting that although LEDs are the user interface output device used by the current SOPHYA implementations, it is possible for other output components to be incorporated into future revisions. For example, in order to assist in locating a container in situations where it may be obscured (e.g., at the bottom of a stack) auditory output could be included (e.g., in the form of a beeper). This would enable users to locate containers in a similar manner to the commonly used method of locating a lost phone by calling it and listening to it ring.

Confusion between multiple users. Another potential problem to be aware of when designing future systems for larger environments is potential confusion between multiple people using the system concurrently (e.g., identifying which LEDs are illuminated for which user). This problem could be countered, for example, by using multiple colours of LEDs and assigning each user a colour or, when displaying search results, informing the user which colour to retrieve. However, a related issue pointed out by a participant was that the choice of LED colours would need to take colour blindness into account.

As mentioned above, it is possible that alternative user interface output devices could be incorporated into containers and storage locations to overcome limitations of LEDs. Some possible future user interface designs are discussed in Section 11.3.2.

Dividers. A design issue that was pointed out by several participants was the space taken up by the dividers in the storage locations (these can be seen in Figures 7.6b and 8.6). "*I was just thinking about the storage space . . . at the moment we've got big filing cabinets, and they're absolutely chock-a-block full of files. So if we were doing a system like this [indicating the dividers on* *the shelves] they'd obviously take up more space*". Similarly, "*infrastructure to support it. [Indicating the dividers in the cabinet] can't be too clunky. When you're keeping records for seven years they take up a lot of space*".

Dividers are not a necessary feature of SOPHYA, and were simply included to aid reliability in the current storage location implementations. The important observation from this feedback is that when designing the physical hardware of the containers and storage locations, it is necessary to ensure they do not take up more space than necessary. This is another issue that can be addressed by industrial design.

10.9 Summary

This chapter described a study that was conducted to evaluate the two SOPHYA prototypes described in Chapters 7 and 8. The study was laboratory-based with sixteen participants chosen for their work experience with document management. The study comprised a tutorial, a set of five tasks, and an interview. The results of the study were positive, with participants rating SOPHYA's ease of use highly. A number of benefits of SOPHYA from the perspective of the participants were identified, as well as potential barriers to its deployment, and suggestions for improvements.

Chapter 11

Conclusions and Future Work

This concluding chapter provides a summary of the thesis (Section 11.1), discusses how its research questions were addressed (Section 11.2), identifies directions to future research (Section 11.3), and provides a final summary (Section 11.4).

11.1 Thesis Summary

This thesis has described the design, implementation, and evaluation of a tangible user interface technology for seamless integration of paper and electronic document management. This research has been motivated by the fact that paper documents are still commonly used in present-day offices, and their continued use alongside electronic documents raises two significant issues that need to be addressed. Firstly, paper documents have severe shortcomings in comparison to electronic documents with respect to their management. Secondly, the co-existence of paper and electronic documents in a modern office setting leads to fragmentation of information and document management between the two forms. Therefore, the aim of this thesis has been to investigate the seamless integration of the management of paper and electronic documents, and to develop and evaluate a system that achieves this objective. As such, the primary research question that this thesis has attempted to answer has been:

> *To what extent can tangible user interaction facilitate a more seamless integration of paper and electronic document management?*

In order to answer this question, four secondary research questions have been identified:

1. *What roles do paper and electronic documents serve, how are they managed, and what problems are encountered with their management in present-day offices?*
2. *What are the requirements of an integrated paper and electronic document management system?*
3. *How could a tangible user interface be designed to provide a more seamless integration of paper and electronic document management?*
4. *To what extent would the proposed tangible user interface fulfil the requirements of an integrated paper and electronic document management system?*

These four secondary questions correspond to the three main parts of the thesis:

- the study of existing offices in order to identify the requirements (questions 1 & 2);
- the design and development of a system for integrating management of paper and electronic documents (question 3);
- and the evaluation of the developed system (question 4).

Each of these parts is described in more detail in the following sections.

11.1.1 Identification of requirements

The first part of this thesis aimed to answer the following research questions:

> *What roles do paper and electronic documents serve, how are they managed, and what problems are encountered with their management in present-day offices?*
>
> *What are the requirements of an integrated paper and electronic document management system?*

These questions were addressed by conducting an observational study of document use in offices, as described in Chapter 4. This study identified that paper documents are still prevalent in offices, and identified the roles that they serve in the offices studied. It also investigated the ways that these documents are managed and the problems that are encountered when managing them. Furthermore, the study led to specification of a set of requirements for such systems that integrate physical and electronic document management.

Summary of the role of paper and electronic documents and their management

The main roles that documents serve in the offices that were studied are keeping records (i.e., documents that are kept for legal or administrative reasons), facilitating communication (both internally within the office and externally with other entities), providing a means of authoring and proofing, serving as reference material, and facilitating data collection. There is some variation between the offices as to how much information is kept, and managed with in paper form. While in some cases documents have a short lifespan (e.g., documents that are used for authoring and proofing), in other cases (notably in the case of records) documents are kept for mid- to long-term.

In terms of management, it is noted that a hierarchical organisational structure is often present, with documents being organised into units (often known as files). Maintenance is an ongoing task which includes maintaining indexes of the collections, filing and re-filing documents, and moving documents from one level of storage to another (e.g., archiving).

Problems with managing documents that were identified include the fact that it is labour intensive and time-consuming, especially in terms of filing paper documents, finding, and retrieving them. Other problems include duplication of paper documents (both intentionally, for organisational purposes, and unintentionally), loss of documents, storage of irrelevant documents, and the fact that organisational schemes can be user-specific, and therefore difficult to use by others who are unfamiliar with them. Issues due to fragmentation of document management between paper and electronic forms were also investigated. It was identified that it is difficult to maintain consistency between the paper and electronic document collections, and link paper and electronic documents.

Summary of the system requirements

Based on the findings of this study, a set of fifteen requirements was identified as a guide for designing an integrated paper and electronic document management system. These requirements were grouped into six categories (corresponding to the framework described in Section 2.2.5): organisation, maintenance, control, storage, retrieval, and other miscellaneous requirements.

11.1.2 Design and implementation of an integrated paper and electronic document management system

The second part of the thesis centred around the design and implementation of a system for integrating paper and electronic document management based on the requirements identified previously. This part aimed to answer the following question:

> *How could a tangible user interface be designed to provide a more seamless integration of paper and electronic document management?*

It was noted in Chapter 3 that existing systems for integrating paper and electronic document management largely focus on using some form of tagging to link paper documents to electronic information and to track their movement. As such, the feedback that they provide to the user is through an on-screen interface, as opposed to output that is embodied in the document management system itself. Therefore, having identified the requirements of the system (as described above), the design and development of a prototype system that implements a tangible/embodied user interface to integrating paper and electronic document management was investigated and the resulting prototype is described in Chapter 6. Chapter 7 described a system called SOPHYA designed to overcome the limitations of the first prototype with respect to flexibility, scalability, reliability, and interaction. A second SOPHYA prototype system was then developed, as described in Chapter 8, which added the ability to detect the ordering of document containers within their physical storage locations.

Design decisions

As with any design process, the system designed as part of this research required making a number of design decisions, which have shaped the developed systems. The main design decisions that were made are discussed below.

Tangible/embodied user interface. Previous research has focused on the tracking of documents, but has focused less on embedding user interfaces in artefacts used as part of the physical document management system itself. By embedding the interface in the physical document management artefacts, a step is taken towards making human-computer interaction seamless by integrating real-world actions with electronic management of documents.

Containers vs. documents. An important design decision of the systems that were implemented was to choose to manage documents at the container level rather than the document level. This decision is due to the implementation of the document management system as a tangible/embodied user interface. It would have been impractical to implement such a user interface at the document level. As was previously noted, in most offices documents are organised in some type of containers (e.g., folders), which matches with the container components managed by the systems developed as part of this thesis.

However, the implementation of the interface at a container level has some implications for the tracking of documents. It means that it is not possible to track individual documents. However, this is a topic that has been addressed in previous research (e.g., Bodhuin et al., 2007; Arregui et al., 2003; O'Neill et al., 2006). Furthermore, if such tracking is required, a range of technology could be used in combination with the system developed as part of this thesis. These possibilities will be discussed in Section 11.3.3.

Wired vs. wireless. Another design decision that has had implications for the functionality of the systems that have been developed was the choice of power supply and communication medium. A wired rather than a wireless connection was chosen for both of these purposes. This has enabled the systems developed to overcome the limitations of wireless communication, such as the inability to precisely determine the positions of containers, and has made it possible to supply power to more advanced interface components on document containers.

This decision, however, results in some limitations to the system. Because communication relies on contact between the container and storage location, when a container is removed from a storage location it is invisible to the system, until it is placed back into a storage location. These limitations can be overcome by integration with some form of wireless tagging in order to track the containers when they are removed from storage locations. This is an option related to the issue of tracking individual documents, and will be discussed in Section 11.3.4.

11.1.3 Evaluation of the system

The final part of the thesis concerned the evaluation of the prototype systems that have been developed. This aimed to answer the following research question:

> *To what extent would the proposed tangible user interface fulfil the requirements of an integrated paper and electronic document management system?*

The first part of the evaluation that was carried out (see Chapter 9) was to measure the extent to which the systems that were developed met the requirements specified in Chapter 4. This was followed by a user study, which is fully described in Chapter 10. The purpose of this study was to evaluate the system by getting feedback from people with past experience of working in the types of offices where such a system would ultimately be deployed.

Fulfilment of requirements

It was demonstrated that, when combined with appropriate software components, SOPHYA would be able to meet all of the identified requirements, except those requiring tracking of individual documents. If tracking of individual documents is required, then SOPHYA can be combined with existing technologies to support this. Similarly, while SOPHYA can determine which storage location a given container is currently placed in, it is not able to track containers while they are outside storage locations. However, if this feature is required then SOPHYA can easily be combined with existing wireless tracking technology to support it.

User evaluation results

The results of the user evaluation of the SOPHYA prototypes were very positive. The study showed that the potential users of this system believe that it would be useful in their offices. Benefits of the system that were identified or confirmed by the study demonstrate that SOPHYA would:

- save time when retrieving and returning documents;
- assist with locating containers;
- ensure accuracy when returning containers to a collection; and

- make it easier to use an unfamiliar filing system.

Although several barriers to adoption of the system were also identified, these are the kinds of barriers that would be met by the introduction of any new technology into existing office environments. However, it is important to be aware of these in order to be able to counter them when a system such as SOPHYA is adopted for document management in such settings.

11.2 Answering The Main Research Question

The previous section summarised the contributions of this thesis and provided an overview of how the four secondary research questions were addressed. These contributions, when combined together, address the main research question of this thesis:

> *To what extent can tangible user interaction facilitate a more seamless integration of paper and electronic document management?*

The conclusion is that there is clearly no single system that would satisfy all the requirements of an integrated system, as there are many ways of approaching the task of integrating paper and electronic document management. In order to develop a solution that would satisfy its users' needs, this thesis has followed a process that investigated the users' needs, and then brought elements of ubiquitous computing and tangible/embodied user interface design to create a series of prototypes. The results of these evaluations demonstrate that it is possible to develop systems that can utilise tangible user interaction techniques to enhance the integration of paper and electronic document management, and thus better bridge the divide between the physical and virtual worlds of documents.

11.3 Future Work

The research in this thesis has investigated the need for better integration between paper and electronic document management. As part of this work three prototype systems have been developed, and an evaluation of the two main prototypes carried out. Although the results of this evaluation were encouraging, several areas that would benefit from further improvements have been identified. This section describes how these improvements could be made to

SOPHYA, and outlines some of the logical extensions of the research presented in this thesis.

11.3.1 Integrating the unordered and ordered SOPHYA systems

As previously discussed, the two SOPHYA systems that have been developed are able to serve complementary purposes, since each supports certain features that the other does not. The differences are discussed in detail in Section 9.1, and can be summarised as follows:

- the unordered system supports stacking of containers, whereas the ordered system does not;
- the ordered system supports detection of container position and ordering, while the unordered system does not;
- the user interface devices (e.g., LEDs) of the unordered system are part of the containers, whereas those of the ordered system are part of the storage location.

Chapter 9 discussed how these two systems could be used for different purposes (e.g., the unordered system could be used for containers of hot and warm documents, while the unordered system could be used for cold documents). There are, however, several drawbacks that result from having two separate systems. For example, if an ordered container is placed into an unordered storage location it will not be detected, and vice versa. It would, therefore, be desirable to integrate the two systems.

Integration of the unordered and ordered SOPHYA systems could be achieved by defining a common conductor layout and communication protocol for all container types. Such an interface would support the features of both the unordered and ordered systems (e.g., stacking, position determination, and user interface devices on containers). However, the level of complexity of containers and storage locations could be varied depending on their use. For example, a desktop storage location may support stacking but not position determination, or an archival container may not include user interface devices.

The container design shown in Figure 11.1 is an example of how container conductors could be laid out so that containers could support stacking, position determination, and user interface devices. This design is similar to that of the unordered container, but has one significant difference. The conductive strip that is used for data line is split so that the conductive pad on the spine of the container is electrically isolated from the strips on the front and back of the container. Thus, the conductive pad on the bottom of the container can be used for position detection without it shorting to those of its neighbouring containers, and the strips on the front and back of the container (which are wired together) can be used when the container is stacked.

The level of circuitry on these containers would depend on their intended use. The minimum required circuitry for such containers would give them a unique ID and allow a storage location to read this ID through either the spine pad or front/back strips (a container should support both the pad and the strips so that it can be identified in any type of storage location). If a container is to have more advanced circuitry (e.g., LEDs) it could make use of the optional power supply strip.

11.3.2 Extending the functionality of SOPHYA

The SOPHYA implementations described in Chapters 7 and 8 have output displays in the form of LEDs. These are mounted on the containers in the case of the unordered system, and on the storage location in the case of the ordered system. These output displays provide the user with the capability to identify the position of containers within a collection and, in the case of the ordered system, show where containers should be returned. However, there are several other options for expanding this capability to provide other functionality. SOPHYA provides the basic infrastructure that would allow support of more

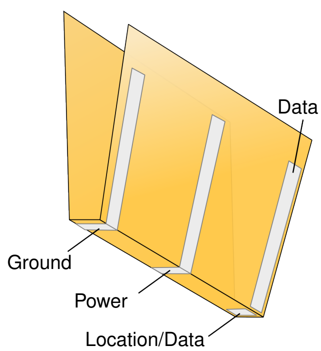

Figure 11.1: Container supporting integration of the unordered and ordered SOPHYA systems.

advanced user interaction with both containers and storage locations. Some of these are discussed below.

Output display on ordered containers

One of the significant differences between the unordered and ordered versions of the SOPHYA prototype was that the output display devices were moved from being part of the containers (in the unordered prototype) to the storage location (in the ordered prototype). It would, however, also be possible to include output displays on the ordered containers, either in addition to or instead of those on the storage location.

Because the system described in Chapter 8 was designed for containers with no user interface devices, it was not necessary to supply the containers with much power. However, since the design of this system requires that no more than one container be present on a container network at a time, each container network controller only has to power a single container, unlike the previous prototypes which had to power an arbitrary number of containers. As such, with some modifications to the design it would be possible to include a small number of low-power user interface devices (e.g., LEDs) on the containers themselves.

For more advanced user interfaces that require more power, a possible approach would be to add a common power rail, as shown in Figure 11.2. Adding this power rail raises some design issues that would need to be overcome in order for this approach to be successful. For example, the system would need to deal with potential short circuits between the power rail and ground, and between the power rail and the location pads. Also, a number of design issues similar to those experienced with previous prototype due to the three points of contact between container and storage location (see Section 7.3.3) would need to be overcome.

Other user interface devices on containers

The prototype systems described so far have included simple output displays in the form of LEDs. This was sufficient because the main focus of this research was the development of a working prototype system that could be evaluated and used as a basis for future developments. However, the implementation of the unordered version of SOPHYA is currently able to support user interface devices with greater power requirements on the containers. Similarly, with the changes described above, such support could also be added to the ordered

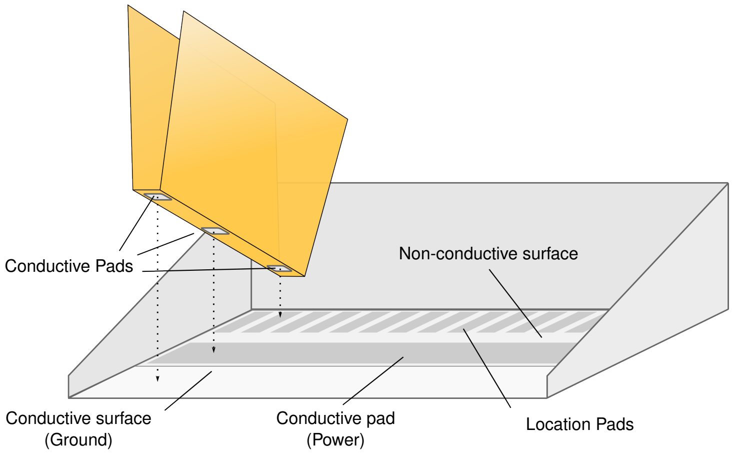

Figure 11.2: Adding a common power rail to support more advanced user interfaces.

version of SOPHYA.

A range of possible interface devices could be added to the containers. Visual display devices include several multi-colour LEDs, and character and graphical LCD screens (although the limitations with implementing such displays would be due to physical rather than electronic constraints). It would also be possible to implement other types of output devices, including audible devices (e.g., beepers or buzzers) and tactile devices (e.g., vibration motors). Non-visual output devices (e.g., beepers) could be useful, for example, when line-of-sight to a container is not available (e.g., when it is obscured in a stack of containers).

Tangible browsing of electronic information

The discussion of user interface elements has so far focused on containers, but other user interface devices could also be embedded into the physical storage location. Figure 11.3 shows an example of a physical storage location with an integrated touchscreen interface that would allow the users to interact with the DMS while at the storage location. This would make it possible for the users to, for example, browse through electronic metadata in order to determine which folder within the storage location their documents are in.

Such an interface could be further extended to support display of electronic information while facilitating tangible browsing through the physical collection. For example, the prototype system discussed in Chapter 6 included a button on each folder to allow such tangible browsing. The button was not implemented in the SOPHYA prototypes because it was not considered to be crucial for the functioning of these prototypes (as discussed in Section 7.4.1),

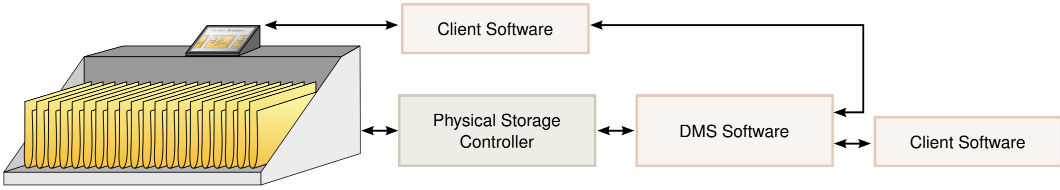

Figure 11.3: Incorporating a more advanced user interface into the storage location that incorporates browsing the electronic content stored by the DMS.

but it may be desirable to include such tangible browsing functionality in future implementations.

One method of implementing an input component for tangible browsing would be to include capacitive touch sensing on the containers so that when a container is touched related information is displayed on an output display screen attached to the storage location. It would even be possible to have separate zones on the containers so that different information can be displayed by touching the container in different places. Another method would be to include a touch-sensitive strip running along the front of the filing cabinet shelf (as shown in Figure 11.4). This would enable the user to scroll through information for each folder in the cabinet. If combined with the ordered SOPHYA system, such an input device would allow the user to display information about individual folders in the cabinet by touching the touch sensitive strip in front of the appropriate container.

Physical access control

Another area in which SOPHYA could prove advantageous is in providing improved access control to physical documents. Although this has not yet been implemented in the current SOPHYA prototypes, it is possible to add the ability for the physical storage controller to control a locking mechanism, which would allow the DMS software to lock and unlock the physical storage location as required. Unlocking could, for example, be linked to the presentation of a user-specific RFID badge at the storage location or, alternatively, an unlock request could be generated by the client software when a user requires a given document as part of their work-flow, as shown in Figure 11.5.

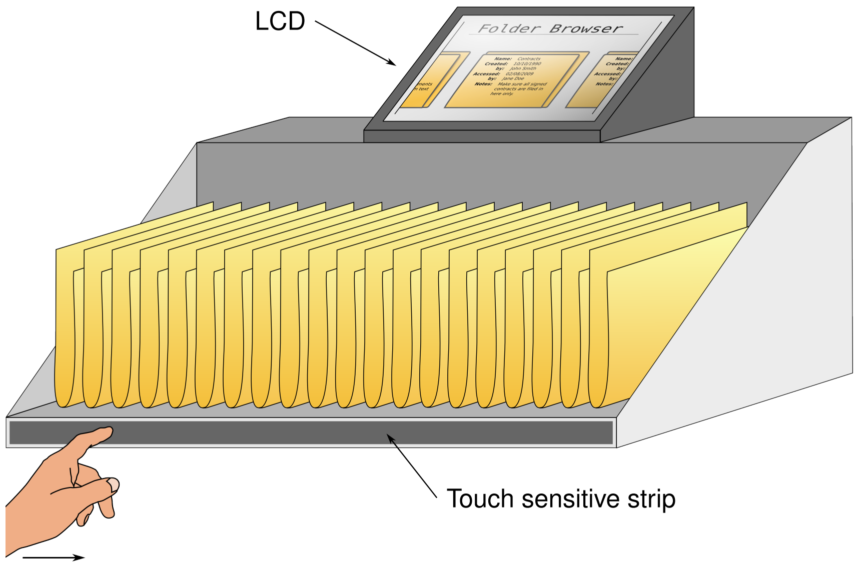

Figure 11.4: Tangible browsing interface, comprising a touch-sensitive strip running along the front of the cabinet shelf and an LCD screen for electronic information display. Users can scroll through information displayed on the screen by running their fingers along the touch-sensitive strip.

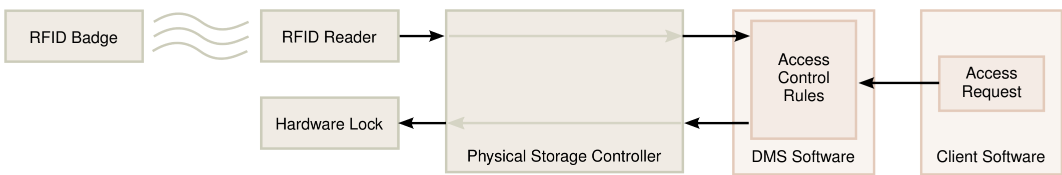

Figure 11.5: Access control to the physical document collection provided by SOPHYA by including a physical locking mechanism that can be electronically controlled by the physical storage controller. Access could be requested, for example, by a user with an RFID badge, or by the client software when a given document is required.

11.3.3 Combining SOPHYA with document tracking

The SOPHYA technology described in this thesis is designed to operate at the document container (e.g., folder) level rather than the individual document level. Such an approach makes sense for offices where the work-flow is at container level (e.g., a container corresponds to a specific job or client). However, in other applications it may be more useful to have tracking at individual document granularity. In such situations containers would still be used to manage collections of related documents (and therefore SOPHYA would still play a part), but some additional technology would also be required to automate the tracking of individual documents within these collections.

The most suitable technology currently available for tracking individual documents is RFID, and several RFID-based document tracking systems were discussed previously in Section 3.2.2. There are a number of ways that the automating of linking between documents and document containers can be approached using RFID. Three possible ways are described below.

RFID readers built into containers. This would enable containers to detect which documents are within them, and communicate this via the physical storage location to the DMS software. This approach may not be feasible at present, however, due to the costs associated with incorporating an RFID reader into every container.

RFID augmented document tray. An alternative method is to use a document tray with RFID tag reading capabilities (e.g., Magellan Technology Pty Ltd, n.d.) to maintain the links between the documents and the containers that they are stored in. When a container of documents is placed into such a document tray, the inbuilt tag reader would read the tags of the documents in the container and notify the DMS software. However, while this approach has the advantage of lower cost, since fewer readers would be required, it has the downside that it requires containers be placed (one at a time) into the document tray each time their contents change, in order for this to recognised by the system.

RFID augmented storage locations. In order to overcome the fact that the augmented document tray method requires an extra step, an alternative option is to incorporate an RFID reader into each physical storage location. This would allow the system to detect which documents are in a container when it is placed into the storage location. However, this approach has the limitation that it requires a type of RFID technology that can cope with a large number of tags being in range of the RFID reader. Furthermore, it also requires the user to place containers into the storage location one at a time, leaving enough time between each to enable the system to determine the contents of each container.

11.3.4 Tracking containers outside of storage locations

Although the choice of a wired connection to supply power to, and communicate with, containers has several benefits, such as enabling the container circuitry to draw more power than, for example, using RFID, a limitation that this imposes is that it only functions when there is a physical connection between a given container and a storage location. There are situations, however, where it would be beneficial to track containers when they are not in a storage location. One method of achieving such tracking would be to combine SOPHYA technology with RFID technology, such that SOPHYA would provide the interaction when containers are within storage locations, and RFID would provide tracking of containers outside storage locations. This would require containers to be tagged with RFID tags, which has been discussed previously in Section 8.3.6, and would require some sort of infrastructure to be provided for tracking the containers. This may be achieved, for example, using existing RFID location determination technology such as that described in Section 3.2.2.

11.3.5 Designing for retrofitting

One of the requirements specified in Section 4.4.1 was that the design of the system should take into account retrofitting existing collections with such technology. Although this is largely an industrial design issue, the current design of SOPHYA already gives some consideration to retrofitting. The electronic

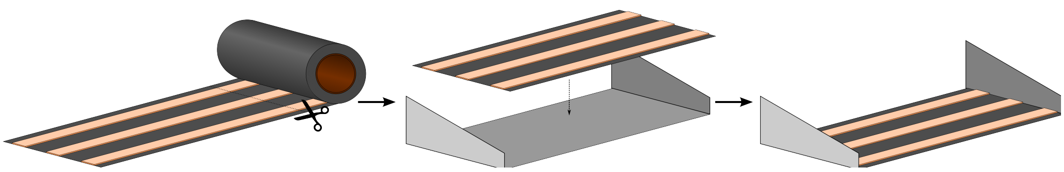

Figure 11.6: Retrofitting storage locations with unordered SOPHYA technology.

hardware of the ordered version of SOPHYA is split into modules connected together in a bus topology. Thus, an arbitrary number of modules can be used per storage location. The modules could be produced with an adhesive backing so that they could be stuck into storage locations. Similarly, a potential method of retrofitting storage locations with the unordered version of SOPHYA would be to buy the conductive material that forms the basis for communication in a roll, this could then be cut to length and attached to storage location along with a physical storage controller (as shown in Figure 11.6).

11.3.6 Longitudinal study of SOPHYA

The evaluation of SOPHYA described in Chapter 10 identified that most of the participants, who work in offices, felt that there would be potential benefit from SOPHYA. However, as laboratory based evaluations are not always able to capture all the subtleties of the real world, in order to better guide future development of SOPHYA and similar systems, and further identify its benefits and weaknesses, the next logical step would be to conduct a longitudinal study of its use in a real-world office.

Conducting a longitudinal study has been beyond the scope of this thesis. Preparation for such a study would include identifying a suitable office to deploy the SOPHYA systems in, constructing a sufficient quantity of SOPHYA containers and storage locations, and developing software to integrate SOPHYA with the existing electronic DMS in the office.

11.4 Summary

This thesis showed that there is need for better integration between paper and electronic document management in offices of today, and demonstrated an approach to developing such systems that implement tangible/embodied user interfaces to document management. In the course of the research conducted for this thesis, several prototype systems were constructed and these were described in the thesis. These systems were evaluated with encouraging results. This thesis has demonstrated that it is possible to develop such tangible/embodied user interfaces for integrating paper and electronic document management, and described one process by which this could be achieved.

Part V

References and Appendices

References

(1975). The office of the future. *BusinessWeek*.

AbuSafiya, M., Mazumdar, S. (2004). Accommodating paper in document databases. In *Proceedings of the 2004 ACM symposium on Document engineering*, pp. 155–162. New York, NY, USA: ACM. doi:10.1145/1030397.1030428.

AIM Inc. (2001). Radio frequency identification: A basic primer. White paper WP-98/002R2, The Association of the Automatic Identification and Data Capture Industry.

Arregui, D., Fernström, C., Pacull, F., Rondeau, G., Willamowski, J., Crochon, E., Favre-Reguillon, F. (2003). Paper-based communicating objects in the future office. In *SOC'03 Smart Object Conference*.

Azuma, R. T. (1997). A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, *6*(4), 355–385.

Baarman, D. W., Schwannecke, J. (2009). Understanding wireless power. Whitepaper, Fulton Innovation LLC.

Baeza-Yates, R., Ribeiro-Neto, B. (1999). *Modern information retrieval*. Addison Wesley.

Bakshi, K., Karger, D. R. (2005). End-user application development for the semantic web. In *Proceedings of the 1st ISWC workshop on the semantic desktop*, pp. 123–137.

Barreau, D. K. (1995). Context as a factor in personal information management systems. *Journal of the American Society for Information Science*, *46*(5), 327–339. doi:10.1002/(SICI)1097-4571(199506)46:5<327::AID-ASI4>3.0.CO;2-C.

Blair, D. C., Gordon, M. D. (1991). The management and control of written information: growing concern amid the failure of traditional methods. *Information & Management*, *20*, 239–246. doi:10.1016/0378-7206(91)90016-U.

Boardman, R. (2004). *Improving tool support for personal information management.* Ph.D. thesis, Imperial College London, London, UK.

Bodhuin, T., Preziosi, R., Tortorella, M. (2007). Supporting document management by using RFID technology. *International Journal of Internet Protocol Technology*, *2*(3/4), 165. doi:10.1504/IJIPT.2007.016218.

Boyette, N., Krishna, V., Srinivasan, S. (2005). Eclipse modeling framework for document management. In *Proceedings of the 2005 ACM symposium on Document engineering*, pp. 220–222. New York, NY, USA: ACM. doi: 10.1145/1096601.1096653.

Buxton, W. (2002). *The invisible future*, chap. Less Is More (More or Less), pp. 145–179. McGraw-Hill, Inc.

Câmara, L., Guerreiro, T., Gonçalves, D., Jorge, J. A. (2008). RealFind: managing personal items in the physical world. In *CHI '08 extended abstracts on Human factors in computing systems*, pp. 3321–3326. New York, NY, USA: ACM. doi:10.1145/1358628.1358851.

Case, D. O. (1986). Collection and organization of written information by social scientists and humanists: a review and exploratory study. *Journal of Information Science*, *12*, 97–104. doi:10.1177/016555158601200302.

Case, D. O. (1991). Conceptual organization and retrieval of text by historians: The role of memory and metaphor. *Journal of the American Society for Information Science*, *42*(9), 657–668. doi:10.1002/(SICI)1097-4571(199110)42:9<657::AID-ASI4>3.0.CO;2-7.

Choi, J., Oh, D., Song, I. (2006). R-lim: an affordable library search system based on rfid. In *Proceedings of the 2006 International Conference on Hybrid Information Technology*, ICHIT '06, pp. 103–108. Washington, DC, USA: IEEE Computer Society. doi:10.1109/ICHIT.2006.197.

Cole, I. (1982). Human aspects of office filing: implications for the electronic office. In *Proceedings of the Human Factors Society 26th Annual Meeting*, pp. 59–63.

Dallas Semiconductor/Maxim (2006a). DS2401: Silicon Serial Number. http://datasheets.maxim-ic.com/en/ds/DS2401.pdf.

Dallas Semiconductor/Maxim (2006b). DS2450: 1-Wire Quad A/D Converter. http://datasheets.maxim-ic.com/en/ds/DS2450.pdf.

Dallas Semiconductor/Maxim (2009a). DS2406: Dual Addressable Switch Plus 1Kb Memory. http://pdfserv.maxim-ic.com/en/ds/DS2406.pdf.

Dallas Semiconductor/Maxim (2009b). DS9490R/DS9490B: USB to 1-Wire/iButton Adapter. http://datasheets.maxim-ic.com/en/ds/DS9490-DS9490R.pdf.

de Ipiña, D. L., Mendonça, P. R. S., Hopper, A. (2002). Trip: A low-cost vision-based location system for ubiquitous computing. *Personal and Ubiquitous Computing*, *6*(3), 206–219. doi:10.1007/s007790200020.

Decker, C., Beigl, M., Eames, A., Kubach, U. (2004). DigiClip: Activating physical documents. In *Proceedings of the 24th International Conference on Distributed Computing Systems Workshops*, pp. 388–393. Washington, DC, USA: IEEE Computer Society.

Dey, I. (1993). *Qualitative data analysis*. London, UK: Routledge.

Dourish, P., Edwards, W. K., LaMarca, A., Lamping, J., Petersen, K., Salisbury, M., Terry, D. B., Thornton, J. (2000). Extending document management systems with user-specific active properties. *ACM Transactions on Information Systems*, *18*, 140–170. doi:10.1145/348751.348758.

Dumais, S. T., Landauer, T. K. (1983). Using examples to describe categories. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, vol. December, pp. 112–115. New York, NY, USA: ACM. doi:10.1145/800045.801592.

Feather, J., Sturges, P. (2003). *International Encyclopedia of Information and Library Science*. London, UK: Routledge (Taylor & Francis Group), 2nd edn.

Fiala, M. (2005). ARTag, a fiducial marker system using digital techniques. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 590–596. Los Alamitos, CA, USA: IEEE Computer Society. doi:10.1109/CVPR.2005.74.

Fishkin, K. P., Gujar, A., Harrison, B. L., Moran, T. P., Want, R. (2000). Embodied user interfaces for really direct manipulation. *Communications of the ACM*, *43*(9), 75–80. doi:10.1145/348941.348998.

Fishkin, K. P., Moran, T. P., Harrison, B. L. (1998). Embodied user interfaces: Towards invisible user interfaces. In *Proceedings of EHCI '98*. Deventer, The Netherlands, The Netherlands: Kluwer, B.V.

Fitzmaurice, G. W., Ishii, H., Buxton, W. A. S. (1995). Bricks: Laying the foundations for graspable user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, May, pp. 442–449. New York, NY, USA: ACM. doi:10.1145/223904.223964.

Frohlich, D., Perry, M. (1994). The paperful office paradox. Tech. Rep. HPL-94-20, HP Laboratories Bristol, Bristol, UK.

Golovchinsky, G. (2008). Reading in the office. In *Proceeding of the 2008 ACM workshop on Research advances in large digital book repositories*, pp. 21–24. New York, NY, USA: ACM. doi:10.1145/1458412.1458420.

Gordon, M. D. (1997). It's 10 A.M. do you know where your documents are? the nature and scope of information retrieval problems in business. *Information Processing and Management*, *33*(1), 107–121. doi: 10.1016/S0306-4573(96)00021-0.

Gorman, G. E., Clayton, P. (2005). *Qualitative Research for the Information Professional*. London, UK: Facet Publishing, 2nd edn.

Gould, J. D., Lewis, C. (1985). Designing for usability: key principles and what designers think. *Communications of the ACM*, *28*(3), 300–311. doi: 10.1145/3166.3170.

Gumstix, Inc. (2007). Gumstix, Inc. introduces faster, more versatile verdex™ product line–doubles memory and increases speed by 50%. http://www.gumstix.com/press/verdex-launch-033007.pdf.

Haller, M., Billinghurst, M., Leithinger, D., Leitner, J., Seifried, T. (2005). Coeno—enhancing face-to-face collaboration. In *Proceedings of the 15th International Conference on Artificial Reality and Telexistence*, pp. 40–47. New York, NY, USA: ACM. doi:10.1145/1152399.1152408.

Haller, M., Leithinger, D., Leitner, J., Seifried, T., Brandl, P., Zauner, J., Billinghurst, M. (2006). The shared design space. In *ACM SIGGRAPH 2006 Emerging technologies*, SIGGRAPH '06. New York, NY, USA: ACM. doi:10.1145/1179133.1179163.

Hark, B. S., Choi, J. H., Leem, C. S. (2008). A database design of RFID document management system with e-Ink technology. In *Proceedings of the Fourth International Conference on Networked Computing and Advanced Information Management*, pp. 14–17. Washington, DC, USA: IEEE Computer Society. doi:10.1109/NCM.2008.155.

Hekimian-Williams, C., Grant, B., Kumar, P. (2010). Accurate localization of rfid tags using phase difference. In *Procedings of the 2010 IEEE International Conference on RFID*, pp. 89–96. Orlando, FL, USA: IEEE Computer Society. doi:10.1109/RFID.2010.5467268.

Henderson, S. (2009). *How Do People Manage Their Documents?* Ph.D. thesis, The University of Auckland, Auckland, NZ.

Heo, J., Pyeon, M. W., Kim, J. W., Sohn, H. (2007). Towards the optimal design of an RFID-based positioning system for the ubiquitous computing environment. In *Proceedings of the 2nd international conference on Rough Sets and Knowledge Technology*, RSKT'07, pp. 331–338. Berlin, Heidelberg, Germany: Springer-Verlag.

Hinske, S., Langheinrich, M. (2008). Using a movable rfid antenna to automatically determine the position and orientation of objects on a tabletop. In *Smart Sensing and Context*, Lecture Notes in Computer Science, pp. 14–26. Heidelberg, Germany: Springer-Verlag. ISBN 978-3-540-88792-8. doi:10.1007/978-3-540-88793-5_2.

Ishii, H., Ullmer, B. (1997). Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 234–241. New York, NY, USA: ACM. doi:10.1145/258549.258715.

Jahoda, G., Hutchins, R., Galford, R. (1966). Characteristics and use of personal indexes maintained by scientists and engineers in one university. *American Documentation*, *17*(2), 71–75. doi:10.1002/asi.5090170204.

Jervis, M., Masoodian, M. (2009). Digital management and retrieval of physical documents. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, pp. 47–54. New York, NY, USA: ACM. doi:10.1145/1517664.1517682.

Jervis, M., Masoodian, M. (2010a). Design of a modular architecture for integrating paper and digital document management systems. In *Workshop Proceedings of PaperComp 2010, The 1st International Workshop on Paper Computing, Part of the 12th ACM International Conference on Ubiquitous Computing*, pp. 44–47. New York, NY, USA: ACM.

Jervis, M., Masoodian, M. (2011). Evaluation of an integrated paper and digital document management system. In *Proceedings of INTERACT 2011, The 13th IFIP TC13 Conference on Human-Computer Interaction*. Springer-Verlag. In print.

Jervis, M. G., Masoodian, M. (2010b). SOPHYA: a system for digital management of ordered physical document collections. In *Proceedings of the fourth international conference on Tangible, Embedded, and Embodied Interaction*, TEI '10, pp. 33–40. New York, NY, USA: ACM. doi:10.1145/1709886.1709894.

Jones, W. (2004). Finders, keepers? the present and future perfect in support of personal information management. *First Monday*, *9*(3).

Jones, W. (2007). Personal information management. *Annual Review of Information Science and Technology*, *41*(1), 453–504. doi:10.1002/aris.2007.1440410117.

Kampffmeyer, U. (2006). ECM: Enterprise content management. White paper, Project Consult, Hamburg, Germany.

Karger, D. R. (2007). Unify everything:it's all the same to me. In Jones, W., Teevan, J. (Eds.), *Personal Information Management*, chap. 8, pp. 127–152. Seattle, WA, USA: University of Washington Press.

Kranz, M., Schmidt, A., Maldonado, A., Rusu, R. B., Beetz, M., Hörnler, B., Rigoll, G. (2007). Context-aware kitchen utilities. In *Proceedings of the 1st international conference on Tangible and Embedded Interaction*, pp. 213–214. New York, NY, USA: ACM. doi:10.1145/1226969.1227013.

Kuon, I., Tessier, R., Rose, J. (2008). FPGA architecture: Survey and challenges. *Foundations and Trends in Electronic Design Automation*, *2*(2), 135–253. doi:10.1561/1000000005.

Lansdale, M. (1988). The psychology of personal information management. *Applied Ergonomics*, *19*(1), 55–66. doi:10.1016/0003-6870(88)90199-8.

Lawrie, D. (1997). Self-organizing file cabinet. Senior Honors Thesis PCS-TR97-313, Dartmouth College.

Lawrie, D., Rus, D. (1999). A self-organized file cabinet. In *Proceedings of the eighth international conference on Information and knowledge management*, pp. 499–506. New York, NY, USA: ACM. doi:10.1145/319950.320056.

Liesaputra, V., Witten, I. H., Bainbridge, D. (2009). Creating and reading realistic electronic books. *Computer*, *42*(2), 72–81. doi:10.1109/MC.2009.46.

Linke, B. (2008). Overview of 1-Wire® Technology and Its Use. http://pdfserv.maxim-ic.com/en/an/AN1796.pdf.

Magellan Technology Pty Ltd (n.d.). Document tray: MDOCR-2505. http://live.magellan-technology.com.tmp.anchor.net.au/res/PDF/Brochures%20and%20Documents/2505%20PB.pdf.

Malone, T. W. (1983). How do people organize their desks?: Implications for the design of office information systems. *ACM Transactions on Information Systems*, *1*(1), 99–112. doi:10.1145/357423.357430.

McKay, D., Conyers, B. (2010). Where the streets have no name: how library users get lost in the stacks. In *Proceedings of the 11th International Conference of the NZ Chapter of the ACM Special Interest Group on Human-Computer Interaction*, pp. 77–80. New York, NY, USA: ACM. doi:10.1145/1832838.1832852.

Meunier, J., Grasso, A. (2002). Method for indexing and retrieval of physical documents. United States Patent US 6,499,665 B1, Xerox Corporation.

Microhip Technology, Inc. (2002). PIC18 Microcontroller Family. http://ww1.microchip.com/downloads/en/DeviceDoc/30327b.pdf.

Minack, E., Paiu, R., Costache, S., Demartini, G., Gaugaz, J., Ioannou, E., Chirita, P.-A., Nejdl, W. (2010). Leveraging personal metadata for desktop search: The beagle++ system. *Web Semantics: Science, Services and Agents on the World Wide Web*, *8*(1), 37–54. doi:10.1016/j.websem.2009.12.001.

Nardi, B., Anderson, K., Erickson, T. (1994). Filing and finding computer files. Tech. rep., Apple Computer, Inc., Cupertino, CA, USA.

Newman, S. (2004). "The tobacco case" document management by professional service providers: the ramifications of McCabe v British American Tobacco Australia Limited. *Heart, Lung & Circulation*, *13 Supplement 3*, S43–S47. doi:10.1016/j.hlc.2004.09.019.

Norman, D. (1990). *The Design of Everyday Things*. Doubleday Business.

Norrie, M. C., Signer, B., Weibel, N. (2006). Print-n-link: weaving the paper web. In *Proceedings of the 2006 ACM symposium on Document engineering*, pp. 34–43. New York, NY, USA: ACM. doi:10.1145/1166160.1166175.

NXP Semiconductor (2007). UM10204: I^2C-bus specification and user manual. http://www.nxp.com/documents/user_manual/UM10204.pdf.

O'Neill, J., Grasso, A., Willamowski, J. (2006). RFID: enhancing paper documents with electronic properties. In *Computer Supported Cooperative Work, Workshop Collaborating Over Paper and Digital Documents*.

Pederson, T. (2001). Magic Touch: A simple object location tracking system enabling the development of physical-virtual artefacts in office environments. *Personal and Ubiquitous Computing*, *5*(1), 54–57.

Porter-Roth, B. (2006). Applying electronic records management in the document management environment: An integrated approach. White Paper May.

Raskar, R., Beardsley, P., van Baar, J., Wang, Y., Dietz, P., Lee, J., Leigh, D., Willwacher, T. (2004). RFIG lamps: interacting with a self-describing world via photosensing wireless tags and projectors. *ACM Transactions on Graphics*, *23*(3), 406–415. doi:10.1145/1015706.1015738.

Raskar, R., Beardsley, P., Dietz, P., van Baar, J. (2005). Photosensing wireless tags for geometric procedures. *Communications of the ACM*, *48*(9), 46–51. doi:10.1145/1081992.1082019.

Rosner, D. K., Oehlberg, L., Ryokai, K. (2008). Studying paper use to inform the design of personal and portable technology. In *CHI '08 extended abstracts on Human factors in computing systems*, pp. 3405–3410. New York, NY, USA: ACM. doi:10.1145/1358628.1358865.

Saxena, A., Ganguly, S., Bhatnagar, S., Izmailov, R. (2007). Rfind: An rfid-based system to manage virtual spaces. In *Proceedings of the Fifth IEEE International Conference on Pervasive Computing and Communications Workshops*, pp. 382–387. Washington, DC, USA: IEEE Computer Society. doi: 10.1109/PERCOMW.2007.99.

Seifried, T. (2007). *Aroo: Integration von realer und virtueller Dokumentenverwaltung.* Master's thesis, Upper Austria University of Applied Sciences.

Seifried, T., Jervis, M., Haller, M., Masoodian, M., Villar, N. (2008). Integration of virtual and real document organization. In *Proceedings of the 2nd international conference on Tangible and embedded interaction*, pp. 81–88. New York, NY, USA: ACM. doi:10.1145/1347390.1347410.

Sellen, A. J., Harper, R. H. R. (2003). *The Myth of the Paperless Office.* Cambridge, MA, USA: MIT Press.

Silberman, S. (2001). The hot new medium: Paper. *Wired Magazine*, *9.04*, 184–191.

Silva, R., Oliveira, J. C., Giraldi, G. A. (2003). Introduction to augmented reality. Tech. rep., National Laboratory for Scientific Computation, Petrópolis, RJ, Brazil.

Smith, J., Long, J., Lung, T., Anwar, M. M., Subramanian, S. (2006). Paperspace: a system for managing digital and paper documents. In *CHI '06 extended abstracts on Human factors in computing systems*, pp. 1343–1348. New York, NY, USA: ACM. doi:10.1145/1125451.1125700.

Stainback, S., Stainback, W. (1988). *Understanding & Conducting Qualitative Research.* The Council for Exceptional Children.

Stanfield, J. D., Branc, J. R., Feldpausch, T. G., Miller, W. L. (2007). Electronic system, components and method for tracking files. United States Patent US 7,301,470 B2, Steelcase Development Corporation, Caledonia, MI, USA.

Strong, K. V. (2008). Enterprise content and records management. *Journal of AHIMA*, *79*(2), 38–44.

Teevan, J., Alvarado, C., Ackerman, M. S., Karger, D. R. (2004). The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '04, pp. 415–422. New York, NY, USA: ACM. doi:10.1145/985692.985745.

The Media Interaction Lab (2007). Aroo picture pack. http://mi-lab.org/files/2007/11/aroo-picture-pack.zip.

Ullmer, B., Ishii, H., Glas, D. (1998). mediaBlocks: physical containers, transports, and controls for online media. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pp. 379–386. New York, NY, USA: ACM. doi:10.1145/280814.280940.

Villar, N. (2007). *Flexible Physical Interfaces*. Ph.D. thesis, Lancaster University.

Villar, N., Gellersen, H. (2007). A malleable control structure for softwired user interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, pp. 49–56. New York, NY, USA: ACM. doi:10.1145/1226969.1226980.

Want, R. (2006). An introduction to rfid technology. *IEEE Pervasive Computing*, *5*(1), 25–33. doi:10.1109/MPRV.2006.2.

Weibel, N., Ispas, A., Signer, B., Norrie, M. C. (2008). Paperproof: a paper-digital proof-editing system. In *CHI '08 extended abstracts on Human factors in computing systems*, pp. 2349–2354. New York, NY, USA: ACM. doi:10.1145/1358628.1358682.

Weiser, M. (1991). The computer for the 21st century. *Scientific American*, *265*(3), 94–104. doi:10.1038/scientificamerican0991-94.

Whittaker, S., Hirschberg, J. (2001). The character, value, and management of personal paper archives. *ACM Transactions on Computer-Human Interaction*, *8*(2), 150–170. doi:10.1145/376929.376932.

Whittaker, S., Sidner, C. (1996). Email overload: exploring personal information management of email. In *Proceedings of the SIGCHI conference on Human factors in computing systems: common ground*, pp. 276–283. New York, NY, USA: ACM. doi:10.1145/238386.238530.

Wild, P. J., McMahon, C., Darlington, M., Liu, S., Culley, S. (2010). A diary study of information needs and document usage in the engineering domain. *Design Studies*, *31*(1), 46–73. doi:10.1016/j.destud.2009.06.002.

Wilson, P. (2001). 20 years in the life of a long-term empirical personal electronic empirical personal filing study. *Behaviour & Information Technology*, *20*(5), 395–409. doi:10.1080/0144929013171884999.

Wilson, T. D. (2003). *International Encyclopedia of Information and Library Science*, chap. Definition of information management. Routledge (Taylor & Francis Group), 2nd edn.

Wood, F. (1986). The history of use of paper in china. http://www.sacu.org/paper.html.

Zhao, Y., Cheung, S., Ni, L. M. (2010). LocaToR: Locating passive rfid tags with the relative neighborhood graph. In *IEEE/IFIP International Conference on Embedded and Ubiquitous Computing*, pp. 154–161. Los Alamitos, CA, USA: IEEE Computer Society. doi:10.1109/EUC.2010.31.

Appendix A

Ethics Approval for Observational Study

This appendix includes the letter of approval from the ethics committee of the Faculty of Computing and Mathematical Sciences, University of Waikato for the study described in Chapter 4.

School of Computing &
Mathematical Sciences
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4021
www.scms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

23rd September 2008

Matthew Jervois
C/- Department of Computer Science
University of Waikato

Dear Matthew

Request for approval to perform several studies involving human participants.

I have considered your request for approval to perform several studies involving human participants from September 2008 to February 2009.

The purpose of these studies is to gain an understanding of how office workers use their digital and physical workspaces by studying the way they work in their office environment.

The procedure described in your request is acceptable. I note your statements that confidentiality and participant anonymity will be strictly maintained, all information gathered will be used for statistical analysis only. No names or other identifying characteristics will be stated in the final or any other reports and collected data (e.g. video and audio recordings and transcripts) will be stored securely in the office of Masoon Masoodian at the University of Waikato (G1.08) for the duration of your PhD research (expected completion date: June 2010) after which time it will be destroyed.

The research participants' Bill of Rights and the Research Consent form comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to undertake the experiment.

Yours faithfully

Mike Mayo
Department of Computer Science
Human Research Ethics Committee
School of Computing and Mathematical Sciences

Appendix B

SOPHYA Container Network Protocol

This chapter describes the communication protocol developed for the container networks of the SOPHYA prototype described in Chapter 7. The design of the protocol can be broken down into three layers, which are (from lowest to highest):

Signalling. Specifies the requirements for the signal and the physical elements that generate and receive it.

Framing. Defines the framing of the signal into packets.

Command. Defines the commands that can be address from master to slaves and the responses that can be received.

Each of these layers is described in more detail in the following sections.

B.1 Signalling layer

The SOPHYA container network implements a bus topology with a single master (the container network controller in the physical storage location) and multiple slaves (the containers), as shown in Figure B.1. The communication protocol requires three lines for communication: power, data, and ground. By convention the these lines are laid out in the following order (from the front of the container to the back): data, ground, power. Appendix C includes the circuitry for the bus master (Section C.1.2) and slaves (Section C.1.1).

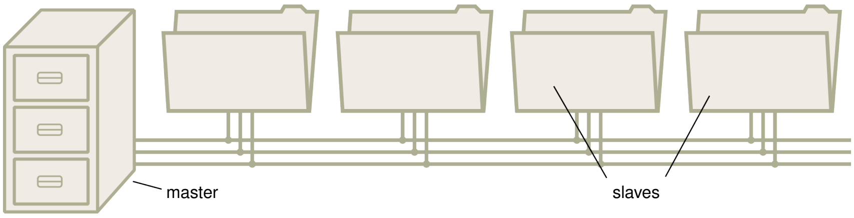

Figure B.1: The containers on a container network form a bus topology.

B.1.1 Bus master circuitry

The bus master is responsible for providing power to the slaves through the power line. The power line is also used for start-of-frame signalling, at which time it is driven low by the master (as described in Section B.2). The design of the driver circuitry for the power line is shown in Figure B.2a. When the `Drive High` input is active, `Q1` turns on allowing current to flow through `R1`. `R1` is a low value resistor (e.g., 10Ω) whose purpose is to limit the current if the power line is shorted to ground. When `Drive High` is inactive, the bus master can drive the power line low by activating `Q2` with the `Drive Low` input. When both inputs are inactive the power line is pulled high by weak pull-up resistor `R2`.

It is expected that the power line driver be able handle a short circuit to ground. Therefore, while the power line control logic is driving the line high it will periodically release the power line and sample the `Status` input to detect whether there is a short circuit condition (i.e., `Status` is low). If no short circuit is detected, then the control logic will again drive the power line high and the cycle will continue to repeat. However, if there is a short circuit the

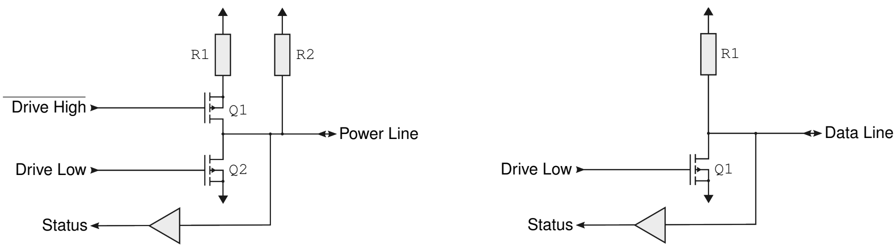

(a) Master power line driver. **(b)** Master data line driver.

Figure B.2: Conceptual diagrams of the signalling circuitry of the master. Note that this is a simplified view showing the concept and does not include some supporting circuitry such as pull-up/-down resistors. For full circuit diagrams see Appendix C.

control logic will continue to poll the status line until the short circuit is no longer detected, at which point it will resume normal driving high operation. When a short circuit is detected the framing layer of the bus master is informed so that it can abort the current transmission and await clearance of the short circuit condition before resuming normal operation.

The data line driver (Figure B.2b) follows a similar design to that of the power line. The difference between the two is that the data line driver lacks the circuitry to drive the line high.

B.1.2 Slave circuitry

The bus circuitry of the slave devices is shown in Figure B.3. This circuit is repeated twice for each slave: once for the power line and once for the data line. It has three main elements: input, output driver, and power diode. Input to the slave logic device is via a resistor (`R1`). The output driver of the slave is used to drive the bus line low. When the `Out` signal is active, `Q1` will drive the bus line low. The power diode `D1` draws power from the bus line. A capacitor (not shown) stores power so that the device can still function for a short period of time when the power line is low.

B.2 Framing layer

The SOPHYA container network protocol uses a standard packet format for all exchanges between master and slaves, which is shown in Figure B.4. All communications are initiated by the master. They begin with the master pulling both the data and power lines low. While holding the power line low

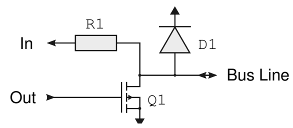

Figure B.3: Conceptual diagrams of the signalling circuitry of the slave.

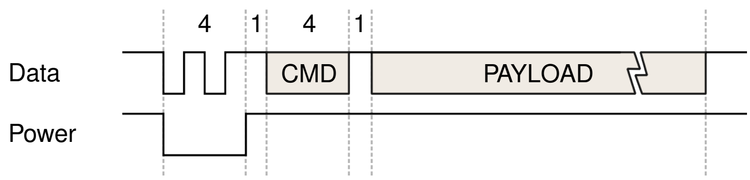

Figure B.4: Framing of standard packet.

the master will then toggle the data line for a total of four bit times. Since the only time the power line is driven low by the master is during the packet preamble, and since the master will abort any communication when a short circuit on the power line is detected, this preamble provides a clear indication to slave devices when a packet is beginning, regardless of the previous state of the bus.

After the preamble, the master drives the power line high again, and waits for one bit-time in this state. It then transmits the command code, which is a four bit code that instructs the slave what to expect next (as described in the next section). The command code is transmitted Most Significant Bit (MSB) first, and a logic high is transmitted by letting the data line float high, while a logic low is transmitted by pulling the data line low.

B.3 Command layer

The protocol currently provides commands that enable it to poll the bus in order to read the IDs of present containers, and control the LEDs on the containers.

The currently defined commands, and their command codes are listed in Table B.1. The command code `0x0` is reserved as invalid. Therefore, if the data line is shorted it cannot be determined to be a command. The remaining codes that have been assigned are classified into three groups: polling, memory, and LED control. The polling and LED control commands are described below. The memory commands are currently reserved, but have not been yet implemented since they have not been needed for the current prototypes. Once these are implemented, further control of the slave devices can be achieved through memory mapped registers.

B.3.1 Polling commands

This protocol employs a novel method for identifying the devices currently connected to the network. This is an important element of the protocol because the system relies on the knowledge of which devices are present on which networks in order to determine the location of containers and where to route commands to. In order to give every device a unique ID the address space for the IDs needs to be large, and it is therefore not possible to iterate through

Table B.1: Command codes that have been defined.

Code	Description
`0x0`	Invalid command
	Poll commands
`0x1`	Poll
`0x7`	Initiate poll cycle
	Memory commands
`0x2`	Write (master → slave)
`0x3`	Read (master ← slave)
	LED commands
`0x4`	Set LEDs
`0x5`	Clear LEDs

each address individually querying if that device is present. Instead a more intelligent search algorithm is needed. The method that has been implemented is described below.

The bus is polled in cycles. Each cycle begins with an *initiate poll cycle* command. This instructs all containers present on the bus that a polling cycle is beginning and resets all slaves to active state. During the poll cycle, the master transmits a *Poll* command. The poll command initiates a single poll transaction, which reads the ID of a single slave. The poll transaction instructs all active containers to begin outputting their IDs to the bus—one bit at a time, MSB first. For each bit of the container ID, the containers first output this bit on the bus (if this bit of their ID is a `0` they drive the bus low, otherwise they let it float high). They then wait one bit time before sampling the bus. Because the bus is wired-AND, if any container drives the bus low, all the others will read back low. For each container, if the sampled bit does not match its ID bit (this can only happen if its ID has a `1` at the given position, but another container on the bus has a `0`), then it will return to an idle state until the the next poll command (at which point it will become active again). The remaining containers will wait for a bit time before repeating with the next bit of their ID.

Eventually, the end of the ID will be reached and only one slave will remain active. This active container then writes a 16 bit Cyclic Redundancy Check (CRC) calculated from its ID so that the bus master can ensure that it has read the ID correctly. Once this is completed, a flag is set in the container's memory which instructs it to enter an invisible state, in which it will not respond to further *poll* commands until a new polling cycle begins. Thus, the master can repeat the poll commands until no slave responds, in which case an ID that is all 1s will be read (this is a reserved ID that can never belong to a real container). Once the master reaches the end of the cycle it transmits the *initiate poll cycle* command, and starts over again.

B.3.2 LED commands

Two commands have been defined for controlling the LEDs on containers. These are the *set LEDs* command, which sets the LEDs of a specific container to a given state, and the *clear LEDs* command, which turns off the LEDs of every container on the bus.

The payload of the set LEDs command begins with the ID of the container that the command is targeted at. All slaves read this ID, but only the slave whose ID matches actually acts on the command. The command is followed first by an eight bit value that is to be written into the LED register of the appropriate container, and then the complement of this. The complement value acts as a safe-guard to ensure that an incorrect value is not written into the LED register. The definition of the LED register is given in Table B.2

Table B.2: LED register definition.

7–3	2	1	0
Unused	Blue	Green	Red

The payload clear LEDs command is `0xB155`. All slaves that receive this command with the correct payload must reset their LEDs register to `0x00`.

Appendix C

SOPHYA Circuitry

This chapter provides the schematics and, where appropriate, PCB layouts for the most recent implementations of the two SOPHYA systems.

Note that this is not a complete documentation of all the hardware that was developed for this thesis. This is only the most recent revision of each circuit. Also, some of the more generic circuitry has been left out.

C.1 Unordered System

This section documents the circuitry developed for the unordered SOPHYA prototype, which is described in Chapter 7.

C.1.1 Container circuitry

Schematic diagram of SOPHYA container circuitry.

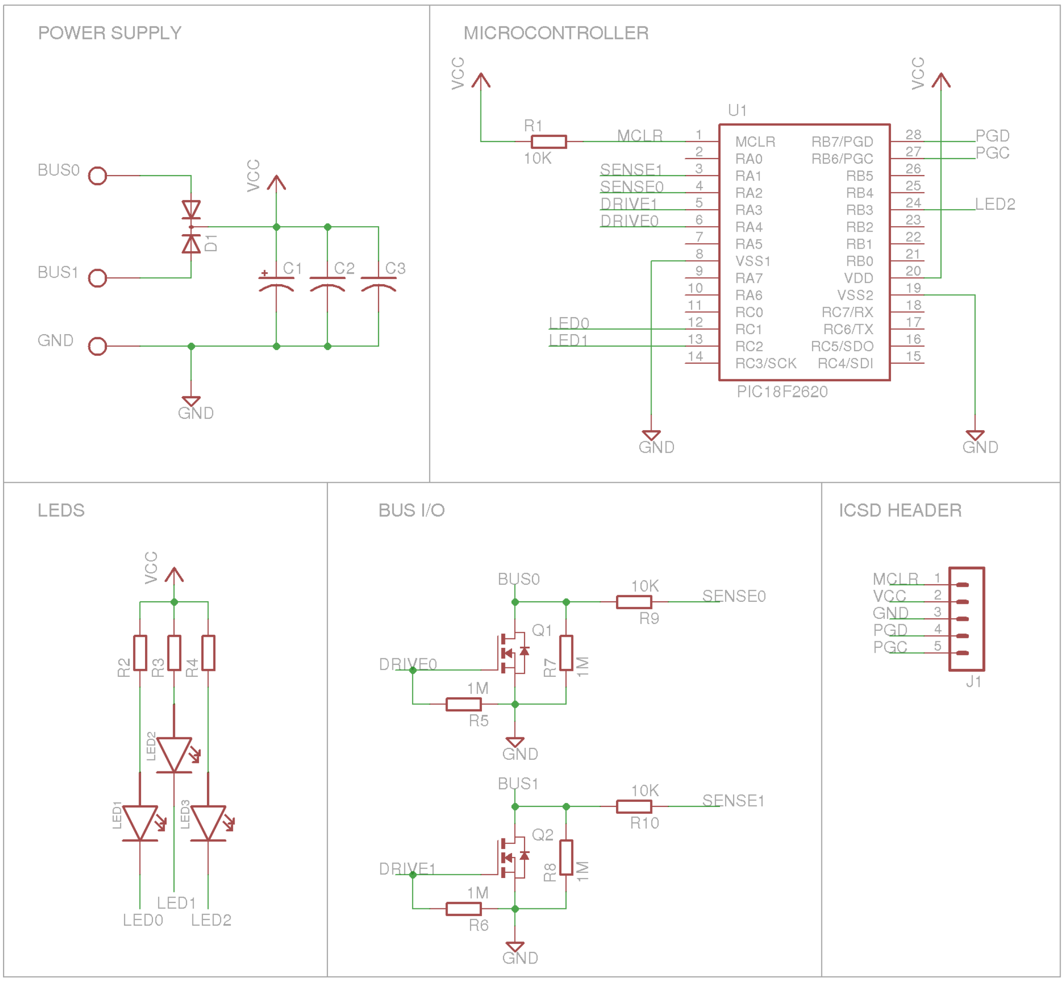

Layout of SOPHYA container circuitry (actual width = 57.9mm).

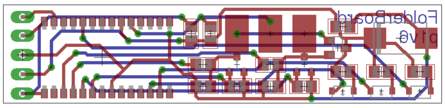

C.1.2 Physical Storage controller

The schematic diagram of the bus driver element of the container network controller circuitry is shown below. This was connected to an FPGA which implemented the control logic of the container network controller.

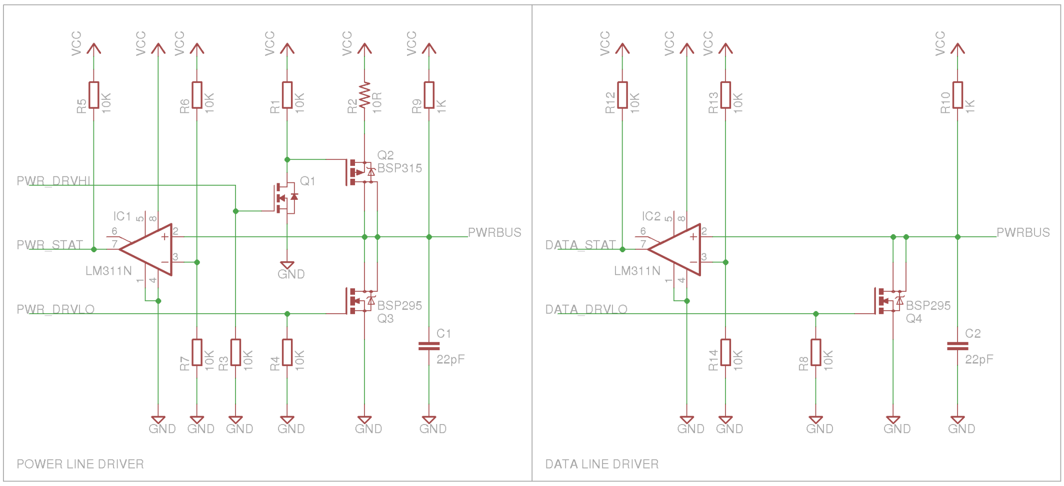

C.2 Ordered System

This section includes the schematics and PCB layouts of the circuitry developed for the ordered SOPHYA prototype, which is described in Chapter 8.

C.2.1 Container circuitry

Schematic diagram of ordered SOPHYA container circuitry.

Layout of ordered SOPHYA container PCB (actual width = 52mm).

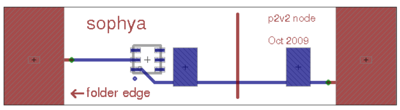

C.2.2 Module circuitry

Schematic diagram of ordered SOPHYA module circuitry.

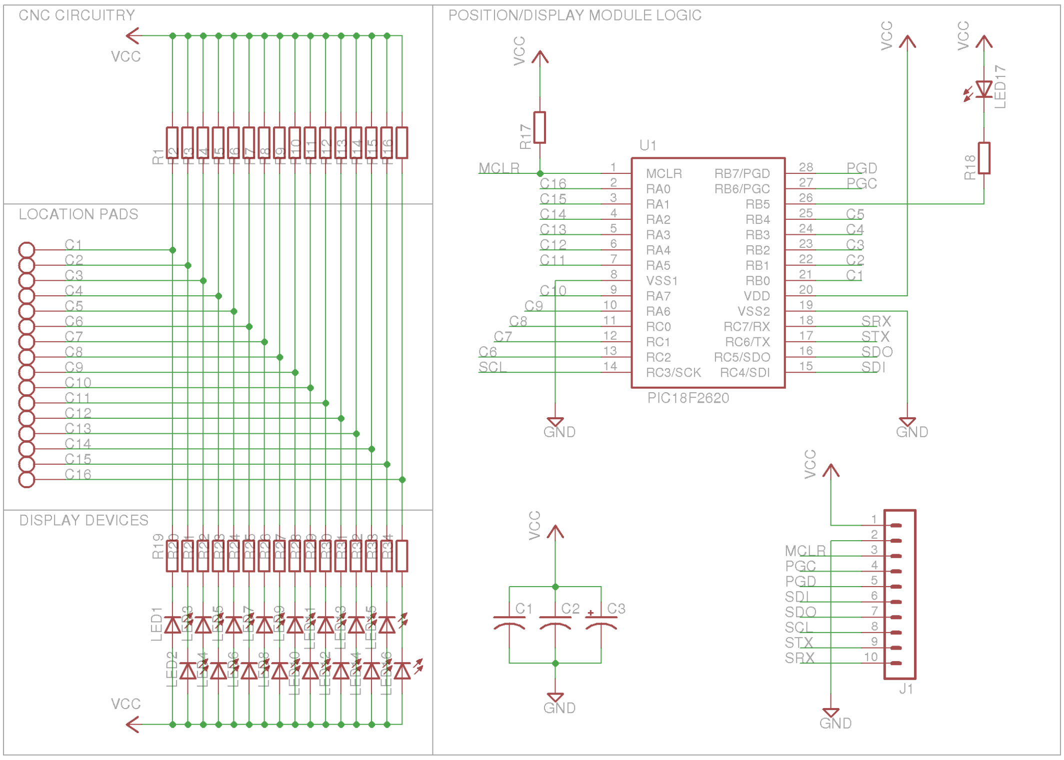

Layout of ordered SOPHYA module PCB (actual width = 96mm).

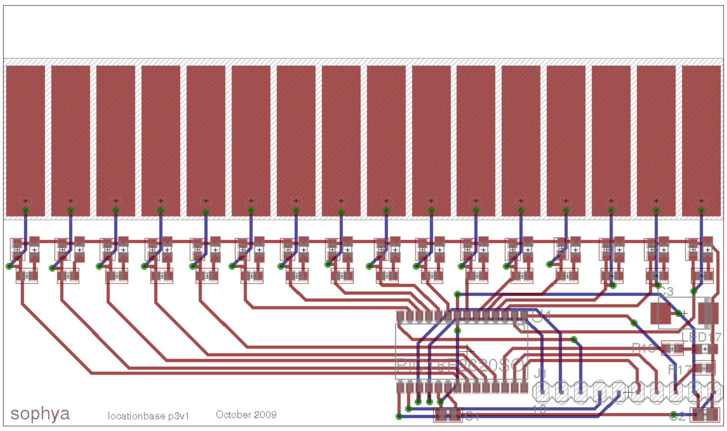

C.2.3 Controller circuitry

The schematic diagram of the controller circuitry is included below. Although the controller logic and serial line driver sections of the schematic are only shown once, these are actually implemented twice: once for the position controller and once for the display controller.

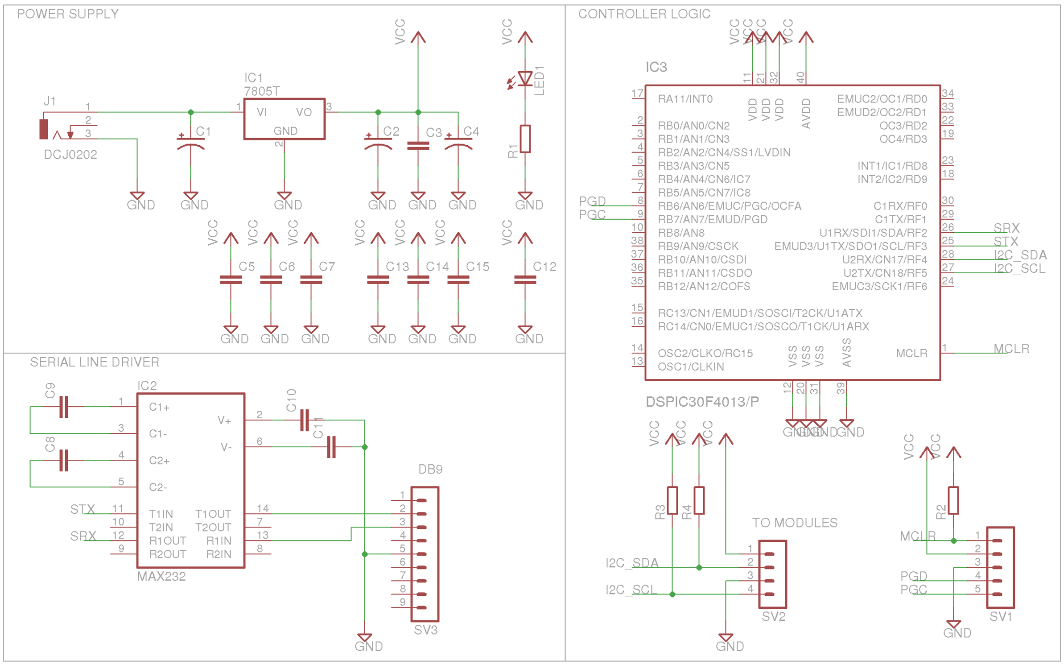

There is no accompanying PCB layout, since a PCB was not manufactured. The controller circuitry was constructed on stripboard, as shown below.

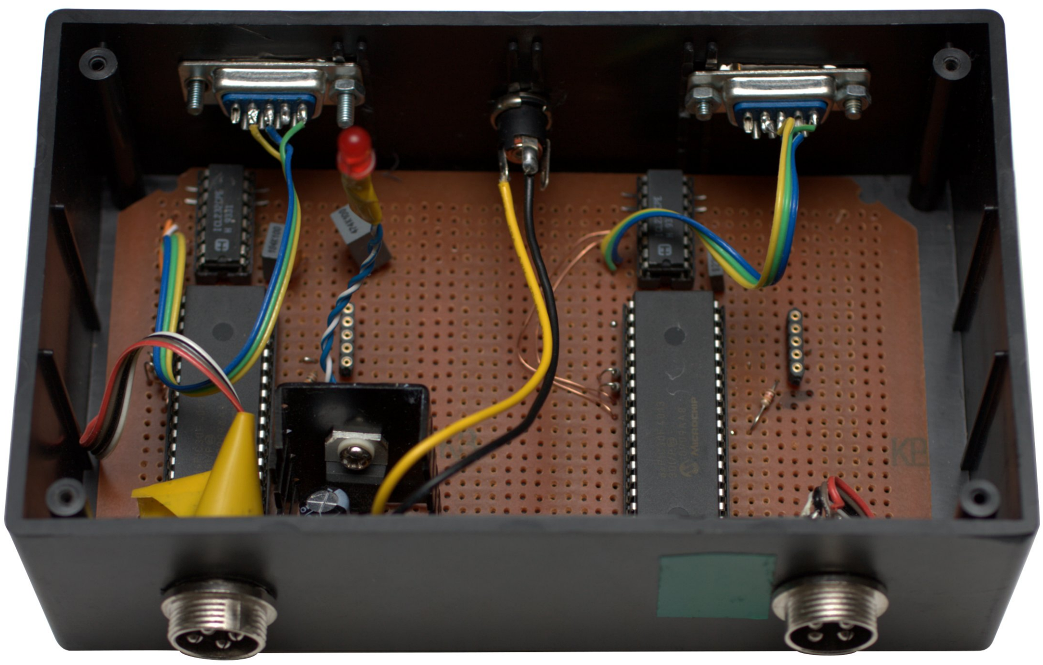

Appendix D

Ethics Approval for Evaluation Study

This appendix includes the letter of approval from the ethics committee of the Faculty of Computing and Mathematical Sciences, University of Waikato for the study described in Chapter 10.

School of Computing &
Mathematical Sciences
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4021
www.scms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

28 October 2009

Matthew Jervis
C/- Department of Computer Science
University of Waikato

Dear Matthew

Request for approval to perform an evaluation involving human participants

I have considered your request for approval to perform an evaluation involving human participants, commencing in October 2009.

We note that this experiment is a partial requirement for your PhD research "Integrating Physical and Digital Workspaces" being conducted to evaluate a system (Sophya) for assisting with paper document management and integrating this with digital document management.

The procedure described in your request is acceptable. I note your statements that you expect to produce a conference or journal paper and that information gathered will be kept securely stored in the SCMS data archive. No identifying information about the participants will be published and any information used in publications will be anonymised. Notes and documents will be destroyed and recordings erased by 31 January 2011.

The research participants' Information Sheet the Research Consent form complies with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to perform the user study.

Yours faithfully

Tony Smith
Department of Computer Science
Human Research Ethics Committee
School of Computing and Mathematical Sciences

Appendix E

Evaluation Study Handbook

The following is the handbook that was given to participants in the evaluation study described in Chapter 10.

Introduction

The purpose of this study is to evaluate SOPHYA, a system designed to integrate paper and electronic document management. In order to evaluate SOPHYA we have simulated an office environment in which it could potentially be used. You will assume the role of a PA at a fictitious landscape design company called Blandscapes Ltd. During the course of this study you will perform several typical filing tasks, which will give you a chance to familiarise yourself with the features of SOPHYA. After each task you will be asked to fill in a questionnaire. The aim of this process is to gather feedback, based on your own experience, as to whether or not such a system would be useful in the office, and your reasons why you feel this way.

Blandscapes

Employees of the fictitious company Blandscapes work with both paper and electronic documents. Documents are stored in folders, organised by job (each job gets its own folder). Job folders exist both as a real world folder and as an electronic virtual folder. When jobs are initially started, and while they remain *in progress*, plastic clamp folders (as in Figure 1(a)), are used. When a job is completed, it is archived and the clamp folder is replaced with a codafile folder (Figure 1(b)).

There are two shelves for storing job folders, one for *in progress* jobs (Figure 1(c)) and one for *completed* (Figure 1(d)) jobs. Folders and shelves have SOPHYA circuitry embedded into them that allows them to communicate with the computer. When a job folder is placed onto or removed from its shelf, this is detected by the SOPHYA circuitry and communicated to the computer. LED indicators are able to show the location of a specific folder when requested by the software. We will cover these features in more detail in the next section.

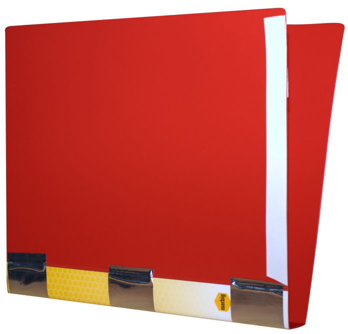

(a) Folder for *in progress* jobs

(b) Folder for *completed* jobs

(c) Shelf for *in progress* job folders

(d) Shelf for *completed* job folders

Figure 1: Filing system used by Blandscapes.

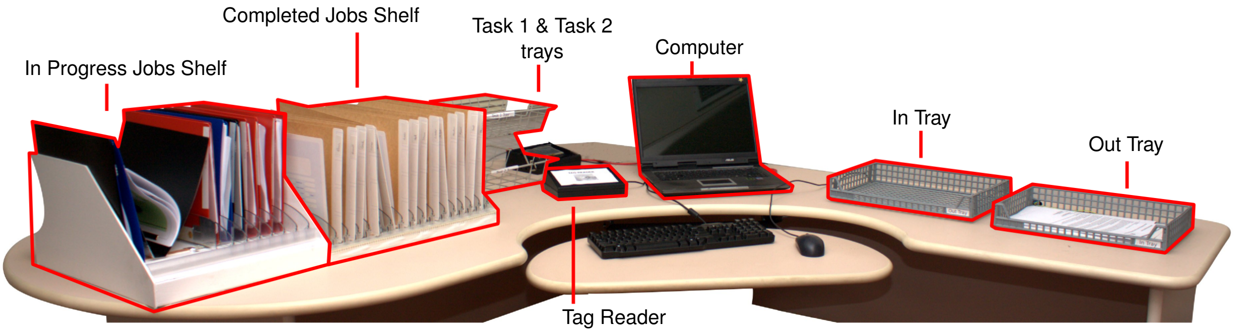

Figure 2: Layout of the desk.

Desk Layout

Figure 2 shows the layout of the desk you will be using during this study. On the left are the shelves where the real-world folders for *in progress* and *completed* jobs are stored. During the study you will work with the folders on these shelves.

You will also need to use the computer and tag reader; the next section provides you with instructions for using these. Also shown are the document trays that you will use during the tasks.

Tutorial

In this section we will run through a practical demonstration of the document management system so that you can familiarise yourself with it. For this tutorial you will be carrying out the following task:

> *You have received in your "in tray" hard copies of the the "Pool and Driveway Maintenance Guide," which need to be filed in the folders of jobs that were started before 8 January 2010 and include a pool and driveway.*

To carry out this task you will need to take the following steps: use the software to search for jobs that match the given criteria, retrieve the appropriate folders, add the documents to the folders, and return the folders to the appropriate places.

Searching

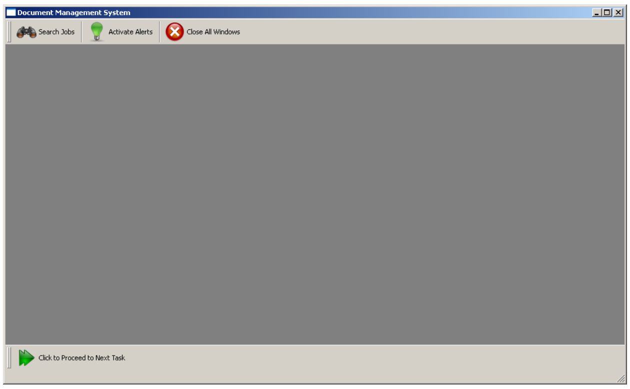

When you are ready to begin click "Start". The image above shows the document management software when it is started. The first step is to use the search function to find jobs that match the given criteria.

- Bring up the search window by clicking the *Search Jobs* button.
- To find all jobs started before 8 January 2010 select *Started before* from the drop down box next to *Date*, and select the date in the date box to its right.
- To find only jobs with a Pool and Driveway, check the boxes next to *Pool* and *Driveway*.

You should now see the following:

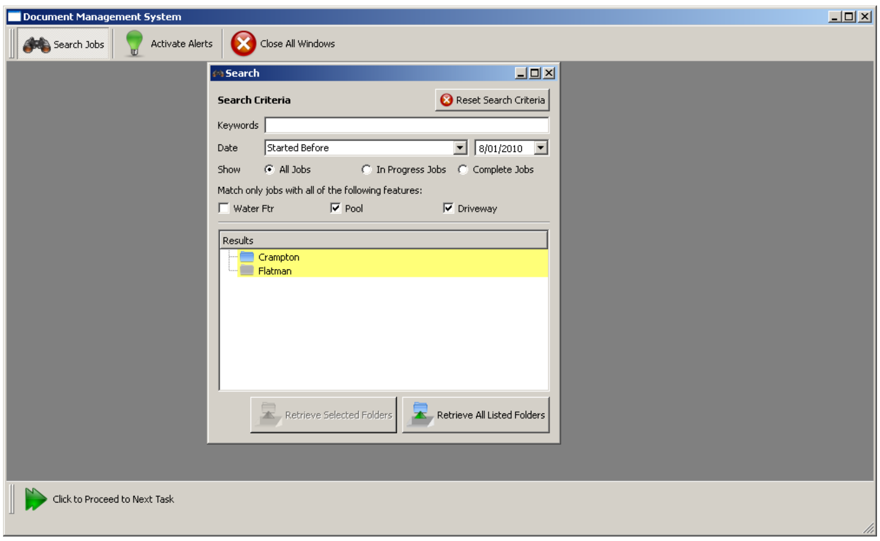

Locating and retrieving folders

Now that the search criteria has been entered, the job folders that need to be retrieved from the shelves are listed. SOPHYA is able to show the location of these folders in the real world by illuminating LEDs on or beneath the folders. The *Retrieve All Listed Folders* button at the bottom of the search window turns on the LEDs of all job folders currently listed. Clicking this button will cause the dialog box depicted below to appear. Once the folders have been retrieved from the shelf the "Done" button can be clicked, this will turn off the LEDs.

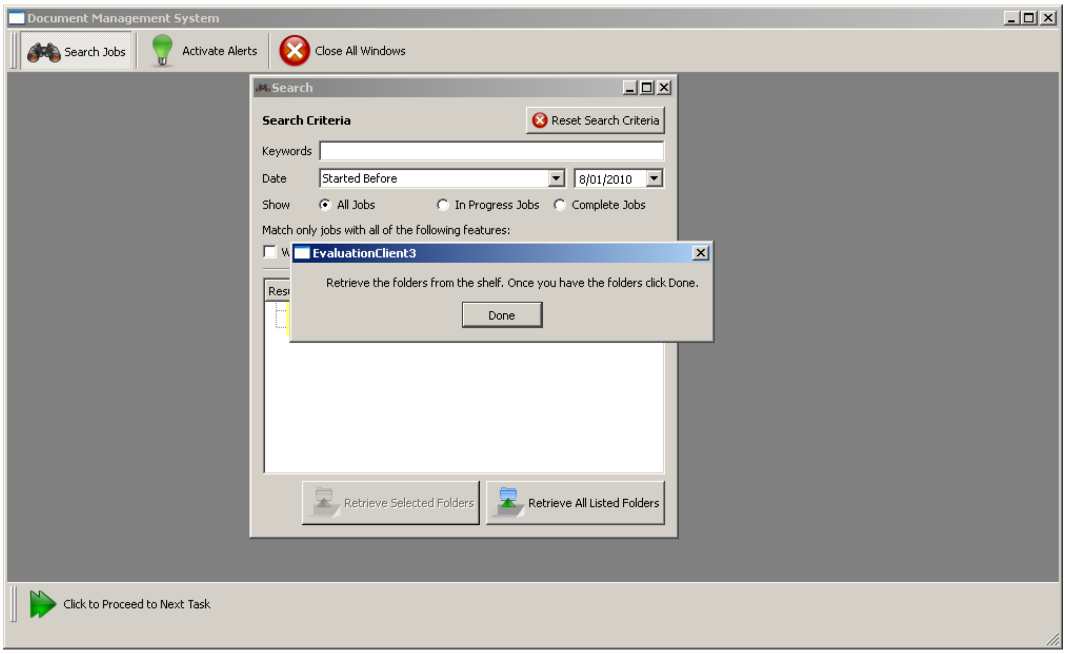

- Click the *Retrieve Folders* button. When the dialog box comes up do not click *Done* yet.
- From the job folder shelves (shown below), retrieve the real world folders whose LEDs have lit up. Note that the folders may be on either or both the *in progress* and *completed* job shelves.

- Click the *Done* button on screen. This will cause the LEDs to go out.
- Place a copy of the "Pool and Driveway Maintenance Guide" (this can be found in the "in tray" on your desk) in each of the folders you retrieved.

Returning folders to their shelves

The process for returning a folder to its appropriate shelf depends on whether the folder is for an *in progress* or *completed* job.

In Progress Jobs

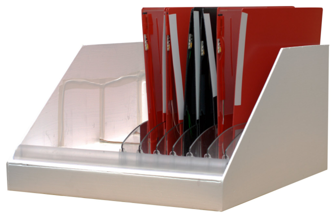

The order of folders on the *in progress* jobs shelf is not important; these folders are constantly moving around the office and the LEDs on the folders are able to quickly display their location as long as they are on the shelf. Therefore, in progress job folders can simply be placed anywhere on the shelf (they must, however, be placed onto the correct shelf, as shown on the right).

- Place the *Crampton* folder back on the *in progress* jobs shelf now.

Completed Jobs

The process for returning folders for *completed* jobs is different. These jobs are accessed less frequently and there is a larger number of folders, therefore they are stored in alphabetical order. SOPHYA is able to assist with returning these folders to the shelf; it does this by illuminating the LED indicators on the shelf at the appropriate location.

Each of the folders has an electronic tag (Figure 3(a)) that can be read by a tag reader connected to the computer (Figure 3(b)). When a folder's tag is read by the tag reader it instructs the software that you wish to return this folder to the cabinet. The software will work out the best location for the folder and light the LEDs there. Once the folder is returned to the shelf LEDs will turn off.

> *Please note:* When returning the folder to the shelf care must be taken to ensure that it sits properly, so that SOPHYA can detect its presence. When placing the folder on the shelf, make sure the back of the folder is hard up against the back of the shelf (the folder may need to be lifted as you slide it back as the contacts tend to catch on the bottom). Please remember that this is an early prototype of the system, and issues like this are to be resolved in future designs.

- Scan the tag of the *Flatman* folder.
- Place the folder back onto the shelf at the position indicated by the LEDs.

(a) Tag (b) Tag Reader

Figure 3: Tag and Tag Reader used to identify documents.

- You will find the folders for clients *Reilly* and *Hobbs* in your "in tray", . Follow the steps above to file them in appropriate places on the shelves.

Moving on

- Click the *Click to Proceed to Next Task* button at the bottom of the screen (circled below).
- Click *Yes* to proceed to the next task.

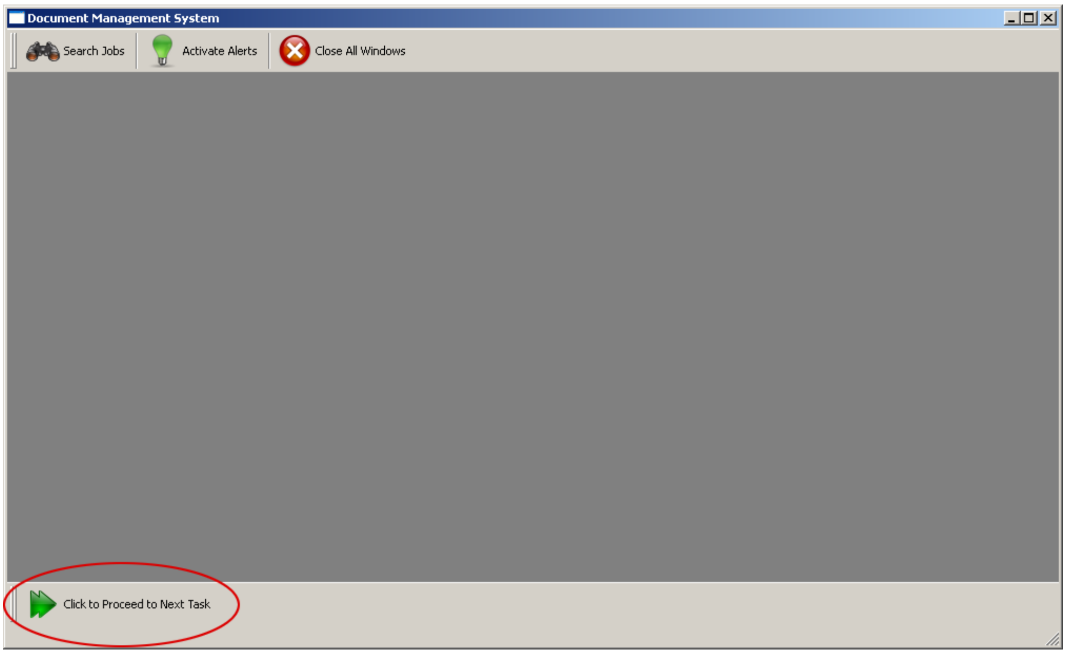

Task 1

The manager of Blandscapes wishes to review some of the jobs that have been in progress since 9 April 2010, and has asked you to retrieve the appropriate folders and place them in a tray to be collected later.

1. Use the search facility of the software to locate the jobs which are still in progress, and which were started after 9 April 2010.
2. Retrieve the appropriate files and place them in the "Task 1" document tray.
3. Click the *Proceed to Next Task* button.
4. This will bring up a questionnaire. Please provide as much detail as you can in each section. Your thoughts and views on SOPHYA are an important part of the information being gathered in this study.
5. Once you have filled in the questionnaire click the *Done* button at the bottom, and proceed to the first task.

Task 2

You have been asked by one of the designers to retrieve the folders for all jobs completed before 20 June, 2009.

1. Use the search facility of the software to locate the jobs which were completed before 20/06/2009.
2. Retrieve the appropriate files and place them in the "Task 2" document tray.
3. Click the *Proceed to next task* button and fill in the questionnaire.

Please ensure that you have filled in the questionnaire for this task before you turn the page.

Task 3

One of the company's sales reps is going to meet with prospective clients. You have been asked to prepare a file of some completed drawings to be shown to clients at the meeting. The sales rep has asked for drawings from all jobs that match *either* of the following criteria:

- *completed* jobs with a pool *and* water feature
- *completed* jobs matching the keyword "Newton"

Repeat the following steps for each of the criteria bullet pointed above:

1. Bring up the *Search Jobs* window.
2. Use the search function to find the job(s) matching the criteria.
3. Use the software to determine the folders' real world positions and retrieve them from the shelf.
4. Retrieve the design from each folder and place it into the meeting folder (you will find this in your "In Tray").
5. Return the folders to the correct location on the shelf (refer back to the tutorial section *Returning folders to their shelves* to remind yourself of the correct procedure for returning folders).
6. Close the *Search Jobs* window.

Once you have completed these steps for each of the given criteria, place the meeting folder in your "out tray" and click the *Proceed to Next Task* button.

Please ensure that you have filled in the questionnaire for this task before you turn the page.

Task 4

During the lifetime of a job, the drawings done by the designers will be viewed by the client, giving the client the opportunity to request changes be made to the drawing. The designer will then make the requested changes. Before the drawing is shown to the client again, you are asked to review the changes that have been made against the list of changes provided by the client.

This type of task involves working with documents both in the real world (the drawings) and on the computer (the change list). The following steps will walk you through the process of performing this task for the clients named **Davie**.

1. Browse or search for the *Davie* job using the search window.
2. Double click on the job name in the list of search results. This will bring up the job information window (shown on the right below). This is the "virtual" representation of the job. It shows information associated with the job that is stored in the database, as well as listing the documents for that job.

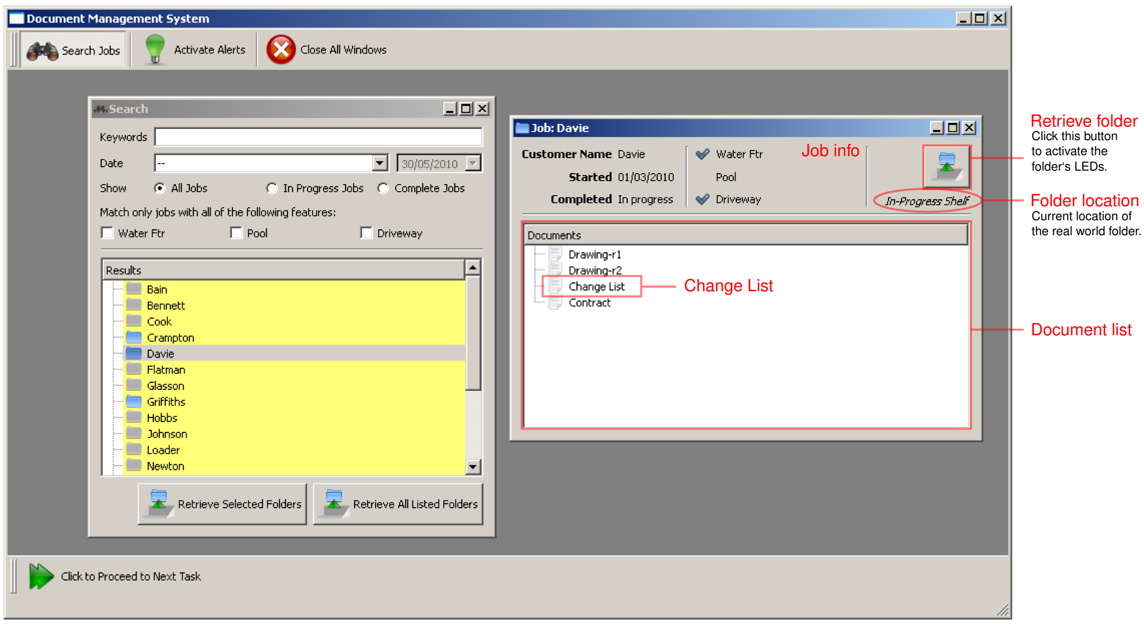

3. The "Change List" is an electronic document, stored on the computer. To view it double click the "Change List" document in the job window's document list (as shown above). This brings up the list of changes requested by the client.
4. The *Retrieve Folder* button (see above) can be used to retrieve the folder. Click it to illuminate the folder's LED and retrieve the folder.
5. In the folder you will find hard copies of the two most recent drawings (you can tell which one is newer by looking at the revision number in the bottom right hand corner of the drawing).
6. Compare the two hard copy versions of the document and check whether all the changes listed in the "Change List" have been applied. You may find the "Guide to Drawing Symbols" in the *in tray* useful.
7. If the changes were not applied correctly place the job folder in your "out tray", otherwise return it to the filing cabinet (if necessary refer back to the tutorial section *Returning folders to their shelves* to remind yourself of the correct procedure for returning folders).
8. Click the *Proceed to Next Task* button and fill in the questionnaire.

Please ensure that you have filled in the questionnaire for this task before you turn the page.

Task 5

The software has the ability to generate alerts at given times or under set conditions. When an alert is triggered the LED for a given folder will light up green (if it is an *in progress* job), or blink (if it is a *completed* job). Potential uses of time based alerts include alerting you when jobs need immediate action (for example getting close to deadline) or indicating jobs that might need attention (such as jobs that have not had any recent activity). Other alerts may be triggered by the software, for example receiving an email from a given client may trigger an alert for their job folder.

Part I

For this task we will assume that an alert has previously been setup to highlight *in progress* jobs that have had no activity within the past seven days. After seven days of a job having no activity the alert will be triggered and the folders green LED will light.

In order to save confusion when performing other tasks, it is possible to turn the alerts off (as they have been in the previous tasks). To enable alerts click the *Activate Alerts* button in the toolbar at the top of the main window (shown below).

Now retrieve the folder(s) whose green LEDs are lit (ie those that have had no activity in the past 7 days) and place them in your "out tray", as you will give them to the manager to review.

Part II

Another alert was previously setup to be triggered when a given document was returned to you; this was to remind you to file the document back into its original folder. The alert is activated when the document's tag is scanned by the tag reader. When this happens it will cause the LED of the appropriate folder to turn on.

You will find the document, (shown below) in your "in tray". Scan its tag (on the back of the document) and return it to the appropriate folder.

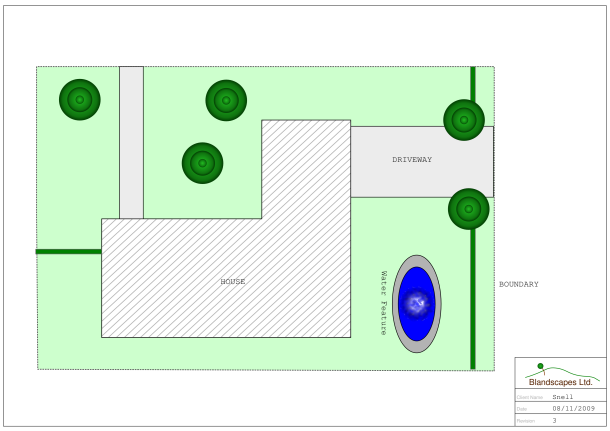

Please ensure that you have filled in the questionnaire for this task before you turn the page.

Appendix F

Video Demonstration

A video demonstration of the two SOPHYA prototypes described in Chapters 7 and 8 has been included with this thesis.